THESE DE DOCTORAT

NNT : 2023UPASG072

université PARIS-SACLAY

Concurrent two-player antagonistic games on graphs

Jeux à deux joueurs antagonistes concurrents sur les graphes

Thèse de doctorat de l'université Paris-Saclay

École doctorale n° 580, sciences et technologies de l'information et de la communication (STIC)
Spécialité de doctorat: Informatique
Graduate School : Informatique et sciences du numérique. Référent : ENS Paris-Saclay

Thèse préparée dans l'unité de recherche **Université Paris-Saclay, CNRS, ENS Paris-Saclay**, sous la direction de **Patricia BOUYER**, Directrice de recherche, la co-direction de **Stéphane LE ROUX**, Maître de conférence

Thèse soutenue à Paris-Saclay, le 12 octobre 2023, par

Benjamin BORDAIS

Composition du jury

Membres du jury avec voix délibérative

Serge Haddad Professeur, ENS Paris-Saclay	Président
Richard Mayr Reader, University of Edinburgh	Rapporteur & Examinateur
Benjamin Monmege Maître de conférences, Aix-Marseille Université, LIS	Rapporteur & Examinateur
Krishnendu Chatterjee Professor, IST Austria	Examinateur
Serge Haddad Professeur, ENS Paris-Saclay	Examinateur
Marta Kwiatkowska Professor, University of Oxford	Examinatrice
Sophie Pinchinat Professeur, Université de Rennes 1, IRISA	Examinatrice
Bruno Ziliotto Chargé de recherche, CNRS, CEREMADE	Examinateur

universite PARIS-SACLAY
ÉCOLE DOCTORALE
Sciences et technologies de l'information et de la communication (STIC)

Titre: Jeux concurrents à deux joueurs antagonistes sur les graphes

Mots clés: Jeu concurrent stochastique, forme de jeu, transfert local-global, stratégie optimale, Détermination des jeux de Blackwell, Jeu de parité

Résumé: On étudie des jeux à deux joueuses (A et B) sur des graphes. À partir d'un état du graphe, les joueuses interagissent pour aller d'un état à un autre. Ceci induit une suite infinie d'états à laquelle une fonction de gain mesurable associe une valeur dans $[0, 1]$. La Joueuse A (resp. B) tente de maximiser (resp. minimiser) l'espérance de cette fonction de gain.

Les jeux à tours, i.e. les jeux tels qu'à chaque état une seule joueuse choisit (une loi de probabilités sur) l'état suivant, ont de nombreuses bonnes propriétés. Par exemple, dans tous les jeux à tours perd/gagne déterministes, une joueuse a une stratégie gagnante. De plus, dans les jeux de parité à tours finis, les deux joueuses ont des stratégies optimales positionnelles. A contrario, les jeux concurrents, i.e. les jeux tels qu'à chaque état les deux joueuses concourent au choix d'une loi de probabilité sur les états suivants, se comportent mal. Ainsi, il existe des jeux concurrents de parité déterministes tels que : aucune des joueuses n'a de stratégie gagnante ; aucune joueuse n'a de stratégie optimale, même stochastique. De plus, lorsque c'est possible, jouer de manière optimale peut nécessiter une mémoire infinie.

Le but de ce manuscrit est d'enrichir notre compréhension du comportement des jeux concurrents. Pour ce faire, on étudie la notion de forme de jeu. Les formes de jeu sont les objets mathématiques qui décrivent les interactions (locales) des joueuses à chaque état d'un jeu concurrent. Les formes de jeu sont définies par un ensemble de stratégies locales par joueuse, un ensemble d'issues et une fonction envoyant une paire d'une stratégie locale par joueuse sur une loi de probabilités sur les issues. Généralement, dans les articles sur les jeux concurrents, les interactions locales sont des formes de jeu standard (finies) : les ensembles de stratégies locales sont des lois de probabilités sur les ensembles (finis) d'actions sous-jacents. Ici, on définit des formes de jeu plus générales, que l'on appelle formes de jeu arbitraires. Certains des résultats établis dans ce manuscrit supposent que les interactions locales sont standard, tandis que les autres ne font pas de telles hypothèses.

Premièrement, on prouve des résultats généraux sur les jeux concurrents, avec très peu d'hypothèses sur les fonctions de gain et les interactions locales. En particulier, on considère un résultat crucial sur les jeux concurrents : la détermination de Blackwell de Martin, qui peut être énoncé comme suit. Soit un jeu concurrent dont toutes les interactions locales sont standards finies. Depuis chaque état, il existe une valeur u dans $[0, 1]$ telle que les stratégies de la Joueuse A (resp. B) peuvent garantir que l'espérance de la fonction de gain est au moins (resp. au plus) égal à n'importe quel seuil en-dessous (resp. au-dessus) de u. On généralise ce résultat aux jeux dont les formes de jeu sont arbitraires et en déduisons d'autres résultats sur les jeux concurrents. On prouve également d'autres résultats sur les jeux concurrents, en particulier sur les stratégies optimales en sous-jeu.

Deuxièmement, on étudie le comportement des jeux de parité concurrents finis en termes d'existence et de nature des stratégies (presque) optimales (en sous-jeu), avec très peu d'hypothèses sur les interactions locales.

Troisièmement, on définit des ensembles de jeux concurrents qui ont certaines des propriétés des jeux à tours tout en étant plus généraux que les jeux à tours. Ainsi, étant donnée une propriété souhaitable sur les jeux concurrents, on caractérise tout d'abord les formes de jeu qui garantissent que tous les jeux simples qui les utilisent comme interactions locales satisfont cette propriété. On caractérise ainsi les formes de jeu qui se comportent bien individuellement. On montre ensuite que tous les jeux concurrents qui utilisent ces formes de jeu comme interactions locales satisfont également cette propriété. Ces formes de jeux se comportent également bien collectivement.

université PARIS-SACLAY
ÉCOLE DOCTORALE
Sciences et technologies de l'information et de la communication (STIC)

Title: Concurrent two-player antagonistic games on graphs

Keywords: Concurrent stochastic game, game form, local-global transfer, optimal strategy, Blackwell determinacy, Parity game

Abstract: We study games played by two players, Player A and Player B, on a graph. Starting from a state of the graph, the players interact to move from state to state. This induces an infinite sequence of states, which is mapped to a value in $[0, 1]$ by a measurable payoff function. Player A (resp. B) tries to maximize (resp. minimize) the expected value of this payoff function.

Turn-based games, i.e. games where at each state only one player chooses a (probability distribution over) successor state, enjoy many nice properties. For instance, in all deterministic win/lose turn-based games, from each state, one of the players has a winning strategy. In addition, in finite turn-based parity games, both players have positional optimal strategies from each state. By contrast, concurrent games, i.e. games where at each state both players interact concurrently, i.e. simultaneously, to generate a probability distribution over successor states, behave much more poorly. Indeed, there are very simple deterministic concurrent parity games such that: neither player has a winning strategy; neither player has an optimal strategy, even a stochastic one. In addition, when optimal strategies do exist, they may require infinite memory.

The goal of this dissertation is to give significant insight on how concurrent games behave. To do so, we study the notion of game form. Game forms are the mathematical objects that describe the (local) interactions of the players at each state of a concurrent game. Game forms are defined by a set of local strategies per player, a set of outcomes and a function mapping a pair of one local strategy per player to a probability distribution over outcomes. Generally, in the literature on concurrent games, local interactions are standard (finite) game forms: the sets of local strategies are distributions over underlying (finite) sets of actions. In this dissertation, we define and study more general game forms, which we call arbitrary game forms. Some of the results we prove hold even with arbitrary local interactions, the others use a standard assumption on the local interactions involved.

First, we prove general results on concurrent games, with very few assumptions on the payoff functions and local interactions involved. In particular, we consider a crucial result on concurrent games: Martin's result on Blackwell determinacy, which can be stated as follows. Consider a concurrent game where all local interactions are standard finite. From each state, there is a value u in $[0, 1]$ such that Player A's (resp. B's) strategies can guarantee that the expected value of the measurable payoff function is above (resp. below) any threshold below (resp. above) u. We generalize this result to games with arbitrary game forms. We deduce from this generalization other results on concurrent games, possibly using standard local interactions, which could not have been obtained directly from the original result by Martin. We also prove other results on concurrent games, in particular results related to subgame optimal strategies.

Second, we study how finite-state concurrent parity games behave in terms of existence and nature of (almost and/or subgame) optimal strategies, with very few assumptions on the local interactions involved.

Third, we define subsets of concurrent games that enjoy some of the nice properties of turn-based games while being more general than turn-based games. These subsets are constructed via local-global transfers, which is a novel approach. Specifically, given a desirable property on concurrent games, we first characterize the game forms that ensure that all simple games using them as local interactions satisfy this property. Thus, we characterize the game forms that behave well individually. We then show that all concurrent games that use these game forms as local interactions also satisfy this property. Thus, we show that these game forms also behave well collectively, hence globally.

Acknowledgments

I am deeply unsettled as to how to write acknowledgments, as it seems unavoidable to forget to thank people without whom I could not have completed my PhD. That being said, I will still try my best, I am sorry for all those that I do not mention.

This goes without saying — but it is important to say so — I was very lucky to have you, Patricia and Stéphane, as my PhD advisors: I could not have accomplished anything without the both of you. I have always looked forward to our weekly meetings as our discussions have constantly been very stimulating and encouraging for me. I would like to thank the both of you for all your time and dedication.

I would like to thank all the jury members of my PhD defense, and especially the two reviewers. I know that my PhD dissertation was long, and probably took way too much of your time to read, or even skim through. I would also like to thank the people in charge of handling the technical aspect of the defense (such as the zoom session), who allowed me to not have to worry about these issues.

I would like to thank all the internship supervisors I had before starting my PhD. I owe you both the opportunity to do this PhD and all the experience I gathered prior to my PhD, which made possible all the research I did afterwards. I would also like to thank my fellow students and colleagues, who brightened my days at ENS, both Rennes and Paris-Saclay.

Finally, I switch to French for the conclusion:

Je voudrais remercier ma famille pour m'avoir bien orienté et soutenu durant toutes ces longues années d'étude jusqu'à avoir été là — en présentiel ou en distanciel — pendant ma soutenance de thèse. Je termine évidemment par remercier ma Clémentine, qui m'a merveilleusement accompagné pendant ces trois années de thèse.

Contents

Introduction

In this introduction, we give the general context motivating the framework used and the questions tackled in this dissertation.

After we discuss game theory, the path that we take to introduce model checking is inspired by what is done in [1, 2]. We then give the purpose of this dissertation and an outline of how this dissertation is structured. We conclude by several remarks regarding how to read this dissertation.

Game Theory

Game theory is a very broad interdisciplinary field that studies interactions between multiple players in a competitive or cooperative setting. Each of these players tries to satisfy an individual or common objective. One of the founding results in game theory dates back to 1928 when von Neumann proved his minimax theorem [3]. It also appeared in [4] in, arguably, one of the seminal books on game theory. This result holds in one-shot two-player zero-sum games, also known as "matrix games". The study of non-antagonistic matrix games, where many players have individual objectives which are independent from one another, essentially started in the early 1950s with Nash's articles [5, 6]. The notion of infinite-duration games, where the outcome is induced by the players' interaction lasting infinitely many rounds, emerged shortly after. For instance, in 1953, the existence of winning strategies in deterministic turn-based games with open or closed objectives was proved in [7]. The proof that this actually holds for all Borel objectives (i.e. Borel determinacy) came much later in [8] by Martin, in 1975. Stochastic games, in which stochastic transitions occur at each round of the game, appeared with Shapley [9] in 1953 and with Everett [10] in 1957. Then, the notion of imperfect information games (which are referred to as concurrent games in this dissertation) was introduced by Blackwell in 1969 in [11]. The determinacy of Blackwell games was established by Martin much later, in 1998 [12].

Game theory is applied in various domains from biology to economics or logic and, more importantly for this dissertation, computer science. The link between game theory and biology is quite old since evolutionary game theory essentially started in 1973 with [13], where the ways animal species behave are modeled in a game theoretic setting. The links between game theory and logic is even older as it dates back to the 1950s where game semantics started in different areas of logic. See for instance [14] for examples of Ehrenfeucht–Fraïssé games, which are games that are used to prove bisimulation.

The relationship between game theory and automata theory essentially started with [15, 16, 17] where decidability issues were tackled. A few years later, game theory was used at the intersection of automata theory and model checking [18] (and therefore computer science). In the remainder of this introduction, and in this dissertation as a whole, we will mostly comprehend game theory from a model checking perspective. This should not impede the use of the results established in this dissertation in other areas.

Context

Nowadays, a very large variety of aspects of our daily lives are controlled by computer systems, which grow rapidly in number and complexity. Our dependence towards these systems often entails handling critical tasks, whether it is managing banking accounts, controlling power plants or coordinating railway systems. For such critical tasks, it is essential to be able to either guarantee the safety of already existing systems or build new ones that are safe by design, thus ensuring that no failure will occur. However, such tasks are inherently complex due to the multitude of possible scenarios in which these systems need to behave properly. In addition, establishing guarantees on computer systems is made all the more difficult by the fact that these systems are often reactive [19], in the sense that they interact with an external environment. This environment encompasses both interactions with human agents, and also unpredictable events. Therefore, it is crucial to be able to synthesize system controllers whose purpose is to keep reactive systems safe against all environments, even possibly hostile ones.

Formal method and model checking. Merely testing a reactive system against many possible environments is not enough to give strong guarantees on its behavior. Indeed, testing may only exhibit failures, it cannot prove their absence against possibly infinitely many environments. The purpose of formal methods is to give formal, mathematical guarantees on how a system behaves against all possible environments. The system is described by a simplified, abstract model. How the system should behave is encoded as a specification on the model. The most natural question we may be asking given a model and a specification is whether the model satisfies the specification. That is, whether the system behaves how it is supposed to. This corresponds to verification, and it is the historical focus of model checking [20, 21]. However, there are other, harder questions we may consider on models and specifications.

Rather than considering verification, one may be interested in controller synthesis. That is, in situations where a controller can impact the real reactive system, the model describing this system is left underspecified. Depending on the state of the system, several actions are available for the controller to choose; and to be complete, the model needs a controller's policy (or strategy), that is, a way to choose actions depending on the state of the system. Controller synthesis then amounts to inquire if there is a controller's policy that makes the model satisfy the specification; and if so, synthesize such a policy.

Finally, one may want to perform model design, which is a (very) conjectural approach. It consists, from a specification describing in a simplified (abstract) way how a real system should behave, in designing a model, or a class of models, that satisfies the specification; or in which there exists a controller's policy that makes the model satisfy the specification. We can then build real systems corresponding to this abstract model, and possibly implement appropriate controller's policies. That way, the system that we obtain behaves as desired by design.

Game theory for model checking. As mentioned above, the use of game theory for model checking is not new, see for instance [22] in which synthesis is seen as a game. In this dissertation, we will be interested in two-player antagonistic games on graphs. In such a setting, one player represents the controller. She has at her disposal the actions available to her in the reactive system, depending on the state of the system. The other player represents the environment. Indeed, the possible ways that the environment can influence the system are known, but the environment is completely unpredictable. That is, the way it will impact the system cannot be described with a known probability distribution. Since the environment may potentially be influenced or controlled by a malevolent agent, the game is antagonistic, or zero-sum, in the sense that any outcome of the game is as positive for a player as it is negative for the other player. The graph on which the game is played then represents the different states in which the reactive system can be. How the controller and the environment impact the system differs depending on the current state of the system. In addition, going from states to states in the system may be described via stochastic transitions: though the process may be random, the underlying probability distributions according to which the system evolves is known a priori. To subsume both terminating and non-terminating behaviors of the system, the games we consider are infinite-duration, i.e. they last indefinitely, as long as no stopping state is reached, in which case the game stops immediately. Furthermore, the specification of the model, which describes how the system should behave, is encoded as the objective in the game; the controller player tries to ensure that the specification holds, while the environment player tries to ensure that is does not.

When no transition between states of the graph is stochastic, the game is said to be deterministic. In that case, the controller seeks a winning strategy, i.e. a way to choose actions (i.e. a strategy) ensuring that the objective holds in all cases, regardless of what the environment player does, and reciprocally for the environment player. However, there are not always winning strategies, for either of the player. In such cases, the controller player seeks an optimal strategy, i.e. a strategy that maximizes the probability that the objective holds, against all environment player's strategies. This notion of optimal strategies is also made all the more relevant when we consider richer specifications that are not encoded by an objective but rather by a payoff function. In that case, the controller player tries to maximize its expected value, while the environment player tries to minimize it.

Finally, for synthesizing purposes, the simpler winning or optimal strategies are to describe, the better. Hence, when designing a model, a criterion that could be used to assess a model's worth could be the existence of simple-to-describe winning or optimal strategies.

Game Formalism

Let us describe the games that we consider in this dissertation a little more formally. Before describing games, we describe arenas, that can be seen as games where the objective or the payoff function is not yet specified. In an arena, two players, that we will always call Player A and Player B, interact on a graph. This graph consists of sets of states and transitions between these states. Furthermore, it is equipped with an arbitrary set of colors. Each state of the graph is given a color and a set of local strategies[1] available to the players at this state. Each time the play reaches this state, the players can choose among these available local strategies and, as a result of the players' choices, a probability distribution over successor states is induced. The process then proceeds (stochastically) to another state, and this is repeated indefinitely. This infinite repetition thus generates an infinite sequence of states, which naturally induces an infinite sequence of colors, which is called a trace. We can then obtain a game from this arena by defining a payoff function mapping every trace, i.e. every infinite sequence of colors, to a value in $[0, 1]$. In the game that we obtain, Player A tries to maximize the expected value of this payoff function, whereas Player B tries to minimize it. The game is said to be win/lose if any trace is mapped by the payoff function to either 1 or 0. In that case, any trace is either winning for Player A and losing for Player B (when it is mapped to 1) or the other way around (when it is mapped to 0).

Among the states of the graph, there may be some stopping states. When they are reached, they immediately stop the play and output a value. Similarly to the expected value of the payoff function, Player A tries to maximize this output value and Player B tries to minimize. For simplicity, in the remainder of this introduction, we assume that there is no stopping state.

Playing in graph arenas is done via graph strategies: they prescribe to the players what to do at each state, depending on the history of the game, i.e. on the finite sequence of states already visited. In other words, graph strategies

[1] In fact, in this dissertation, these local strategies will be called GF-strategy. It will be made clear why in the next section of this introduction.

map finite sequences of states to local strategies. In the following, we will use the word "strategy" to refer to graph strategies. They are not to be confused with local strategies, which are the ones played at each state of the graph. In win/lose games, the players seek winning strategies. A strategy is winning for a player if all traces compatible with this strategy are winning for her. However, in the games that we will study in this dissertation, it will often be the case that neither player has a winning strategy. Furthermore, the notion of winning strategy is not applicable in games with richer payoff functions than win/lose. In such cases, as mentioned above, the players turn to optimal strategies, which are strategies maximizing, or minimizing depending on the player considered, the expected value of the payoff function.

As for the notion of simplicity of strategies hinted above (in page 14), the simplest kind of strategies are positional strategies. What positional strategies prescribe only depends on the current state of the game, not on the whole history of visited states. That is, they play only one local strategy per state of the game. They are therefore much easier to describe than arbitrary strategies. This is in sharp contrast with what we call infinite-choice strategies, which are strategies that may play infinitely many different local strategies at some states of the game. The latter are much harder to describe than positional strategies. In particular, all infinite-choice strategies are infinite-memory. The notion of infinite-choice strategies is novel, and will be formally defined in this dissertation.

Objectives studied. As mentioned above, when a payoff function maps all traces to either 0 or 1, the game is win/lose. In such cases, the payoff function is entirely defined by the set of traces mapped to 1. This constitutes the winning objective (for Player A). In this dissertation, we will often focus on win/lose games. We will mostly consider prefix-independent objectives, i.e. objectives disregarding all finite prefixes in infinite traces. Furthermore, among prefix-independent win/lose games, when considering an explicit objective, it will always be a parity objective, which is a special kind of prefix-independent objective. Let us describe the set of traces winning for Player A with a parity objective. In this case, the set of colors considered is a finite set of integers. Then, given a trace, which is an infinite sequence of colors (i.e. of integers), we consider the maximum of the colors seen infinitely often in that sequence. This maximum exists since there are only finitely many colors. The identity of the player for whom this trace is winning depends on the parity of this maximal color, hence the parity terminology. The set of traces winning for Player A are exactly the ones for which this maximum is even.

The benefit of parity objectives lies in the expressiveness of these objectives as well as their relative simplicity. Indeed, parity objectives are well-suited to express ω-regular expressions sets [23]. These ω-regular sets are very useful when defining specifications on (reactive) systems, see [24].

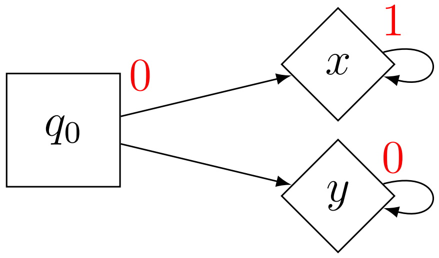

Figure 1: A turn-based game.

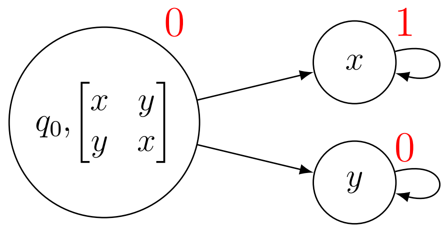

Figure 2: A concurrent game.

Concurrent Vs Turn-based games. Crucially, when describing how the players interact in the arenas, we have not detailed how the local strategies that the players use impact the induced probability distribution over successor states. On the one hand, if at each state of the arena, only one player is really playing while the other has no impact on the decision, the game is turn-based. On the other hand, if there are some states where both players may have an impact, then the game is truly concurrent, i.e. no longer turn-based. The difference is exemplified in the two arenas depicted in Figures 1 and 2. The arena depicted in Figure 1 is turn-based. At the state q_0, Player A plays alone and may choose to go to either x or y. The fact that it is Player A that plays alone at q_0 can be spotted by the fact that the state q_0 is squared-shaped. In fact, she can also choose any probability distribution between the two transitions leading to x or y. The arena depicted in Figure 2 is concurrent. Indeed, at state q_0, Player A and Player B interact. Player A chooses among the (probability distributions over the) rows of the bi-dimensional table, while Player B chooses among the (probability distributions over the) columns. As the result of their concurrent choices, a next state, either x or y, is reached.

Let us make two additional remarks on the arenas depicted in Figures 1 and 2. First, the red numbers appearing near each state of the arenas correspond to their colors. Second, both of these arenas are deterministic because there is no stochastic transition between states unless the players decide to play stochastically. In general, we will assume that there are intrinsically stochastic transitions between the states.

Concurrent games subsume turn-based games as turn-based interactions can be seen as special cases of concurrent interactions. That is, concurrent games have more expressive power than their turn-based counterpart. However, turn-based games have been widely studied in the literature, especially compared to concurrent games. Take for instance the book [25] dating back to 2011 that gives an overview of game theory (on graphs). A large part of this book is dedicated to the study of turn-based games. This is also the case of the very recent book preprint [26], where most models are turn-based. The fact that turn-based games are much more studied than concurrent games can be explained as follows: turn-based games have many nice properties that even simple concurrent games fail to have. We present two striking examples below.

First example: Borel determinacy [8] ensures that, in all deterministic win/lose turn-based games with Borel objectives, from any state, either of the players has a winning strategy, i.e. can ensure winning regardless of what the other player does. This does not hold in concurrent games in general. In fact, it does not even hold on a concurrent game obtained from the concurrent arena of Figure 2. Indeed, consider a win/lose game obtained from this arena such that Player A wins if and only if the color 1 is seen, i.e. the state x is reached. Then, no player has a winning strategy. Indeed, regardless of what row Player A chooses, Player B may choose a column that leads to y. Symmetrically, regardless of what column Player B chooses, Player A may choose the row that leads to x.

Second example: the difference in behavior between turn-based and concurrent games is also abundantly clear when considering parity games and optimal strategies. Consider parity games with finitely many states and possibly stochastic transitions. In that case, whether the game is turn-based or not, the players do not a priori have a winning strategy. However, it does hold that if the game is turn-based, from every state, both players have an optimal strategy that is positional [27, 28]. This means that the players have strategies maximizing their probability to win. In concurrent games, however, optimal strategies may not exist even in very simple games, as will be shown in Figure 3.1. Furthermore, when optimal strategies do exist, they may require infinite choice, and therefore infinite memory. In other words, this means that the strategies considered need to play infinitely many different local strategies at some states of the game.

However, the inherent intricate behavior of concurrent games should not deter us from studying them. Indeed, real life systems, in which synchronicity is involved, are best described with the help of concurrency [29, 30]. Although turn-based games are more studied than concurrent games, that is not to say that concurrent games have not been studied. Along this dissertation, we will cite several papers dealing with concurrent games either to use a result as is, or to generalize it. Below for the record, we would like to cite additional papers for their algorithmic and/or complexity contributions related to concurrent games, which are issues we do not tackle at all in this dissertation. For instance, in [31], the authors provide algorithms to compute, in concurrent reachability games, the set of states from which Player A can win surely, almost-surely, and limit-surely in reachability games. In [32] it is shown that the values, i.e. the probability with which Player A's strategies can win, of concurrent parity games can be computed with quantitative μ-calculus. In the same setting, the author of [33] show that the problem of computing the value is in TFNP[NP]. On a more practical note, see [34] for a (very recent) implementation of model

$$\begin{bmatrix} x & y \\ y & x \end{bmatrix}$$

Figure 3: The local interaction at state q_0 of the concurrent arena of Figure 2.

checking algorithms on concurrent games.

Purpose of this dissertation

In this dissertation, we study two-player antagonistic concurrent games and our goal is to provide insight on how these concurrent games behave. As mentioned above, concurrent games behave poorly, especially compared to their turn-based counterpart. What differentiates turn-based games from concurrent games is the type of local interactions of the players at each state of the games. Formally, a local interaction is a game form, that is a set of actions available to both players, a set of outcomes, and a function mapping each pair of one action per player to a probability distribution over the outcomes. The local strategies available to the players are then the probability distributions over their available actions. Game forms are usually represented by bi-dimensional tables where Player A chooses the rows, while Player B chooses the columns. For example, the game form in Figure 3 describes the interaction of the players at the state q_0 of the concurrent arena of Figure 2.

The notion of game form that we have informally defined above was first introduced in [35] in the context of social choice theory. However, the first result established on game forms was proved in [36].

This dissertation has three distinctive traits. However, note that they do not correspond to the three parts of this dissertation. The first distinctive trait of this dissertation is that we consider game forms, and by extension local interactions, as first-class citizens. This mathematical object will be studied in two different contexts. First, inside concurrent arenas and games where it appears at each state. Second, for itself outside of a concurrent game context. In this dissertation, the interaction of the players at each state of a concurrent arena will be described by game forms. That is, a concurrent arena is a graph where each state is endowed a game form describing the interaction of the players at that state. Then, the local interaction at any given state refers to the game form that this state is endowed with. This is to be compared with what is usually done in the literature on concurrent games. Indeed, usually a concurrent arena is a graph where, at each state, both players have a set of available actions. There is in addition a function mapping each pair of actions for both players to a probability distribution over successors states. With such a definition, the notion of game form is only implicit, and is often not formally defined (since this notion is not useful to establish the results of these papers). In this dissertation, what we have called until now local strategies, which is what (graph) strategies prescribe at each state, are referred to as GF-strategies, where GF stands for game form.

The second distinctive trait of this dissertation is the game forms that we consider. Indeed, the game forms that we have described above correspond to what we call standard game forms in this dissertation. In fact, we consider arbitrary game forms, which subsume standard game forms. Indeed, an arbitrary game form can be defined as follows. It is a set of GF-strategies for both players, a set of outcomes and a function mapping each pair of one GF-strategy per player to a probability distribution over the outcomes. The crucial difference with standard game forms is that the set of GF-strategies need not be defined as the set of probability distributions over an underlying set of available actions. In particular, this set need not be convex. When all local interactions of a game are standard, the game itself is said to be standard, otherwise the game is said to be arbitrary. Standard games are the concurrent games studied in the literature.

We would like to mention a notable exception to this last statement. In his seminal paper on concurrent (reachability) games [10], Everett implicitly used non-standard game forms. The only assumption made in that paper is that the game forms he uses "possess a minimax-solution". These game forms that "possess a minimax-solution" correspond exactly to the valuable game forms defined in this dissertation. Note however that the notion of game form is not formally introduced in [10].

The purpose of this dissertation is to provide a way to design some well-behaved concurrent arenas, i.e. concurrent arenas enjoying some of the nice properties that turn-based games enjoy, while being significantly more general than turn-based games. Along the way, we will prove several results of various kinds, including a generalization of Blackwell determinacy. This will be discussed in the next section.

Restricting the set of game forms used in concurrent games. Our idea to define such well-behaved concurrent arenas is roughly as follows. Consider a desirable property φ that we want to hold on concurrent arenas. For instance, the property φ could be:

- the existence of winning strategies in all obtained win/lose games (with Borel objectives);

- the existence of optimal (positional) strategies in obtained reachability games.

If φ does not hold in turn-based arenas, there is no chance that it will hold on classes of concurrent arenas that include turn-based arenas. Hence, assume that φ holds in turn-based arenas. This is the case of the two examples cited above. Then, as discussed earlier, without any assumption on the local interactions involved in the arena, this property φ may not hold. However, if we assume that all local interactions are turn-based — and in that case, the arena, and the games obtained from it, are therefore turn-based — then this property φ holds, by assumption. Therefore, we know that there exists a restriction on game forms that make the property φ hold in all games whose local interactions all satisfy this restriction. One of the main goal of this dissertation, which is also the third distinctive trait of this dissertation, is to establish local-global transfers, i.e. to define restrictions on game forms ensuring that:

- they encompass more interactions than only turn-based ones;
- the property φ holds in all arenas where all local interactions satisfy these restrictions.

Furthermore, the restrictions on local interactions that we define may depend on the property φ at hand. The general method that we will use to define these restrictions proceeds in two steps:

- First, we define the game forms $\mathcal{F}$ that are individually well-behaved w.r.t. φ. Such game forms $\mathcal{F}$ are such that all simple arenas built on $\mathcal{F}$ satisfy the property φ. The notion of "simple arenas built on $\mathcal{F}$" is formally defined in this dissertation. Informally, this corresponds to the arenas where the only source of concurrency comes from $\mathcal{F}$.
- Second, we check that the property φ holds in all the arenas (possibly, only those with finitely many states) with all local interactions that are individually well-behaved w.r.t. the property φ. That is, we check that these individually well-behaved game forms also behave well collectively.

Assuming that we have achieved both of these steps, this provides a way to build arenas that are safe (w.r.t. the property φ) by design. In this dissertation, we will be particularly interested in properties φ involving the existence of optimal positional strategies in parity games. We believe that this way of defining safe by design arenas constitutes the most important contribution of this dissertation.

Outline

Let us give an outline of the kinds of results that we show in this dissertation. We highlight some of these results in this section — those we believe are the most important — but we do not provide an exhaustive overview of all our contributions. As mentioned in the next section, more details are given at the beginning of each part and chapter.

This dissertation contains three parts, preceded by Chapter 1 that presents the formalism used throughout this dissertation.

Then, in Part I, we establish results on concurrent games with almost no assumptions on the local interactions and payoff functions. This part contains two chapters: Chapters 2 and 3. Chapter 2 is dedicated to the determinacy of Blackwell games [12]. This is a very important result on standard concurrent games. Indeed, concurrent games, contrary to turn-based games, do not enjoy Borel determinacy [8], i.e. the existence of winning strategies in win/lose games. However, informally, concurrent games ensure the following: the supremum of what Player A can guarantee is equal to the infimum of what Player B can guarantee. More specifically, from each state of the game, a Player-A strategy along with a Player-B strategy induce an expected value of the payoff function. Then, the value of a Player-A strategy is equal to the infimum of the expected value of the payoff function against all Player-B strategies. The Player-A value of a state is then equal to the supremum of the values of Player-A strategies from this state. This is symmetrical for Player B. Then, the determinacy of Blackwell games ensures, with a quite mild assumption on the standard local interactions, that from each state, the values of the game for both players is equal. This is why we summarize this as follows: the supremum of what Player A can guarantee is equal to the infimum of what Player B can guarantee. In Chapter 2, we generalize this result, in particular, in arbitrary (i.e. not necessarily standard) concurrent games. This generalization is then used several times in this dissertation to establish results on concurrent games. These results could not have been deduced from the original statement of the determinacy of Blackwell games.

On the other hand, the other chapter of Part I, Chapter 3, deals with general properties on concurrent games related to subgame optimal strategies, a strengthening of optimal strategies. In particular, we would like to mention a result that we believe gives significant insight as to why concurrent games behave much more poorly than turn-based games. It is stated as Theorem 3.17. This theorem is stated in the context of games with finitely many states, with some assumptions made on the payoff function including prefix-independence. An informal takeaway from this theorem, from Player-A's perspective, is the following. There exists subgame optimal strategies if and only if there exists a strategy:

- whose value is positive from every state whose Player-A value is positive;
- that never makes a definitive mistake.

A strategy makes a definitive mistake if at some point, after some finite history, it plays a GF-strategy that is sub-optimal[2]. The notion of sub-optimality is defined formally in this dissertation. In other words, the reason why, there are not always (subgame) optimal strategies in concurrent games (even in very simple games) is that to approach the Player-A value of a state, Player-A strategies may need to make a definitive mistake. In fact, in this chapter we exhibit a concurrent game (depicted in Figure 3.2) where playing optimally requires making a definitive mistake after the opponent has made one. This phenomena cannot happen in turn-based games (recall that the games considered have finitely many states). That is, if one only uses strategies that do not make a definitive mistake, then the Player-A values of the states do not change, i.e. they do not drop. This chapter also introduces the notion of infinite-choice strategies, mentioned above. The results of Part I will then be used in the two other parts.

In Part II, containing two chapters, we study arbitrary, i.e. not-necessarily standard, concurrent parity games with finitely many states. As mentioned above, it is not always possible to play optimally in concurrent parity games, and when it is, it may require infinite-choice strategies. In fact, the situation is rather heterogeneous when considering various parity objectives (in terms of the number of colors involved) and flavors of optimality, such as almost-optimality or subgame optimality. We give an (almost) complete overview of the situation in terms of the simplest kind of strategies among which we can find a desirable effect of strategies. A summary is given in page 185. In particular, one can notice a positional/infinite choice dichotomy on the strategies necessary or sufficient to achieve any specific flavor of optimality. This observation is proved in Chapter 3 (in Part I) in standard games.

Finally, Part III is entirely dedicated to the study of game forms and contains four chapters. Three chapters focus on local-global transfers and use the two-step method described in the previous section to define individually well-behaved game forms and prove that the induced arenas also behave well. To achieve these two steps, we use results proved in Chapters 2 and 3, while also making use of some results proved in Part II. Among these three chapters, two of them, Chapter 7 and 8, are dedicated to parity objectives: with arbitrary game forms for Chapter 7, and standard game forms for Chapter 8. A fourth chapter is dedicated to results about complexity/decidability and expressiveness of game forms. We believe that this part is the most important of this dissertation: it tackles the novel method described above to appropriately restrict the class of concurrent arenas.

[2]This notion is actually called being "thrifty " in [37].

On another note, four papers [38, 39, 40, 41] have been published in the scope of this PhD. Most of what is proved in these papers can be found in this dissertation. However, this dissertation contains several results that are not yet published (nor submitted). Note that there is no one-to-one correspondance between chapters and published articles. In addition, several results proved in these papers are generalized to a more general context in this dissertation. It has to be noted that Chapter 8 is entirely based on a single unpublished paper that will be resubmitted soon.

How to read this dissertation

We would like to conclude this introduction by giving a few tips on how this dissertation is meant to be read.

Detailed outline. We have given in the previous section a rather brief outline of what is done in this dissertation. In the body of the dissertation we also provide more local overviews: First, the purpose of each part is explained at the beginning of each of these three parts, and we further explain how each part fits in the broader context of the dissertation. Second, at the beginning of each chapter, we provide a detailed overview of the chapter and of its main contributions.

Finally, the general conclusion recalls the main goals and contributions of the dissertation, and highlights a few research directions that seem promising.

Framed results. We make several contributions in this dissertations. The results that we believe are most important can be spotted by the frames around the corresponding environments. Some new results are not framed. Often, this is because they are stepping stones towards the proof of (what we believe are) more important results. At the end of this dissertation, one can find a list of the main new results proved in this dissertation.

Appendices. Due to the stochasticity of the games and strategies considered in this dissertation, the proof of several results are very technical. That is why we have included a section called “Appendix” at the end of all chapters but Chapter 5. These sections contain technical proofs of results stated in the core of these chapters. Nevertheless, we provide (very often) proof sketches of the results whose formal proofs are put in an “Appendix” section.

Chapters dependence. Chapter 1 gives definitions and notations that will be used throughout this dissertation. Chapter 3 is independent of Chapter 2, except that the main result of Chapter 2 (i.e. Theorem 2.3) is used once in Chapter 3. These two chapters give general results used in subsequent chapters:

- Chapters 4 and 5 depend mostly on Chapters 3;

- Chapter 6 depends mostly on Chapter 2;
- Chapters 7 and 8 depend exclusively on Chapters 3;
- Chapter 9 depends on Chapters 6 and 8 (and a little bit on Chapter 2).

In addition, for the readers reading this dissertation on a PDF file, note that all numerical references (including page references) are clickable. Furthermore, from the numerical denomination of all definitions, theorems, remarks, etc., one can infer the chapters they come from. For instance, Theorem 2.3 comes from Chapter 2 and Definition 7.3 comes from Chapter 7.

Arbitrary Vs Standard terminology. Finally, in this dissertation, we will consider standard game forms which are a special kind of arbitrary game forms, and similarly for concurrent games. When the terminology "arbitrary" is used on game forms (and local interactions) or concurrent games (or arenas), it should actually be read as "not-necessarily standard".

1 - Concurrent games: the formalism

In this chapter, we give the definitions and notations we will use throughout this dissertation. Later, we will introduce additional definitions, however we will do so when we need them.

In Section 1.1, we introduce some notations. Most of them are very classical. In Section 1.2, we give relevant background on probability theory and stochastic trees (generalization of Markov chains). In Section 1.3, we recall the central notion of game forms. We state important properties satisfied by game forms. Finally, in Section 1.4, we present the formalism we use for concurrent games.

1.1 Notations

We denote by $\mathbb{N}, \mathbb{Z}, \mathbb{Q}$ and $\mathbb{R}$ the sets of all non-negative integers, integers, rationals and reals respectively. For two sets E, E' with $E \cap E' = \emptyset$, we denote by $E \uplus E'$ the disjoint union of E and E'. Furthermore, for all pairs of integers $(i, j) \in \mathbb{Z}^2$, we denote by $[\![i, j]\!] := \{k \in \mathbb{Z} \mid i \leq k \leq j\}$ the set of integers between i and j, non-strict. We say that a set is countable if it is either finite or in bijection with $\mathbb{N}$.

Set of sequences. Consider a non-empty set Q. The notations Q^*, Q^+, and Q^ω refer to the set of finite sequences, of non-empty finite sequences and of infinite sequences of elements of Q respectively. Furthermore, $Q^\uparrow := Q^* \cup Q^\omega$ denotes the set of all finite or infinite sequences of elements of Q. The notation $\epsilon \in Q^*$ refers to the empty sequence (which is defined regardless of the underlying set Q considered). For $\bullet \in \{+, \omega\}$, $A \subseteq Q^\bullet$ and $\pi \in Q^+$, the notation $\pi \cdot A \subseteq Q^\bullet$ refers to the set $\pi \cdot A := \{\pi \cdot \rho \mid \rho \in A\}$.

For all $n \in \mathbb{N}$, the notation Q^n (resp. $Q^{\leq n}$) refers to the set of all sequences of n (resp. at most n) elements of Q. The length, denoted $|\pi|$, of a sequence $\pi = \pi_0 \cdots \pi_{n-1} \in Q^n$, is equal to $|\pi| := n$. Furthermore, if $\pi \in Q^\omega$, $|\pi| := \infty$. Given a sequence $\pi = \pi_0 \cdots \in Q^+ \cup Q^\omega$, for $i < |\pi|$, $\pi_{\leq i} \in Q^*$ refers to the finite sequence $\pi_0 \ldots \pi_i$. If $i < 0$, then $\pi_{\leq i} = \epsilon$. For a non-empty sequence $\pi = \pi_0 \cdots \pi_n \in Q^+$, we denote by π_{lt} the last element of the sequence: $\pi_{\mathsf{lt}} = \pi_n$ and by $\mathsf{tl}(\pi)$ the path π but its last element: $\mathsf{tl}(\pi) = \pi_0 \cdots \pi_{n-1}$.

A finite sequence $\pi \in Q^*$ is a prefix (resp. strict prefix) of another finite sequence $\pi' \in Q^*$, denoted $\pi \sqsubseteq \pi'$ (resp. $\pi \sqsubset \pi'$) if there is some $\rho \in Q^*$ (resp. $\rho \in Q^+$) such that $\pi' = \pi \cdot \rho$. The sequence π' is called a suffix of π.

Furthermore, for all $\bullet \in \{*, +, \omega, \uparrow\}$ and arbitrary non-empty sets E, any function $f : E \to Q$ is extended into a function $f^\bullet : E^\bullet \to Q^\bullet$ such that for all $\pi = \pi_0 \cdots \in E^\bullet$, we have $f^\bullet(\pi) := f(\pi_0) \cdots \in Q^\bullet$.

We also define the notion of residual functions. Consider some $\bullet \in \{*, +, \omega\}$.

Given two non-empty sets E, F, for all functions $f : E^\bullet \to F$ and finite paths $\pi \in E^*$, we denote by $f^\pi : E^\bullet \to F$ the residual function defined by, for all $\rho \in E^\bullet$, we have $f^\pi(\rho) := f(\pi \cdot \rho) \in F$.

Real functions and sequences. Consider two arbitrary non-empty sets E and F and a function $f : E \to F$. For all $F' \subseteq F$, we let $f^{-1}[F'] := \{e \in E \mid f(e) \in F'\}$ be the reverse image of F' by f. For G an arbitrary non-empty set and $f : E \to F$ and $g : F \to G$, we denote by $g \circ f : E \to G$ the composite of the functions f and g.

Assume now that $F \subseteq \mathbb{R}$. Consider two real functions $f, g : E \to F$. We write $f \leq g$ if, for all $x \in E$, we have $f(x) \leq g(x)$. Furthermore, we let $\|f\|_\infty \in \mathbb{R} \cup \{\infty\}$ be equal to $\|f\|_\infty := \sup_{e \in E} |f(e)|$. In addition, if the support $\{x \in E \mid f(x) \neq 0\}$ of f is countable, we also let $\|f\|_1 \in \mathbb{R} \cup \{\infty\}$ be equal to $\|f\|_1 := \sum_{e \in E} |f(e)|$.

Consider an arbitrary infinite sequence $(f_n)_{n \in \mathbb{N}}$ of real functions indexed by $\mathbb{N}$ with $f_n : E \to F$ and $F \subseteq \mathbb{R}$[1]. A function $f : E \to F$ is a limit of f_n w.r.t. $\| \cdot \|_\infty$ if:

$$\forall \varepsilon > 0, \ \exists n \in \mathbb{N}, \forall k \geq n, \ \|f - f_k\|_\infty \leq \varepsilon$$

We can define similarly a limit w.r.t. $\|\cdot\|_1$ when E is countable. A subsequence $(f_n)_{n \in \mathbb{N}}$ is a sequence $(f_{\varphi(n)})_{n \in \mathbb{N}}$ for an increasing function $\varphi : \mathbb{N} \to \mathbb{N}$. A subsequential limit of $(f_n)_{n \in \mathbb{N}}$ is the limit of some subsequence of $(f_n)_{n \in \mathbb{N}}$. Furthermore, note that if E is finite and $F = [0, 1]$ then Bolzano-Weierstrass' theorem ensures that such a subsequential limit exists, as stated below in Theorem 1.1.

Theorem 1.1. *Let E be a finite non-empty set and let $x = (f_n)_{n \in \mathbb{N}}$ be an infinite sequence of real functions with $f_n : E \to [0, 1]$. Then, x has a subsequential limit (w.r.t. $\| \cdot \|_1$ and $\| \cdot \|_\infty$).*

1.2 Probability measures, distributions and stochastic trees

In this section, we fix a non-empty set Q.

1.2.1 . Probability distribution

The support $\mathsf{Sp}(\mu)$ of a function μ into $[0, 1]$ is the set of elements whose image by μ is nonzero: $\mathsf{Sp}(\mu) = \mu^{-1}[\ (0, 1]\ ]$. A discrete probability distribution (or simply distribution) over Q is a function $\mu : Q \to [0, 1]$ with countable support such that $\sum_{x \in Q} \mu(x) = 1$. The set of all distributions over the set Q is denoted $\mathcal{D}(Q)$. A distribution μ is deterministic if $|\mathsf{Sp}(\mu)| = 1$. In addition, for all functions $f : Q \to [0, 1]$, with an abuse of notations, the sum $\sum_{q \in Q} \mu(q) \cdot f(q)$ refers to the countable sum $\sum_{q \in \mathsf{Sp}(\mu)} \mu(q) \cdot f(q)$. In the following, an element

[1] If E is a singleton, the functions f_n can be seen as real numbers.

$q \in Q$ will be seen as the deterministic (Dirac) distribution $\mu : Q \to [0,1]$ such that $\mu(q) = 1$ (and therefore $\mu(q') = 0$ for all $q' \neq q$).

Given any valuation over Q, $v : Q \to [0,1]$, and distribution $d \in \mathcal{D}(Q)$, we consider the expected value of v w.r.t. d: $\mathbb{E}_d(v) := \sum_{q \in Q} d(q) \cdot v(q)$. We can also consider the expected value of distributions. Consider two sets Q and Q' with $d \in \mathcal{D}(Q)$ and $d' : Q \to \mathcal{D}(Q')$. The expected value of d' w.r.t. d, denoted $\mathbb{E}_d(d') \in \mathcal{D}(Q')$ is such that, for all $q' \in Q'$:

$$\mathbb{E}_d(d')(q') := \sum_{q \in Q} d(q) \cdot d'(q)(q')$$

1.2.2 . Topology on Q^ω and probability measure

A σ-algebra $\mathcal{Q}$ on a set Q^ω is such that $\mathcal{Q}$ is a set of subsets of Q^ω where $Q^\omega \in \mathcal{Q}$, and $\mathcal{Q}$ is closed under complementation and countable union. That is, for all $E \in \mathcal{Q}$ we have $Q^\omega \setminus E \in \mathcal{Q}$ and for all $(E_n)_{n \in \mathbb{N}} \in \mathcal{Q}^{\mathbb{N}}$, we have $\cup_{n \in \mathbb{N}} E_n \in \mathcal{Q}$. A probability measure on a σ-algebra $\mathcal{Q}$ is a function $\Upsilon : \mathcal{Q} \to [0,1]$ such that $\Upsilon(\emptyset) = 0$, $\Upsilon(Q^\omega) = 1$, and Υ is σ-additive over $\mathcal{Q}$, that is for all $(E_n)_{n \in \mathbb{N}} \in \mathcal{Q}^{\mathbb{N}}$ pairwise disjoint, we have $\Upsilon(\uplus_{n \in \mathbb{N}} E_n) = \sum_{n \in \mathbb{N}} \Upsilon(E_n)$.

Let us now recall the definition of cylinder sets and consequently of Borel sets. For all finite sequences $\pi \in Q^*$, the cylinder set $\mathsf{Cyl}(\pi)$ generated by π is the set $\mathsf{Cyl}(\pi) = \{\pi \cdot \rho \in Q^\omega \mid \rho \in Q^\omega\}$. We denote by Cyl_Q the set of all cylinder sets on Q^ω. The open sets of Q^ω are the sets equal to an arbitrary union of cylinder sets. The set of Borel sets on Q^ω, denoted $\mathsf{Borel}(Q)$, is then equal to the smallest σ-algebra containing all open sets. By Carathéodory's theorem, a probability measure over Q^ω is entirely defined by the measure of all cylinder sets as stated in Theorem 1.2 below.

Theorem 1.2. *Consider a function $v : \mathsf{Cyl}_Q \to [0,1]$ such that $v(Q^\omega) = 1$ (note that $Q^\omega = \mathsf{Cyl}(\epsilon)$) and v is σ-additive over Cyl_Q. Then, there exists a unique probability measure $\Upsilon : \mathsf{Borel}(Q) \to [0,1]$ such that, for all $C \in \mathsf{Cyl}_Q$, we have $\Upsilon(C) = v(C)$,*

We also mention the monotone continuity of the probability measure.

Proposition 1.3. *Consider any probability measure $\Upsilon : \mathsf{Borel}(Q) \to [0,1]$. For all $(A_n)_{n \in \mathbb{N}} \in (\mathsf{Borel}(Q))^{\mathbb{N}}$ such that, for all $n \in \mathbb{N}$, we have $A_n \subseteq A_{n+1}$ (resp. $A_n \supseteq A_{n+1}$), then we have $\Upsilon(\cup_{n \in \mathbb{N}} A_n) = \lim_{n \to \infty} \Upsilon(A_n)$ (resp. $\Upsilon(\cap_{n \in \mathbb{N}} A_n) = \lim_{n \to \infty} \Upsilon(A_n)$).*

Finally, in the following, any subset of finite sequences $E \subseteq Q^*$ may be seen, in Q^ω, as the set $E \cdot Q^\omega \subseteq Q^\omega$, which is equal to the open set $\cup_{\pi \in E} \mathsf{Cyl}(\pi) \in \mathsf{Borel}(Q)$.

1.2.3 . Measurable functions and integrals

The definitions we present in this subsection are classical, but we specifically used the presentation given in [42] where the proofs we do not give here can be found.

We define what the measurable functions from Q^ω to $[0,1]$ are (for the Borel topology we consider).

Definition 1.1 (Measurable function). *A function $f : Q^\omega \to [0,1]$ is* measurable *if, for all $\alpha \in [0,1]$, we have $f^{-1}([0,\alpha]) \in \mathsf{Borel}(Q)$.*

In fact, all residual functions obtained from a measurable function are also measurable.

Proposition 1.4 (Proof 1.5.1). *Consider a non-empty set Q, a measurable function $f : Q^\omega \to [0,1]$ and finite path $\pi \in Q^*$. For all Borel sets $B \in \mathsf{Borel}(Q)$, the sets $\pi \cdot B \in Q^\omega$ and $\pi^{-1} \cdot B := \{\rho \in Q^\omega \mid \pi \cdot \rho \in B\}$ are also Borel. Consequently, the residual function $f^\pi : Q^\omega \to [0,1]$ is measurable.*

Given a measure of the Borel sets of Q^ω, we can define the integral (or equivalently, the expected value) of a measurable function from Q^ω to $[0,1]$. The integral of a function is first defined on the simple step functions defined below and then extended to arbitrary measurable functions.

Definition 1.2 (Step functions). *A* step function *on Q^ω is a function $f : Q^\omega \to [0,1]$ such that there exists a finite collection of pairwise disjoint Borel sets $(E_i)_{1\leq i\leq n} \in (\mathsf{Borel}(Q))^n$ such that $\cup_{i=1}^n E_i = Q^\omega$ and $f = \sum_{i=1}^n \alpha_i \cdot \mathbb{1}_{E_i}$ with $\alpha_i \in [0,1]$ for all $1 \leq i \leq n$. The notation $\mathbb{1}_{E_i}$ refers to the indicator of the set E_i: $\mathbb{1}_{E_i}[Q^\omega] \subseteq \{0,1\}$ and for all $\rho \in Q^\omega$, we have $\mathbb{1}_{E_i}(\rho) = 1 \Leftrightarrow \rho \in E_i$.*

Interestingly, we have the following proposition.

Proposition 1.5. *Consider a probability measure $\mathbb{P}$ on $\mathsf{Borel}(Q)$, a step function $f : Q^\omega \to [0,1]$ and two finite collections $(E_i)_{1\leq i\leq n} \in (\mathsf{Borel}(Q))^n$ and $(E'_j)_{1\leq j\leq n} \in (\mathsf{Borel}(Q))^m$ of pairwise disjoint sets such that $\cup_{i=1}^n E_i = Q^\omega = \cup_{j=1}^m E'_i$ and $f = \sum_{i=1}^n \alpha_i \cdot \mathbb{1}_{E_i} = \sum_{j=1}^m \alpha'_j \cdot \mathbb{1}_{E'_j}$ with $\alpha_i \in [0,1]$ for all $1 \leq i \leq n$ and $\alpha'_j \in [0,1]$ for all $1 \leq j \leq m$. Then, for all Borel sets $B \in \mathsf{Borel}(Q)$:*

$$\sum_{i=1}^n \alpha_i \cdot \mathbb{P}[E_i \cap B] = \sum_{j=1}^m \alpha'_j \cdot \mathbb{P}[E'_j \cap B]$$

This proposition above justifies the definition below of the integral of a step function w.r.t. to a probability measure.

Definition 1.3. *Consider a probability measure $\mathbb{P}$ and a step function $f = \sum_{i=1}^n \alpha_i \cdot \mathbb{1}_{E_i}$. Then, for all Borel sets $B \in \mathsf{Borel}(Q)$:*

$$\int_B f d\mathbb{P} := \sum_{i=1}^n \alpha_i \cdot \mathbb{P}[E_i \cap B]$$

Let us now define the integral of an arbitrary measurable function. To do so, we will approximate it with step functions. Specifically:

Definition 1.4 (Approximating a measurable function). *Consider a measurable function $f : Q^\omega \to [0,1]$. A sequence of step functions $(f_n)_{n \in \mathbb{N}}$ such that, for all $n \in \mathbb{N}$: $f_n : Q^\omega \to [0,1]$* approximates the function *f if, for all $\rho \in Q^\omega$:*

- *for all $n \in \mathbb{N}$, we have $f_n(\rho) \leq f_{n+1}(\rho)$;*
- $\lim_{n \to \infty} f_n(\rho) = \sup_{n \in \mathbb{N}} f_n(\rho) = f(\rho)$.

Then, we have the very useful theorem below: the supremum of the integral of a sequence of step functions approximating a measurable function does not depend on the sequence of step functions considered (it only depends on the measurable function approximated).

Theorem 1.6. *Consider a probability measure $\mathbb{P}$ and a measurable function $f : Q^\omega \to [0,1]$. Then:*

- *There exists a sequence of step functions approximating f;*
- *For all Borel sets $B \in \mathsf{Borel}(Q)$, there exists a value $v_B \in [0,1]$ such that, for all sequences of step functions $(f_n)_{n \in \mathbb{N}}$ approximating f, we have $\lim_{n \to \infty} \int_B f_n d\mathbb{P} = v_B$.*

We are now able to define the integral of any measurable function: it is defined as the limit of any sequence of step functions approximating it.

Definition 1.5 (Integral of a measurable function). *Consider a non-empty set Q, a probability measure $\mathbb{P}$ and a measurable function $f : Q^\omega \to [0,1]$. Consider a sequence of step functions $(f_n)_{n \in \mathbb{N}}$ approximating f. Then, for all Borel sets $B \in \mathsf{Borel}(Q)$:*

$$\int_B f d\mathbb{P} := \lim_{n \to \infty} \int_B f_n d\mathbb{P}$$

Note that in particular if $\mathbb{P}(B) = 0$, then $\int_B f d\mathbb{P} = 0$ (since this holds for step functions).

1.2.4 . Stochastic trees and Markov chains

Let us now define the crucial notion of stochastic tree and Markov chain. It will be extensively used in the remainder of this dissertation as stochastic trees is what is obtained from a concurrent arena when both players have fixed their strategies. In Definition 1.6 below, we define formally the notion of stochastic trees, and the special case of Markov chains.

Definition 1.6 (Stochastic Tree and Markov chain). A stochastic tree $\mathcal{T}$ *is a pair $\langle Q, \mathbb{P} \rangle$ where $Q \neq \emptyset$ is a non-empty set of states and $\mathbb{P} : Q^+ \to \mathcal{D}(Q)$ is a probability distribution function.*

When the function $\mathbb{P}$ ensures that, for all $q \in Q$ and $\rho \in Q^+$, we have $\mathbb{P}(q) = \mathbb{P}(\rho \cdot q)$, then the tree $\mathcal{T}$ is actually a Markov chain *$\mathcal{M} = \langle Q, \mathbb{P} \rangle$. In that case, the probability distribution function $\mathbb{P}$ can be seen as a function $\mathbb{P} : Q \to \mathcal{D}(Q)$.*

Alternatively, Markov chains can be defined as sequences of random variables where the value of the next variable only depends on the value of the current variable, not all preceding ones. In this dissertation, we use the definition that we gave above of Markov chains as it is better-suited for the game formalism we will use in the following.

We now define, in stochastic trees, the probability of occurrence of any finite sequence, and consequently of any Borel set and deduce the expected value of any measurable function.

Definition 1.7 (Probability of Borel sets, Expected value in stochastic trees)**.** *Consider a stochastic trees $\mathcal{T} = \langle Q, \mathbb{P} \rangle$. Fix a finite sequence of states $\rho \in Q^+$. We define the function $\mathbb{P}_\rho : Q^* \to [0,1]$ which we then extend into a probability measure $\mathbb{P}_\rho : \mathsf{Borel}(Q) \to [0,1]$. First, the probability of occurrence of a finite path $\pi \in Q^*$ is equal to:*

$$\mathbb{P}_\rho(\pi) := \prod_{i=0}^{|\pi|-1} \mathbb{P}(\rho \cdot \pi_{\leq i-1})(\pi_i)$$

In particular, $\mathbb{P}_\rho(\epsilon) = 1$. Then, for all finite paths $\pi \in Q^*$, the probability of a cylinder set $\mathsf{Cyl}(\pi)$ is:

$$\mathbb{P}_\rho[\mathsf{Cyl}(\pi)] := \mathbb{P}_\rho(\pi)$$

This induces a probability measure over Borel sets, that we denote by $\mathbb{P}^{\mathcal{T}}_\rho : \mathsf{Borel}(Q) \to [0,1]$, via Theorem 1.2.

For all measurable functions $f : Q^\omega \to [0,1]$, we denote by $\mathbb{E}^{\mathcal{T}}_\rho[f] \in [0,1]$ the expected value of the function f w.r.t. the probability measure $\mathbb{P}_\rho$: $\mathbb{E}^{\mathcal{T}}_\rho[f] := \int_{Q^\omega} f^\rho d\mathbb{P}^{\mathcal{T}}_\rho$ (recall the notion of residual function Page 26).

In the following, we will sometimes need to relate two stochastic trees, whose behaviors are different but similar. More specifically, in the following we will need to add intermediate states in the games we will consider. This will be useful either to encode some additional information in the states or to be able to use results on what we will call turn-based games (or both). In that case, we would like to be able to show that a stochastic tree obtained from the original game behaves similarly to the one obtained from the modified game. To do so, we introduce the notion of stochastic tree alternating between two sets of states Q and Q': there is probability 0 to go from a state in Q to another state in Q and similarly for Q'. We can then obtain a payoff function $f_{Q,Q'} : (Q \cup Q')^\omega \to [0,1]$ from a payoff function $f : Q^\omega \to [0,1]$.

Definition 1.8 (Alternating stochastic trees, Projecting measurable function)**.** *Consider two disjoint sets of states Q and Q'. A stochastic tree $\mathcal{T} = \langle Q \uplus Q', \mathbb{P} \rangle$ is (Q,Q')-*alternating *if, for all $q \in Q$ and $\rho \notin Q' \cdot ((Q \cdot Q')^* \cup (Q \cdot Q')^* \cdot Q)$, we have: $\mathbb{P}_q(\rho) = 0$.*

We denote by $\phi_{Q,Q'} : (Q \cdot Q')^\uparrow \to Q^\uparrow$ the function that extracts the elements in Q by considering every other element. Consider a function $f : Q^\omega \to [0,1]$.

We denote by $f_{Q,Q'} : (Q \cup Q')^\omega \to [0,1]$ the function such that, for all $\rho \in (Q \cup Q')^\omega$, we have:

$$f_{Q,Q'}(\rho) := \begin{cases} 0 & \text{if } \rho \notin (Q \cdot Q')^\omega \\ f \circ \phi_{Q,Q'}(\rho) & \text{otherwise} \end{cases}$$

Note that we consider these definitions even if Q and Q' are not disjoint.

We have the following lemma.

Lemma 1.7 (Proof 1.5.2)**.** *Consider two non-empty sets of states Q and Q'. For all measurable functions $f : Q^\omega \to [0,1]$, the function $f_{Q,Q'} : (Q \cup Q')^\omega \to [0,1]$ is also measurable. Consider now two stochastic trees $\mathcal{T} = \langle Q, \mathbb{P} \rangle$ and $\mathcal{T}' = \langle Q \cup Q', \mathbb{P}' \rangle$ and let $q \in Q$. For all $\pi \in Q^*$, we consider the set $\mathsf{T}(\pi) := Q' \cdot (\phi_{Q,Q'})^{-1}[\{\pi\}]$. Now, assume that $\mathcal{T}'$ is (Q, Q')-alternating and that, for all $\pi \in Q^*$, we have:*

$$\mathbb{P}_q[\mathsf{Cyl}(\pi)] = \mathbb{P}'_q[\cup_{\pi' \in \mathsf{T}(\pi)} \mathsf{Cyl}(\pi')]$$

Then, for all measurable functions $f : Q^\omega \to [0,1]$, we have:

$$\mathbb{E}_q[f^q] = \mathbb{E}'_q[(f_{Q,Q'})^q]$$

This lemma will be used in the following in Chapters 2 and 3.

Finally, we define the notion of Bottom strongly connected component in a Markov chain which informally is a set of states that is strongly connected and that cannot be exited.

Definition 1.9 (Bottom strongly connected component)**.** *Consider a Markov chain $\mathcal{M} = \langle Q, \mathbb{P} \rangle$. A* bottom strongly connected component *(BSCC for short) of $\mathcal{M}$ is a subset of states $B \subseteq Q$ such that:*

- *B is strongly connected, that is for all $(q, q') \in B^2$, there is a finite path $\pi \in B^*$ such that $\mathbb{P}_q(\pi \cdot q') > 0$;*
- *B cannot be exited, that is for all $q \in B$ and $q' \in Q$, we have $\mathbb{P}_q(q') > 0$ implies $q' \in B$.*

We denote by $\mathcal{B}_\mathcal{M}$ the set of all BSCCs in the Markov chain $\mathcal{M}$.

In fact, when the set of states Q of a Markov chain is finite, the set of BSCCs is not empty and almost-surely, regardless of the starting state considered, the Markov chain eventually settles in a BSCC B. Furthermore, every state in B is seen infinitely often. This is a well-known result, see for instance [43, Theorem 10.27], that we recall below in Theorem 1.8.

Theorem 1.8. *Consider a Markov chain $\mathcal{M} = \langle Q, \mathbb{P} \rangle$ and assume that Q is finite. Then, $\mathcal{B}_\mathcal{M} \neq \emptyset$ and, for all $q \in Q$:*

$$\mathbb{P}_q\left[\bigcup_{B \in \mathcal{B}_\mathcal{M}} \left((Q^* \cdot B^\omega) \cap \left(\bigcap_{q \in B} (Q^* \cdot \{q\})^\omega \right) \right) \right] = 1$$

Disclaimer: game terminology for both players. In the remainder of this dissertation, we will consider games with two players we will call Player A and Player B. Often, definitions, lemmas and theorems can be applied to the two players, but sometimes with slight modifications. We will often state them only for one of the players (usually Player A). However, the case for the other player may be either **similar**, in which case (almost) no modification needs to be done to be applied to the other player, or **symmetrical**, in which case one has to reverse inequalities and supremum and infimum (when applicable) to obtain the appropriate definition for the other player.
Moreover, we will define properties w.r.t. one player. Unless otherwise stated, these can also be ensured w.r.t. the other player. Furthermore, when we say that such a property is ensured without mentioning any player, it means that this property holds for both players (see for instance Definition 1.15).

1.3 Game Forms

In this section, we define and discuss the crucial notion of game forms. A game form represents an interaction between two players — that we call Player A and Player B — with a set of strategies for Player A, a set of strategies for Player B, a set of outcomes and a function mapping from a pair of strategies (one per player) to a probability distribution over the set of outcomes. A special class of game forms — that we will call standard game forms — consists in interactions where the set of strategies of a player is equal to the set of distributions over an underlying set of actions. We call them standard game forms because they in fact correspond to the interactions that are almost always used in games. In particular, in all the articles published in the scope of this PhD [38, 39, 40, 41], we have used standard game forms. We define them formally below in Definition 1.10.

Definition 1.10 ((Standard) Game Form)**.** *Consider a non-empty set of outcomes O. A* game form *(GF for short) $\mathcal{F}$ on O is a tuple $\mathcal{F} = \langle \Sigma_{\mathsf{A}}, \Sigma_{\mathsf{B}}, \mathsf{O}, \varrho \rangle$ where Σ_{A} (resp. Σ_{B}) is the non-empty set of strategies available to Player A (resp. B) and $\varrho : \Sigma_{\mathsf{A}} \times \Sigma_{\mathsf{B}} \to \mathcal{D}(\mathsf{O})$ maps a pair of strategies to a probability distribution over outcomes. Recall that all probability distributions that we consider have a countable support. We denote by $\mathsf{Form}(\mathsf{O})$ the set of game forms on the set of outcomes O.*

A game form $\mathcal{F}$ on O is standard *if $\Sigma_{\mathsf{A}} = \mathcal{D}(\mathsf{Act}_{\mathsf{A}})$ and $\Sigma_{\mathsf{B}} = \mathcal{D}(\mathsf{Act}_{\mathsf{B}})$ for some underlying non-empty sets of actions $\mathsf{Act}_{\mathsf{A}}$ (resp. $\mathsf{Act}_{\mathsf{B}}$) available to Player A (resp. B) and $\varrho : \mathsf{Act}_{\mathsf{A}} \times \mathsf{Act}_{\mathsf{B}} \to \mathcal{D}(\mathsf{O})$. Furthermore, the map $\mathbb{E}(\varrho) : \mathcal{D}(\mathsf{Act}_{\mathsf{A}}) \times \mathcal{D}(\mathsf{Act}_{\mathsf{B}}) \to \mathcal{D}(\mathsf{O})$ is such that, for all $\sigma_{\mathsf{A}} \in \mathcal{D}(\mathsf{Act}_{\mathsf{A}})$ and*

$$A \overset{B}{\begin{bmatrix} x & y \\ y & x \end{bmatrix}} \qquad A \overset{B}{\begin{bmatrix} x \\ y \end{bmatrix}} \qquad A \overset{B}{\begin{bmatrix} x & y \end{bmatrix}}$$

Figure 1.1: Three standard finite game forms on the set of outcomes $\{x, y\}$.

$\sigma_\mathsf{B} \in \mathcal{D}(\mathsf{Act}_\mathsf{B})$, *for all* $o \in \mathsf{O}$, *we have (recall the beginning of Section 1.2):*

$$\mathbb{E}(\varrho)(\sigma_\mathsf{A}, \sigma_\mathsf{B})(o) := \mathbb{E}_{\sigma_\mathsf{A},\sigma_\mathsf{B}}(\varrho)(o) = \sum_{a \in \mathsf{Sp}(\sigma_\mathsf{A})} \sum_{b \in \mathsf{Sp}(\sigma_\mathsf{B})} \sigma_\mathsf{A}(a) \cdot \sigma_\mathsf{B}(b) \cdot \varrho(a, b)(o)$$

Such a standard game form is described by the tuple $\mathcal{F} = \langle \mathsf{Act}_\mathsf{A}, \mathsf{Act}_\mathsf{B}, \mathsf{O}, \varrho \rangle_s$*, where* $_s$ *specifies that the game form is standard.*

Below we illustrate the definition of standard game forms.

Example 1.1 (Standard game forms)**.** *Some standard game forms are represented in Figure 1.1. They are represented as bi-dimensional tables where Player* A*'s actions are the rows, and Player* B*'s are the columns. For instance, for the leftmost game form* $\mathcal{F}$ *in Figure 1.1, denoting* t, b *the top and bottom rows respectively and* l, r *the left and right columns, we have* $\mathcal{F} := \langle \{t, b\}, \{l, r\}, \{x, y\}, \varrho \rangle_s$ *with* $\varrho(t, l) := \varrho(b, r) = x$ *and* $\varrho(b, l) := \varrho(t, r) = y$*. The colors used for the outcomes* x *and* y *are only there to ease the readability of the game forms. For the remainder of this dissertation, we will omit, when drawing game forms, the* A *on the left and the* B *on the top of the game forms.*

We would now like to discuss the benefit of considering non-standard game forms. Clearly, they are more general than standard game forms. Furthermore, this more general framework allows to define relevant situations. For instance, with non-standard game forms, one can express that, given a set of actions, the only possible strategies available to a player are exactly deterministic probability distributions, or are exactly probability distributions with a fixed precision. In addition, there are some players interactions that can be defined more concisely with non-standard game forms than with standard game forms, as discussed below.

Example 1.2 (Non-standard game form)**.** *Consider the standard game form* $\mathcal{F}$ *in the middle of Figure 1.1. Let us now consider the non-standard game form* $\mathcal{F}'$ *from* $\mathcal{F}$ *by considering that the strategies available to Player* A *are all the distributions that do not play the bottom row with probability 1. To describe such an interaction with a standard game form, one needs infinitely many Player-*A *actions. For instance, a standard game form defined by* $\mathcal{F}_\mathbb{N} := \langle \mathbb{N}, \{*\}, \{x, y\}, \varrho \rangle_s$ *where, for all* $n \in \mathbb{N}$*, we have* $\varrho(n, *) \in \mathcal{D}(\{x, y\})$ *such that* $\varrho(n, *)(x) := \frac{1}{2^n}$ *and* $\varrho(n, *)(y) := 1 - \frac{1}{2^n}$*. One can see that the strategies* *available to Player* A *are the same in* $\mathcal{F}'$ *and* $\mathcal{F}_{\mathbb{N}}$*, when considering the corresponding distribution over outcomes. However, this is only possible because there are infinitely many actions for Player* A *with an increasing probability to see* y*, without it being equal to 1.*

For the remainder of this dissertation and unless otherwise stated, when the set of outcomes O is clear from context, the notation $\mathcal{F}$ will always refer to the tuple $\langle \Sigma_{\mathsf{A}}, \Sigma_{\mathsf{B}}, \mathsf{O}, \varrho \rangle$ if we consider an arbitrary game form. If the game form is standard, it will also refer to the tuple $\langle \mathsf{Act}_{\mathsf{A}}, \mathsf{Act}_{\mathsf{B}}, \mathsf{O}, \varrho \rangle_s$ with $\Sigma_{\mathsf{A}} = \mathcal{D}(\mathsf{Act}_{\mathsf{A}})$ and $\Sigma_{\mathsf{B}} = \mathcal{D}(\mathsf{Act}_{\mathsf{B}})$. For the remainder of this section, we fix an arbitrary non-empty set of outcomes O.

We would like to mention several special classes of game forms. Firstly, trivial game forms, i.e. game forms such that there is only one possible distribution over outcomes, regardless of what the players do. Then, among standard game forms, there are finite game forms, i.e. game forms where the sets of actions of both players are finite. Turn-based game forms are such that the set of actions of either of the players is a singleton. Finally, deterministic game forms are such that each pair of actions is mapped to an outcome with probability 1. This is defined formally below.

Definition 1.11 (Trivial, turn-based, standard finite, deterministic game forms)**.** *Consider an arbitrary game form* $\mathcal{F}$*. It is* trivial *if the function* $\varrho : \Sigma_{\mathsf{A}} \times \Sigma_{\mathsf{B}} \to \mathcal{D}(\mathsf{O})$ *is constant. Assume now that the game form* $\mathcal{F}$ *is standard. We say that* $\mathcal{F}$ *is* finite *if both sets of actions* $\mathsf{Act}_{\mathsf{A}}$ *and* $\mathsf{Act}_{\mathsf{B}}$ *are finite.*

We say that it is a Player-A game form *if* $|\mathsf{Act}_{\mathsf{B}}| = 1$ *(no assumption is made on* $\mathsf{Act}_{\mathsf{A}}$*). In a Player-*$\mathsf{A}$ *game form* $\mathcal{F}$*, the only Player-*B *action is denoted* $*$*, and analogously for a Player-*B *game form. When* $\mathcal{F}$ *is either a Player-*A *or a Player-*B *game form, it said to be a* turn-based *game form.*

The game form $\mathcal{F}$ *is deterministic if, for all* $(a, b) \in \mathsf{Act}_{\mathsf{A}} \times \mathsf{Act}_{\mathsf{B}}$*, we have* $|\mathsf{Sp}(\varrho(a, b))| = 1$*. Furthermore, a Player-*A *strategy* $\sigma_{\mathsf{A}} \in \mathcal{D}(\mathsf{Act}_{\mathsf{A}})$ *is* deterministic *if* $|\mathsf{Sp}(\sigma_{\mathsf{A}})| = 1$*. This is similar for Player* B*.*

In Figure 1.1, all three standard game forms are finite and the two rightmost game forms are turn-based: the middle one is a Player-A game form and the rightmost one is a Player-B game form.

When the set of outcomes is equal to $[0, 1]$, we call the game form a game in normal form (this terminology comes from [44]). Note that we can obtain games in normal form from game forms with a map from the set of outcomes to $[0, 1]$, as described in Definition 1.12 below.

Definition 1.12 (Game in normal form)**.** *A game form* $\mathcal{F} \in \mathsf{Form}([0, 1])$ *is a* game in normal form*. Given a game form* $\mathcal{F} = \langle \Sigma_{\mathsf{A}}, \Sigma_{\mathsf{B}}, \mathsf{O}, \varrho \rangle$ *and a valuation* $v : \mathsf{O} \to [0, 1]$*, the notation* $\langle \mathcal{F}, v \rangle$ *refers to the game in normal form* $\langle \mathcal{F}, v \rangle := \langle \Sigma_{\mathsf{A}}, \Sigma_{\mathsf{B}}, [0, 1], \mathbb{E}_{\varrho(\cdot,\cdot)}(v) \rangle$ *with* $\mathbb{E}_{\varrho(\cdot,\cdot)}(v) : \Sigma_{\mathsf{A}} \times \Sigma_{\mathsf{B}} \to [0, 1]$ *such that, for all* $(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \in \Sigma_{\mathsf{A}} \times \Sigma_{\mathsf{B}}$*, we have* $\mathbb{E}_{\varrho(\cdot,\cdot)}(v)(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) = \mathbb{E}_{\varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})}(v) \in [0, 1]$*.*

Remark 1.1. *In this dissertation, we only consider game in normal forms as defined above. In particular, the outcomes in games in normal forms are reals in* $[0, 1]$*. However, we could consider instead any bounded set of reals with (almost) identical definitions. Subsequently, the valuations of outcomes we will consider will also take their values in* $[0, 1]$*. However, we could also consider valuations of outcomes taking their values in an arbitrary bounded real set.*

Consider such a game in normal form $\mathcal{F}$. Given a pair of strategies $(\sigma_\mathsf{A}, \sigma_\mathsf{B}) \in \Sigma_\mathsf{A} \times \Sigma_\mathsf{B}$ — one per player — the corresponding distribution in $\mathcal{D}([0, 1])$ is given by ϱ. The expected values of the outcomes defines the outcome of the game in normal form $\mathcal{F}$ under the pair of strategies $(\sigma_\mathsf{A}, \sigma_\mathsf{B})$, as defined below in Definition 1.13.

Definition 1.13 (Outcome given two strategies)**.** *Consider a game in normal form* $\mathcal{F}$ *and a pair of strategies* $(\sigma_\mathsf{A}, \sigma_\mathsf{B}) \in \Sigma_\mathsf{A} \times \Sigma_\mathsf{B}$*. The* outcome $\mathsf{out}[\mathcal{F}](\sigma_\mathsf{A}, \sigma_\mathsf{B})$ *of the game in normal form* $\mathcal{F}$ *under* $(\sigma_\mathsf{A}, \sigma_\mathsf{B})$ *is equal to (recalling the notation* $\mathbb{E}$ *from the beginning of Section 1.2):*

$$\mathsf{out}[\mathcal{F}](\sigma_\mathsf{A}, \sigma_\mathsf{B}) := \mathbb{E}(\varrho(\sigma_\mathsf{A}, \sigma_\mathsf{B}))$$

In games in normal form, Player A tries to maximize the outcome whereas Player B tries to minimize it. Then, what we call the value of a Player-A strategy σ_A in a game in normal form $\mathcal{F}$ is the best that this strategy can achieve against all Player B strategies. Hence, it is equal to the infimum, over all Player B strategies σ_B, of the outcome of $\mathcal{F}$ under $(\sigma_\mathsf{A}, \sigma_\mathsf{B})$. Following, the Player-$\mathsf{A}$ value of $\mathcal{F}$ is the supremum of the values of her strategies. This is defined formally below in Definition 1.14.

Definition 1.14 (Value in games in normal form)**.** *Consider a game in normal form* $\mathcal{F}$ *and a Player-*A *strategy* $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}$*. The* value $\mathsf{val}[\mathcal{F}](\sigma_\mathsf{A})$ *of the strategy* σ_A *in the game in normal form* $\mathcal{F}$ *is equal to:*

$$\mathsf{val}[\mathcal{F}](\sigma_\mathsf{A}) := \inf_{\sigma_\mathsf{B} \in \Sigma_\mathsf{B}} \mathsf{out}[\mathcal{F}](\sigma_\mathsf{A}, \sigma_\mathsf{B})$$

Then, the Player-A *value of the game in normal form* $\mathcal{F}$ *is equal to:*

$$\mathsf{val}[\mathcal{F}](\mathsf{A}) := \sup_{\sigma_\mathsf{A} \in \Sigma_\mathsf{A}} \mathsf{val}[\mathcal{F}](\sigma_\mathsf{A})$$

For all $\varepsilon > 0$*, a Player-*A *strategy* $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}$ *ensuring* $\mathsf{val}[\mathcal{F}](\sigma_\mathsf{A}) \geq \mathsf{val}[\mathcal{F}](\mathsf{A}) - \varepsilon$ *is said to be* ε-optimal *in* $\mathcal{F}$*. When* $\varepsilon = 0$*, the strategy* σ_A *is simply said to be* optimal*. We denote by* $\mathsf{Opt}_\mathsf{A}(\mathcal{F}) \subseteq \Sigma_\mathsf{A}$ *the set of Player-*A *strategies optimal in* $\mathcal{F}$*. The definitions and notations are symmetrical for Player* B*. Then, when the values of the game in normal form* $\mathcal{F}$ *for both players are equal, i.e.* $\mathsf{val}[\mathcal{F}](\mathsf{A}) = \mathsf{val}[\mathcal{F}](\mathsf{B})$*, this defines the value of the game in normal form* $\mathcal{F}$*:* $\mathsf{val}[\mathcal{F}] := \mathsf{val}[\mathcal{F}](\mathsf{A}) = \mathsf{val}[\mathcal{F}](\mathsf{B})$*.*

Although the values of a game in normal form is not necessarily the same for both players, it always holds that the value for Player A is at most the value for Player B, as stated in Lemma 1.9 below.

Lemma 1.9 (Proof Subsection 1.5.3)**.** *Consider a game in normal form $\mathcal{F}$. We have* $\mathsf{val}[\mathcal{F}](\mathsf{A}) \leq \mathsf{val}[\mathcal{F}](\mathsf{B})$.

Since the outcome of a game in normal form is an expected value, it is in fact linear in the valuation of the outcomes. Inequality between valuations of outcomes also propagates to value of games in normal form.

Lemma 1.10 (Proof Subsection 1.5.4)**.** *Consider a game form $\mathcal{F}$ and a strategy per player $(\sigma_\mathsf{A}, \sigma_\mathsf{B}) \in \Sigma_\mathsf{A} \times \Sigma_\mathsf{B}$. Consider also $(v_n)_{n \in \mathbb{N}} \in ([0,1]^\mathsf{O})^\mathbb{N}$, $(\lambda_n) \in ([0,1])^\mathbb{N}$ such that $\sum_{n \in \mathbb{N}} \lambda_n \cdot v_n : \mathsf{O} \to [0,1]$. Then:*

$$\sum_{n \in \mathbb{N}} \lambda_n \cdot \mathsf{out}[\langle \mathcal{F}, v_n \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) = \mathsf{out}[\langle \mathcal{F}, \sum_{n \in \mathbb{N}} \lambda_n \cdot v_n \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B})$$

Consider any two valuations $v, v' : \mathsf{O} \to [0,1]$, $\lambda > 0$ and $x \in \mathbb{R}$ such that $\lambda \cdot v + x \leq v'$. We have:

$$\lambda \cdot \mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) + x \leq \mathsf{out}[\langle \mathcal{F}, v' \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B})$$

and, for all $s \in \{\mathsf{A}, \mathsf{B}, \sigma_\mathsf{A}, \sigma_\mathsf{B}\}$:

$$\lambda \cdot \mathsf{val}[\langle \mathcal{F}, v \rangle](s) + x \leq \mathsf{val}[\langle \mathcal{F}, v' \rangle](s)$$

If in addition we have $\lambda \cdot v + x : \mathsf{O} \to [0,1]$, then:

$$\lambda \cdot \mathsf{val}[\langle \mathcal{F}, v \rangle](s) + x = \mathsf{val}[\langle \mathcal{F}, \lambda \cdot v + x \rangle](s)$$

In the following, we will be especially interested in the game forms such that all games in normal forms that can be induced from them have a value. Such game forms are called valuable, this is defined below in Definition 1.15.

Definition 1.15 (Valuable game form)**.** *Consider a game form $\mathcal{F}$. It is* valuable *if for all valuations of the outcomes $v : \mathsf{O} \to [0,1]$, the game in normal form $\langle \mathcal{F}, v \rangle$ has a value.*

Furthermore, for a player $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$ and a subset of Player-C strategies $\mathsf{S}_\mathsf{C} \subseteq \Sigma_\mathsf{C}(\mathcal{F})$, the game form $\mathcal{F}$ is supremized *(resp.* maximized*) by S_C w.r.t. Player C if for all valuations of the outcomes $v : \mathsf{O} \to [0,1]$, for all $\varepsilon > 0$, there is a Player-C strategy $\sigma_\mathsf{C} \in \mathsf{S}_\mathsf{C}$ that is ε-optimal (resp. optimal) in the game in normal form $\langle \mathcal{F}, v \rangle$. A game form is* maximizable *w.r.t. Player C if it maximized by a set of strategies w.r.t. Player C.*

The above-defined notion of valuable game form is crucial. It will in particular appear in Chapter 2 in the statement of (the new version of) Blackwell determinacy.

In other words, the fact that a set of strategies S_C supremizes a game form $\mathcal{F}$ w.r.t. Player C means that the values of the games in normal form that can be induced from $\mathcal{F}$ are identical if Player C restrict herself only to strategies in S_C. In particular, any valuable game form $\mathcal{F}$ is supremized by $\Sigma_\mathsf{C}(\mathcal{F})$ w.r.t. to Player C. We make a straightforward observation below: in any game form $\mathcal{F}$, a finite set of strategies supremizes $\mathcal{F}$ if and only if it maximizes $\mathcal{F}$.

Observation 1.1 (Proof Subsection 1.5.5). *Consider a game form $\mathcal{F}$, a player $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$ and a finite set of Player-C strategies $\mathsf{S}_\mathsf{C} \subseteq \Sigma_\mathsf{C}(\mathcal{F})$. The set S_C supremizes $\mathcal{F}$ w.r.t. Player C if and only if it maximizes $\mathcal{F}$ w.r.t. Player C.*

Let us now focus on standard game forms. It is a well-known result that all standard deterministic finite game forms (recall, game form with finitely many actions for both players) are valuable and maximizable. This comes from Von Neuman's minimax theorem [44] and it is stated below.

Theorem 1.11. *All standard deterministic finite game forms are valuable and maximizable.*

Below in Proposition 1.12, we establish that this holds even if the set of actions available to one of the players is not finite (it can be arbitrary, even uncountable). However, note that the strategies we consider still have a countable support. It was already proved in an unpublished work [45], see also Sion's minimax theorem [46].

Proposition 1.12 (Proof Subsection 1.5.6). *Consider a standard deterministic game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ on a set of outcomes O. If Act_A or Act_B is finite then $\mathcal{F}$ is valuable and for any player $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$ with Act_C finite, the game form $\mathcal{F}$ is maximizable w.r.t. Player C.*

Let us now give a proof sketch below of Proposition 1.12, the full proof (that is quite lengthy due to technical details) can be found in Appendix 1.5.6.

Proof sketch. If both Act_A and Act_B are finite, we are in the scope of Theorem 1.11. Assume now that only Act_A is finite while Act_B is not, the other case being analogous. The proof is in two steps.

First, we show that for any valuation of the outcomes $v : \mathsf{O} \to [0, 1]$ taking finitely many values (i.e. such that $v[\mathsf{O}]$ is finite), the game in normal form $\langle \mathcal{F}, v \rangle$ has a value and Player A has an optimal strategy. This comes from the fact that since Player A has finitely many actions and the outcomes, valued by v, can take only finitely many values then there are in fact finitely many different actions for Player B. (Two Player-B actions being different if there is a Player-A action for which the corresponding outcomes, valued by v, are different. This is defined formally with an equivalence relation.) Hence, the game in normal form obtained can be seen as finite.

Second, given an arbitrary valuation of the outcomes $v : \mathsf{O} \to [0, 1]$, we approximate what happens in the game in normal $\langle \mathcal{F}, v \rangle$ with what happens

$$\begin{bmatrix} x & x & x & x & \dots \\ y & x & x & x & \dots \\ y & y & x & x & \dots \\ y & y & y & x & \dots \\ \dots & \dots & \dots & \dots & \dots \end{bmatrix}$$

Figure 1.2: A game form that is not valuable with the set of actions for both players being countable.

$$\begin{bmatrix} x & y & y & y & \dots \\ y & x & y & y & \dots \\ y & y & x & y & \dots \\ y & y & y & x & \dots \\ \dots & \dots & \dots & \dots & \dots \end{bmatrix}$$

Figure 1.3: A game form that is valuable but not maximizable w.r.t. any player.

in games in normal form $\langle \mathcal{F}, v_n \rangle$ where $v_n : \mathsf{O} \to [0,1]$ takes finitely many values and is closer and closer to v as $n \to \infty$. Specifically, by using the inequalities from Lemma 1.19, we can show that the limit of the values of the games $\langle \mathcal{F}, v_n \rangle$ is in fact equal to the value of the game $\langle \mathcal{F}, v \rangle$. Furthermore, we exhibit a Player-A optimal strategy in $\langle \mathcal{F}, v \rangle$ by considering the limit of Player-A strategies σ_n optimal in $\langle \mathcal{F}, v_n \rangle$. The existence of this limit is ensured by Theorem 1.1 (page 26), since Act_A is finite. □

Remark 1.2. *Proposition 1.12 is already known in the case where the non-finite set of actions is countable (for instance, this is indirectly mentioned in [12]). In that case, the result can be obtained by successively approximating what happens in the whole game in normal form by considering more and more Player-*B *actions, but always finitely many. However, this cannot be extended to the case of uncountably many Player-*B *actions.*

In some way, Proposition 1.12 is tight in the sense that, in standard game forms, as soon as one allows both players' actions set to be infinite, then it is possible to exhibit a game form that is not valuable. This can be witnessed with a game form with countably many actions for both players. Furthermore, there are also game forms that are valuable but that are not maximizable w.r.t. any player. Again, this can be witnessed with a game form with countably many actions for both players. In both cases, only two different outcomes are sufficient. We give such examples below in the next example.

Example 1.3 (Non valuable or maximizable game forms among standard game forms)**.** *The game form* $\mathcal{F}_1$ *represented in Figure 1.2 is equal to* $\mathcal{F}_1 := \langle \mathbb{N}, \mathbb{N}, \{x, y\}, \varrho_1 \rangle$ *where, for all* $(i,j) \in \mathbb{N}^2$*, we have* $\varrho_1(i,j) := x$ *if and only if* $i \leq j$*, otherwise* $\varrho_1(i,j) := y$*. Let us show that it is not valuable. Consider the valuation of the outcomes* $v : \{x, y\} \to [0,1]$ *with* $v(x) := 0$ *and* $v(y) := 1$ *and the induced game in normal form* $\mathcal{F}'_1 := \langle \mathcal{F}_1, v \rangle$*. This game in normal form does not have a value, as we have* $\mathsf{val}[\mathcal{F}'_1](\mathsf{A}) = 0$ *and* $\mathsf{val}[\mathcal{F}'_1](\mathsf{B}) = 1$*. Indeed,* *consider any Player-*A *strategy* $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathcal{F}_1)$. *For all* $j \in \mathbb{N}$, *we have:*

$$\mathsf{out}[\mathcal{F}_1'](\sigma_\mathsf{A}, j) = \sum_{i>j} \sigma_\mathsf{A}(i)$$

Since $\sum_{i\in\mathbb{N}} \sigma_\mathsf{A}(i) = 1$, *it follows that* $\mathsf{out}[\mathcal{F}_1'](\sigma_\mathsf{A}, j) \to_{j\to\infty} 0$. *Hence,* $\mathsf{val}[\mathcal{F}_1'](\sigma_\mathsf{A}) \leq \inf_{j\in\mathbb{N}} \mathsf{out}[\mathcal{F}_1'](\sigma_\mathsf{A}, j) = 0$. *As this holds for all Player-*A *strategies* $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathcal{F}_1)$, *we have* $\mathsf{val}[\mathcal{F}_1'](\mathsf{A}) = \sup_{\sigma_\mathsf{A}\in\Sigma_\mathsf{A}(\mathcal{F}_1)} \mathsf{val}[\mathcal{F}_1'](\sigma_\mathsf{A}) = 0$. *The arguments are similar to show that* $\mathsf{val}[\mathcal{F}_1'](\mathsf{B}) = 1$. *Note that this example of game form that is not valuable is a folk result.*

Consider now the game form $\mathcal{F}_2$ *from Figure 1.3. It is equal to* $\mathcal{F}_2 := \langle \mathbb{N}, \mathbb{N}, \{x, y\}, \varrho_2 \rangle$ *where, for all* $(i, j) \in \mathbb{N}^2$, *we have* $\varrho_2(i, j) := x$ *if and only if* $i = j$, *otherwise* $\varrho_2(i, j) := y$. *Let us show that this game form is valuable but not maximizable w.r.t. any player. Consider a valuation* $v : \{x, y\} \to [0, 1]$. *Let* $\mathcal{F}_2' := \langle \mathcal{F}_2, v \rangle$. *If* $v(x) = v(y)$, *straightforwardly* $\mathsf{val}[\mathcal{F}_2'](\mathsf{A}) = \mathsf{val}[\mathcal{F}_2'](\mathsf{B}) = v(x) = v(y)$. *Assume now that* $v(x) \neq v(y)$. *In that case, we claim that* $\mathsf{val}[\mathcal{F}_2'](\mathsf{A}) = \mathsf{val}[\mathcal{F}_2'](\mathsf{B}) = v(y)$. *Indeed, for all* $n \in \mathbb{N}$, *consider a Player-*A *strategy* σ_A^n *playing uniformly over the* $(n+1)$*-first integers such that for all* $i \in [\![0, n]\!]$ *we have* $\sigma_\mathsf{A}^n(i) := \frac{1}{n+1}$. *In fact, for all Player-*B *strategies* $\sigma_\mathsf{B} \in \Sigma_\mathsf{B}(\mathcal{F})$, *we have:*

$$|\mathsf{out}[\mathcal{F}_2'](\sigma_\mathsf{A}^n, \sigma_\mathsf{B}) - v(y)| \leq \frac{|v(x) - v(y)|}{n+1}$$

Hence, $\mathsf{val}[\mathcal{F}_2'](\mathsf{A}) \geq v(y)$. *By symmetry of the game in normal form* $\mathcal{F}_2'$, *it follows that we also have* $\mathsf{val}[\mathcal{F}_2'](\mathsf{B}) \leq v(y)$. *Since, in any case, by Lemma 1.9,* $\mathsf{val}[\mathcal{F}_2'](\mathsf{A}) \leq \mathsf{val}[\mathcal{F}_2'](\mathsf{B})$, *it follows that* $v(y) \leq \mathsf{val}[\mathcal{F}_2'](\mathsf{A}) \leq \mathsf{val}[\mathcal{F}_2'](\mathsf{B}) \leq v(y)$. *Hence,* $\mathsf{val}[\mathcal{F}_2'](\mathsf{A}) = \mathsf{val}[\mathcal{F}_2'](\mathsf{B}) = v(y)$. *However, if* $v(y) > v(x)$, *Player* A *does not have any optimal strategy in* $\mathcal{F}_2'$ *since she cannot avoid a positive probability of* x. *The same issue arises for Player* B *when* $v(y) < v(x)$.

Note that it is simpler to come up with not valuable game forms or not maximizable game forms with non-standard game forms, we give examples below.

Example 1.4 (Non valuable or maximizable game forms among non-standard game forms)**.** *Consider the standard game form on the left of Figure 1.1. If we consider the non-standard game forms where both players can only play deterministic probability distributions, then this game form is not valuable. Indeed, consider the valuation* $v : \{x, y\} \to [0, 1]$ *such that* $v(x) := 1$ *and* $v(y) := 0$. *Then, any Player-*A *strategy has value 0 since Player* B *can choose to see* y *with probability 1 and symmetrically, any Player-*B *strategy has value 1 since Player* A *can choose to see* x *with probability 1. Furthermore, the game form we have described in Example 1.2 is not maximizable w.r.t. Player* A. *This is witnessed by any valuation mapping* y *to a greater value that* x.

Finally, we transfer results on a specific game form $\mathcal{F}$ to any game form that can be obtained from $\mathcal{F}$ by mapping every outcome to a distribution over another set of outcomes. We define this change formally and state the corresponding lemma below.

Definition 1.16 (Map from outcome to distribution over outcomes). *Consider a non-empty set of outcomes* O *and a game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$. *Consider a non-empty set of outcomes* O' *and a map* $d : \mathsf{O} \to \mathcal{D}(\mathsf{O}')$. *We denote by* $\mathcal{F}^d \in \mathsf{Form}(\mathsf{O}')$ *the game form* $\mathcal{F}^d := \langle \Sigma_\mathsf{A}, \Sigma_\mathsf{B}, \mathsf{O}', \mathbb{E}_{\varrho(\cdot,\cdot)}(d) \rangle$.

Lemma 1.13 (Proof 1.5.7). *Consider a non-empty sets of outcomes* O *and a game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$. *Assume that* $\mathcal{F}$ *is valuable (resp. maximizable w.r.t. to Player* $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$*). Then, for all non-empty sets of outcomes* O' *and map* $d : \mathsf{O} \to \mathcal{D}(\mathsf{O}')$, *so is the game form* $\mathcal{F}^d$.

As a corollary of Proposition 1.12 and Lemma 1.13, we obtain a statement very close to Proposition 1.12 except that we dropped the deterministic assumption:

Corollary 1.14 (Proof 1.5.8). *Consider a standard game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ *on a set of outcomes* O. *If* Act_A *or* Act_B *is finite then* $\mathcal{F}$ *is valuable and for any player* $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$ *with* Act_C *finite, the game form* $\mathcal{F}$ *is maximizable w.r.t. Player* C.

For the remainder of this dissertation, strategies in game forms will be called GF-strategies in order not to confuse them with strategies in concurrent games on graphs. Furthermore, they will usually be denoted by the letter σ, typically σ_A (resp. σ_B) for a Player-A (resp. Player-B) strategy. (As opposed to strategies in concurrent games, that we will usually denoted by s.)

1.4 Concurrent arenas and games

Before defining concurrent games, we need to define the notion of concurrent arenas. To gain intuition on what these are, take a look at the (standard) arena depicted in Figure 1.4. Consider for instance the leftmost state q_0. From there, two players — that we still call Player A and Player B — are going to interact. The result of their interaction will be a (distribution over) successor states. Interacting at state q_0 in fact means playing in the game form depicted in that state, where the outcomes are states of the arena. In the arena we have depicted, all game forms are standard. Hence, this means that Player A chooses a distribution over the rows and concurrently, Player B chooses a distribution over the columns. A new state is then reached with some probability, in that case either q_1 or q_2, and then the process repeats itself indefinitely thus creating an infinite path (i.e. an infinite sequence of states). Finally, in addition, we consider colors (i.e. labels) over the states which will be used to define the payoff function or the winning condition. Note that they do not relate at all with the colors appearing in the game forms in Figure 1.4.

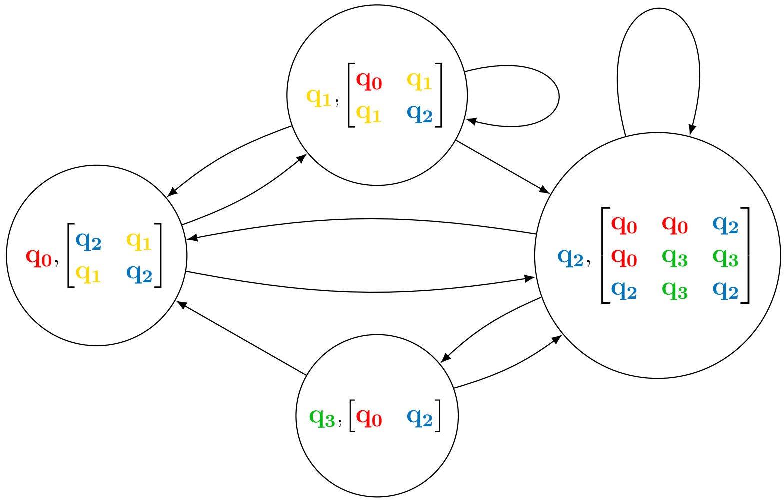

Figure 1.4: A standard deterministic concurrent arena. The colors are there only to facilitate the readability of the arena.

Definition 1.17 (Concurrent arena). *A concurrent arena $\mathcal{C}$ is a tuple $\mathcal{C} = \langle Q, \mathsf{F}, \mathsf{K}, \mathsf{col} \rangle$ where Q is a non-empty countable set of states, $\mathsf{F} : Q \to \mathsf{Form}(Q)$ maps each state to its induced game form, which describes the interaction of the players at that state, K is a non-empty set of colors and $\mathsf{col} : Q \to \mathsf{K}$ is the coloring function that maps each state to a color. For all $q \in Q$, the game form $\mathsf{F}(q)$ is called the* local interaction *at state q.*

In the following, when the set of states considered is clear from context, sequences of states will be called paths.

In the following, we will consider a slightly more general formalism by considering stopping states with output value, i.e. states that, when visited, immediately stop the game and induce a specific value in $[0, 1]$. This is formally defined below in Definition 1.18.

Definition 1.18 (Stopping states). *Consider a concurrent arena $\mathcal{C}$. A* stopping state *$q \in Q$ is a state such that, when reached, the game stops and outputs a value $\mathsf{val}(q) \in [0, 1]$. This will be formalized in Definition 1.30 below. The local interactions at stopping states are trivial and they are self-looping. The coloring function col need not be defined on stopping states.*

We denote by $Q_{\mathsf{s}} \subseteq Q$ the set of stopping states and by $Q_{\mathsf{ns}} := Q \setminus Q_{\mathsf{s}}$ the set of states that are non-stopping.

For the remainder of this dissertation, the notation $\mathcal{C}$ will refer to the arena $\langle Q, \mathsf{F}, \mathsf{K}, \mathsf{col} \rangle$, unless otherwise stated. Furthermore, in such an arena, for all $q \in Q$, the set of Player-A GF-strategies available at state q will be referred to as Σ^q_{A} and similarly for Player B. In addition, if the arena $\mathcal{C}$ is standard, the set of Player-A actions available at state q will be denoted $\mathsf{Act}^q_{\mathsf{A}}$ and similarly for Player B.

We would like to mention several relevant special cases of concurrent arenas: standard arenas, in which all local interactions are standard. Arbitrary arenas will refer to non-necessarily standard arenas. Furthermore, we consider deterministic arenas — such as the one depicted in Figure 1.4 — turn-based arenas, i.e. arenas where each game form is turn-based (such as the game form at state q_3 in the arena of Figure 1.4). We will also consider finite-state arenas, i.e. with finitely many states. Finally, we will consider standard finite arenas, that is finite-state arenas with standard finite local interactions. This is defined formally below in Definition 1.19.

Definition 1.19. *Consider a concurrent arena $\mathcal{C}$. It is:*

- standard*: if every local interaction is standard;*
- arbitrary*: it stands for non-necessarily standard;*
- deterministic*: if it is standard and all of its local interactions are deterministic;*
- turn-based*: for all $q \in Q$, the game form $\mathsf{F}(q)$ is turn-based (in particular, all local interactions are standard);*
- finite-state*: if there are finitely many states;*
- finite*: if it is finite-state and, if the game is standard, we additionally require that all standard local interactions are finite (i.e. both players have finitely many actions).*

Remark 1.3. *One can see that turn-based arenas are defined only with standard game forms (since turn-based game forms are by definition standard, recall Definition 1.11). We choose to do this instead of defining turn-based games with non-standard game forms because turn-based games are widely studied and usually standard in the literature. In the following, we will transfer already existing and prove new results on them. Hence, we do not want any confusion as to the object we consider.*

Below, we consider the notion of valuable (resp. maximizable) arena, that is an arena where all local interactions are valuable (resp. maximizable).

Definition 1.20 (Local interactions and valuable arena)**.** *Consider a concurrent arena $\mathcal{C}$. If, for all $q \in Q$, the game form $\mathsf{F}(q)$ is valuable, then the arena $\mathcal{C}$ is said to be* valuable.

On the other hand, if for all $q \in Q$, the game form $\mathsf{F}(q)$ is supremized w.r.t. Player A by a set $S^q_{\mathsf{A}} \subseteq \Sigma^q_{\mathsf{A}}$ of GF-strategies, the arena $\mathcal{C}$ is said to be supremized *by the collection $(S^q_{\mathsf{A}})_{q \in Q}$. This is similar for Player B. In addition,* *if for all $q \in Q$, the game form $\mathsf{F}(q)$ is maximizable w.r.t. to any player, then the arena $\mathcal{C}$ is said to be* maximizable *w.r.t. that same player. If this holds for both players, we will simply say that it is maximizable.*

Note that, in this dissertation, all the properties we have defined on concurrent arenas will be used to refer to concurrent games whose underlying arenas satisfy these properties.

Finally, we define the notion of concurrent arenas built from a set of game forms, that is such that all local interactions in that games are obtained from a game form in that set. We first define below the notion of game forms obtained from another game form.

Definition 1.21 (Game forms obtained from another game form)**.** *Consider a set of outcomes O and a game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$. We say that a game form $\mathcal{F}' \in \mathsf{Form}(\mathsf{O}')$ on another set of outcomes O' is* obtained *from $\mathcal{F}$ if there is some map $m : \mathsf{O} \to \mathsf{O}'$ such that $\mathcal{F}' = \mathcal{F}^m := \langle \mathsf{Act}_\mathsf{A}, \mathsf{Act}_\mathsf{B}, \mathsf{O}', \mathbb{E}_m(\varrho) \rangle$.*

We can now define the notion of concurrent arena built from a set of game forms.

Definition 1.22 (Arena built from a set of Game Forms)**.** *Consider a set of game form E. We say that an arena $\mathcal{C}$ is* built *from E if all local interactions in $\mathcal{C}$ are obtained from a game form in E.*

This notion will be particularly useful in Part III since, informally, the goal of this part is to define subsets of game forms such that all the games built from them behave well.

1.4.1 . Drawing concurrent arenas

Below, we make a remark about how Definition 1.17 of concurrent arenas above — which makes use of the notion of game form to describe the interactions of the players at each state — may affect how we draw concurrent arenas as opposed to how they have been drawn in the literature so far.

Remark 1.4. *In other papers studying concurrent arenas (for instance, [47, 48, 34, 31, 32, 49, 50]), the formalism used to describe them is different. Especially, since the notion of game form is not apparent, not defined, the interaction of the players at each state is not described with a game form. Instead, both players have a set of available actions and there is a transition function mapping each state and pair of actions at that state to a (distribution over) successors states. One can realize that this exactly corresponds to our formalism with standard game forms, only without having introduced the notion of game form. However, this way of defining concurrent arenas probably has an impact on how these arenas are drawn. Indeed, since the notion of game form is not used, the interactions of the players are not drawn with bi-dimensional tables — as is done in Figure 1.4. This leads to a representation of concurrent arenas where the interaction of the players is described by the pairs of actions* *leading to the different states of the arena (see for instance [47, Figure 3]) labeling the edges of the arena. We believe — but that is obviously debatable — that drawing concurrent arenas with game forms represented as bi-dimensional tables increases the readability of these arenas. This allows to consider local interactions as first-class citizens in concurrent games. Note that not all of our results explicitly use the notion of game forms, but many do and in any case game forms are always an underlying object useful to have at hand.*

1.4.2 . Concurrent games

A concurrent game is obtained from a concurrent arena by specifying what Player A and Player B want to achieve in these arenas. This is done by adding a payoff function mapping each infinite sequence of colors to a value in $[0, 1]$. In this dissertation, we will mostly focus on the special case of win/lose objectives, that is to payoff functions taking their values in $\{0, 1\}$. This is defined below.

Definition 1.23 (Concurrent game). *A* concurrent game *is a pair $\mathcal{G} = \langle \mathcal{C}, f \rangle$ where $\mathcal{C}$ is a concurrent arena and $f : \mathsf{K}^\omega \to [0, 1]$ is a measurable payoff function.*

When $f[\mathsf{K}^\omega] \subseteq \{0, 1\}$ the game $\mathcal{G}$ is called win/lose*. Win/lose games are defined by $\mathcal{G} = \langle \mathcal{C}, W \rangle$ where the measurable set $W := f^{-1}[\{1\}] \subseteq \mathsf{K}^\omega$ is called the* objective *for Player A. Indeed, an infinite sequence of colors $\rho \in W$ is winning for Player A (and losing for Player B) whereas an infinite path $\rho \in \mathsf{K}^\omega \setminus W$ is winning for Player B (and losing for Player A), hence the win/lose terminology. We denote by $W^Q \subseteq Q^\omega$ the measurable set $W^Q := (\mathsf{col}^\omega)^{-1}[W]$.*

For the remainder of this dissertation, the notation $\mathcal{G}$ will refer to the game $\langle \mathcal{C}, f \rangle$, unless otherwise stated. Furthermore, all the payoff functions we consider are measurable and into $[0, 1]$.

We define in Definition 1.24 a special kind of payoff function that will be of particular interest for us in Chapter 3: prefix-independent payoff functions. Informally, these are payoff functions whose values do not depend on any finite prefix. This is defined formally below in Definition 1.24.

Definition 1.24 (Prefix-independent games). *Consider a set of colors K and a payoff function $f : \mathsf{K}^\omega \to [0, 1]$. It is* prefix-independent *(PI for short) if, for all $\rho \in \mathsf{K}^\omega$ and $\pi \in \mathsf{K}^*$, we have $f(\rho) = f(\pi \cdot \rho)$. An objective $W \subseteq \mathsf{K}^\omega$ is* prefix-independent *if the corresponding payoff function is. That is, for all $\rho \in \mathsf{K}^\omega$ and $\pi \in \mathsf{K}^*$, we have $\rho \in W \Leftrightarrow \pi \cdot \rho \in W$.*

We say that a concurrent game $\mathcal{G} = \langle \mathcal{C}, f \rangle$ is prefix-independent if its payoff function is.

Below in Definition 1.25, we define several kinds of win/lose objective of interest for us, namely parity objectives. They are presented from Player A's point of view, the objective for Player B would be the complement (for instance, when Player A has a Büchi objective, Player B has a co-Büchi objective).

Definition 1.25 (Parity objectives)**.**

- *With a* parity objective *the colors are non-negative integers, the goal of Player* A *is to ensure that the maximum of the colors seen infinitely often is even. To ensure that a maximum always exists, we consider a finite subset of non-negative integers:* $\mathsf{K} := [\![m, n]\!] \subseteq \mathbb{N}$ *for some* $m \leq n \in \mathbb{N}$. *For all* $\rho \in \mathsf{K}^\omega$, *we let* $\mathsf{InfOft}(\rho) := \{k \in \mathsf{K} \mid \forall i \in \mathbb{N},\ \exists j \geq i,\ \rho_j = k\} \subseteq \mathsf{K}$ *be the set of colors seen infinitely often in* ρ. *Then, the parity objective with colors* $\mathsf{Parity}_\mathsf{K} \subseteq \mathsf{K}^\omega$ *is equal to* $\mathsf{Parity}_\mathsf{K} := \{\rho \in \mathsf{K}^\omega \mid \max \mathsf{InfOft}(\rho)$ *is even* $\}$.

- *With a* Büchi objective, *the goal of Player* A *is that a given set of states* $S \subseteq Q$ *is seen infinitely often. There are two distinct colors, say* $\mathsf{K} := \{1, 2\}$, *and the Büchi objective is equal to* $\mathsf{Buchi} := \{\rho \in \mathsf{K}^\omega \mid \forall i \in \mathbb{N},\ \exists j \geq i,\ \rho_j = 2\}$ *(hence, S is exactly the 2-colored states). Note that this exactly corresponds to the parity objective* $\mathsf{Parity}_{[\![1,2]\!]}$.

- *With a* co-Büchi objective, *the goal of Player* A *is that a given set of states* $S \subseteq Q$ *is seen only finitely often. There are two distinct colors, say* $\mathsf{K} := \{0, 1\}$, *and the co-Büchi objective is equal to* $\mathsf{coBuchi} := \{\rho \in \mathsf{K}^\omega \mid \exists i \in \mathbb{N},\ \forall j \geq i,\ \rho_j = 0\}$ *(hence S is exactly the 1-colored states). Note that this exactly corresponds to the parity objective* $\mathsf{Parity}_{[\![0,1]\!]}$.

- *With a* Reachability objective, *the goal of Player* A *is that a given set of states* $S \subseteq Q$ *is seen once. There are two distinct colors, say* $\mathsf{K} := \{1, 2\}$, *and the Reachability objective is equal to* $\mathsf{Reach} := \{\rho \in \mathsf{K}^\omega \mid \exists i \in \mathbb{N},\ \rho_i = 2\}$ *(hence S is exactly the 2-colored states). This does not correspond to a parity objective in general. However, this exactly corresponds to the Büchi objective in arenas where all states q in S are self-looping sinks i.e. the only outgoing edge of q leads to q: that is, seeing S once means seeing it infinitely often.*

- *With a* Safety objective, *the goal of Player* A *is that a given set of states* $S \subseteq Q$ *is avoided. There are two distinct colors, say* $\mathsf{K} := \{0, 1\}$, *and the Safety objective is equal to* $\mathsf{Safe} := \{\rho \in \mathsf{K}^\omega \mid \forall i \in \mathbb{N},\ \rho_i = 0\}$ *(hence S is exactly the 1-colored states). This does not correspond to a parity objective in general. However, this exactly corresponds to the co-Büchi objective in arenas where all states q in S are self-looping sinks.*

Note that the parity (therefore also Büchi and co-Büchi) objectives are prefix-independent, but the reachability and safety are not. However, in this dissertation, we will always consider reachability and safety objectives as special cases of parity objectives, for the reason described above.

1.4.3 . Strategies, induced stochastic trees and values

In concurrent arenas, strategies describe how the players play. More specifically, strategies are functions that map the history of the game (i.e. the finite sequence of states visited so far) to a GF-strategy in the game form corresponding to the current state of the game. This is formally defined below in Definition 1.26.

Definition 1.26 (Strategies). *Consider a concurrent arena $\mathcal{C}$. A* strategy *for Player A is a function $\mathsf{s_A} : Q^+ \to \cup_{q \in Q} \Sigma_{\mathsf{A}}^q$ such that, for all $\pi \in Q^+$, we have $\mathsf{s_A}(\pi) \in \Sigma_{\mathsf{A}}^{\pi_{\mathsf{lt}}}$. We denote by $\mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ the set of all strategies in the arena $\mathcal{C}$ for Player A. A strategy $\mathsf{s_A}$ is* deterministic *if, for all $\rho \in Q^+$, the GF-strategy $\mathsf{s_A}(\rho)$ is deterministic (recall that this is only defined if the game form $\mathsf{F}(\rho_{\mathsf{lt}})$ is standard). The definitions are similar for Player B.*

A strategy is generated by a collection indexed by Q of sets of GF-strategies if it always plays a GF-strategy among one of the set of this collection. We define formally this notion below.

Definition 1.27 (Strategies generated by sets of GF-strategies). *Consider a concurrent arena $\mathcal{C}$ and, for each state $q \in Q$ consider a subset of GF-strategies $S_q^{\mathsf{A}} \subseteq \Sigma_{\mathsf{A}}^q$. We say that a Player-$\mathsf{A}$ strategy $\mathsf{s_A}$ is* generated *by the collection $(S_{\mathsf{A}}^q)_{q \in Q}$ if, for all $\rho \in Q^+$, we have $\mathsf{s_A}(\rho) \in S_{\mathsf{A}}^{\rho_{\mathsf{lt}}}$. The definition is similar for Player B.*

The outcome of a game, given a strategy per Player, is a probability measure over infinite paths. To formalize this, we first define below the probability to go from a state q to a state q' given two GF-strategies in $\mathsf{F}(q)$.

Definition 1.28 (Probability transition given two strategies). *Consider a concurrent arena $\mathcal{C}$, a state $q \in Q$ and two strategies $(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \in \Sigma_{\mathsf{A}}^q \times \Sigma_{\mathsf{B}}^q$. Let $q' \in Q$. The probability to go from q to q' if the players plays, in q, σ_{A} and σ_{B}, denoted $\mathbb{P}_{\mathcal{C}}^{\sigma_{\mathsf{A}},\sigma_{\mathsf{B}}}(q, q')$, is equal to (recalling the last sentence of the first paragraph Subsection 1.2.1):*

$$\mathbb{P}_{\mathcal{C}}^{\sigma_{\mathsf{A}},\sigma_{\mathsf{B}}}(q, q') := \mathsf{out}[\langle \mathsf{F}(q), q' \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})$$

Below, we define, given a strategy per player, the probability of finite paths. Then, the definition of stochastic tree induced by a pair of strategies follows.

Definition 1.29 (Probability distribution given two strategies). *Consider a concurrent arena $\mathcal{C}$ and two arbitrary strategies $(\mathsf{s_A}, \mathsf{s_B}) \in \mathsf{S}_{\mathcal{C}}^{\mathsf{A}} \times \mathsf{S}_{\mathcal{C}}^{\mathsf{B}}$. We denote by $\mathbb{P}_{\mathcal{C}}^{\mathsf{s_A},\mathsf{s_B}} : Q^+ \to \mathcal{D}(Q)$ the function giving the probability distribution over the next state of the arena given the sequence of states already seen. That is, for all finite path $\pi \in Q^+$ and $q \in Q$, we have:*

$$\mathbb{P}_{\mathcal{C}}^{\mathsf{s_A},\mathsf{s_B}}(\pi)[q] := \mathbb{P}_{\mathcal{C}}^{\mathsf{s_A}(\pi),\mathsf{s_B}(\pi)}(\pi_{\mathsf{lt}}, q)$$

The stochastic tree $\mathcal{T}_{\mathcal{C}}^{\mathsf{s_A},\mathsf{s_B}}$ induced *by the pair of strategies $(\mathsf{s_A}, \mathsf{s_B})$ is then equal to $\mathcal{T}_{\mathcal{C}}^{\mathsf{s_A},\mathsf{s_B}} := \langle Q, \mathbb{P}_{\mathcal{C}}^{\mathsf{s_A},\mathsf{s_B}} \rangle$.*

Given two strategies and the stochastic tree induced by them, we have, by using Definition 1.7, the expected value of the (measurable) payoff function. Analogously to what happens in game forms (recall Definition 1.14), in a concurrent game on graph, Player A tries to maximize this payoff function whereas Player B tries to minimize it. The value of a Player-A strategy s_A is the best that this strategy s_A can achieve against all Player-B strategies. Therefore, it is equal to the infimum over all Player-B strategies s_B of the expected value of the payoff function given that pair of strategies $(\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B})$. Then, the Player-$\mathsf{A}$ value of the game is equal to the supremum of the values of her strategies. Before giving the formal definitions in Definition 1.31, we define exactly the payoff function we consider, that takes into account the stopping states of the arena.

Definition 1.30 (Payoff function on the sequences of states)**.** *Consider a concurrent game $\mathcal{G} = \langle \mathcal{C}, f \rangle$. We denote by $f_\mathcal{C} : Q^\omega \to [0, 1]$ the function such that, for all $\rho \in Q^\omega$:*

$$f_\mathcal{C}(\rho) := \begin{cases} f \circ \mathsf{col}^\omega(\rho) & \text{if } \rho \in (Q_\mathsf{ns})^\omega \\ \mathsf{val}(q) & \text{if } \rho \in (Q_\mathsf{ns})^* \cdot q \cdot Q^\omega, \quad \text{for } q \in Q_\mathsf{s} \end{cases}$$

Interestingly, such a function is measurable.

Proposition 1.15 (Proof 1.5.9)**.** *For any concurrent game $\mathcal{G} = \langle \mathcal{C}, f \rangle$ and for all $\rho \in Q^*$, the residual function $(f_\mathcal{C})^\rho : Q^\omega \to [0, 1]$ is measurable.*

We can now define formally the value of a concurrent game.

Definition 1.31 (Value of strategies and of the game)**.** *Let $\mathcal{G} = \langle \mathcal{C}, f \rangle$ be a concurrent game and $\mathsf{s}_\mathsf{A} \in \mathsf{S}^\mathsf{A}_\mathcal{C}$ be a Player-A strategy. The function $\chi_\mathcal{G}[\mathsf{s}_\mathsf{A}] : Q \to [0, 1]$ mapping each state to the value of the strategy s_A from that state is such that, for all $q_0 \in Q$, we have:*

$$\chi_\mathcal{G}[\mathsf{s}_\mathsf{A}](q_0) := \inf_{\mathsf{s}_\mathsf{B} \in \mathsf{S}^\mathsf{B}_\mathcal{C}} \mathbb{E}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}}_{\mathcal{C}, q_0}[(f_\mathcal{C})^{q_0}]$$

The function $\chi_\mathcal{G}[\mathsf{A}] : Q \to [0, 1]$ mapping each state to the value for Player A from that state is such that, for all $q_0 \in Q$, we have:

$$\chi_\mathcal{G}[\mathsf{A}](q_0) := \sup_{\mathsf{s}_\mathsf{A} \in \mathsf{S}^\mathsf{A}_\mathcal{C}} \chi_\mathcal{G}[\mathsf{s}_\mathsf{A}](q_0)$$

The vector $\chi_\mathcal{G}[\mathsf{B}] : Q \to [0, 1]$ giving the value of the game for Player B is defined symmetrically. When $\chi_\mathcal{G}[\mathsf{A}] = \chi_\mathcal{G}[\mathsf{B}]$, this defines the value of the game*: $\chi_\mathcal{G} := \chi_\mathcal{G}[\mathsf{A}] = \chi_\mathcal{G}[\mathsf{B}]$.*

For all states $q \in Q$, a Player-A strategy s_A such that (resp. for some positive $\varepsilon > 0$), we have $\chi_\mathcal{G}[\mathsf{A}](q) = \chi_\mathcal{G}[\mathsf{s}_\mathsf{A}](q)$ (resp. $\chi_\mathcal{G}[\mathsf{A}](q) \leq \chi_\mathcal{G}[\mathsf{s}_\mathsf{A}](q) + \varepsilon$) is optimal *(resp. ε-*optimal*) from the state q. When this holds from all states $q \in Q$, the strategy s_A is simply said to be* optimal *(resp. ε-*optimal*). This is symmetrical for Player B.*

Observation 1.2. *In all concurrent games $\mathcal{G}$, we have $\chi_{\mathcal{G}}[\mathsf{A}] \leq \chi_{\mathcal{G}}[\mathsf{B}]$. Furthermore, by definition, for all $\varepsilon > 0$, both players have ε-optimal strategies.*

Below, we state that, for all prefix-independent payoff functions, replacing states with stopping states of the same value (w.r.t. either of the player) does not change the value of any state (w.r.t. the same player).

Lemma 1.16 (Proof Subsection 1.5.10). *Consider an arbitrary prefix-independent concurrent game $\mathcal{G}$, a subset of states $S \subseteq Q$ and a Player $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$. We denote by $\mathcal{G}^{S,\mathsf{C}}$ the game where all states $q \in S$ are stopping states with $\mathsf{val}(q) \leftarrow \chi_{\mathcal{G}}[\mathsf{C}](q)$. Then, the Player-$\mathsf{C}$ values of all states are the same in the games $\mathcal{G}$ and $\mathcal{G}^{S,\mathsf{C}}$: $\chi_{\mathcal{G}^{S,\mathsf{C}}}[\mathsf{C}] = \chi_{\mathcal{G}}[\mathsf{C}]$. Hence, if the game $\mathcal{G}$ has a value, so has the game $\mathcal{G}^{S,\mathsf{C}}$.*

Furthermore, if Player C has an optimal strategy in the game $\mathcal{G}$, then she also has one in the game $\mathcal{G}^{S,\mathsf{C}}$.

Let us introduce below a notation for the set of values occurring in a game. Furthermore, for all values $u \in [0, 1]$, we also consider the set of states of value u that we call a value slice.

Definition 1.32 (Set of values, value slice). *Consider a PI concurrent game $\mathcal{G}$. We let $V_{\mathsf{A}}^{\mathcal{G}} := \chi_{\mathcal{G}}[\mathsf{A}][Q] \subseteq [0, 1]$ be the set of Player-A values occurring in the game and, for all $u \in V_{\mathsf{A}}$, we let $Q_u^{\mathsf{A}} := (\chi_{\mathcal{G}}[\mathsf{A}])^{-1}[\{u\}]$ be the u*-value slice, *i.e. the set of states whose Player-A values are equal to u. The notation is analogous for Player B. Furthermore, we omit the notation for the player if they are the same for both players.*

Let us focus on the special case of Player-A strategies of value 1 (symmetrically, we could focus on Player-B strategies of value 0). Such a Player-A strategy is said to be almost-surely winning since, regardless of Player-B strategy and almost surely, the produced infinite path has value 1 w.r.t. the payoff function. When this happens surely, and in a win/lose game, such a strategy is said to be winning, as defined below in Definition 1.33.

Definition 1.33 (Compatible paths, Winning strategies). *Consider a concurrent game $\mathcal{G}$ without stopping states, a Player-A strategy $\mathsf{s_A}$ and a state $q \in Q$. A (finite or infinite) path $\rho \in Q^{\uparrow}$ is* compatible *with $\mathsf{s_A}$ from q if there is a Player-B strategy $\mathsf{s_B}$ such that, for all $i < |\rho|$, we have:*

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}(\rho_{\leq i}) > 0$$

We denote by $\mathsf{CP}_{\mathcal{C},q}(\mathsf{s_A}) \subseteq Q^\omega$ the set of infinite paths compatible with $\mathsf{s_A}$ from $q \in Q$.

If the game $\mathcal{G} = \langle \mathcal{C}, W \rangle$ is win/lose, we say that the strategy $\mathsf{s_A}$ is winning *from q if $\mathsf{CP}_{\mathcal{C},q}(\mathsf{s_A}) \subseteq W^Q$. The definition is symmetrical for Player B.*

1.4.4 . Positional and finite-memory strategies

A desirable property in concurrent games is that optimal or winning strategies exist, and that they are as simple as possible. In this context, simpler strategies are the strategies that can be implemented with a finite automaton with fewer states. To define how strategies can be implemented with a finite automaton, we first need to introduce the notion of memory skeleton (see, for instance, [51]), informally a finite automaton taking as input a finite sequence of colors. This is formally defined below.

Definition 1.34 (Memory skeleton)**.** *For a set of colors* K*, a finite memory skeleton on* K *is a triple* $\mathsf{M} = \langle M, m_{\text{init}}, \mu \rangle$*, where* M *is a non-empty finite set called the memory,* $m_{\text{init}} \in M$ *is the initial state of the memory and* $\mu : M \times \mathsf{K} \to M$ *is the update function. Note that the update function* μ *can be extended inductively into a function* $\mu^* : M \times \mathsf{K}^* \to M$ *in the following way: for all* $m \in M$*,* $\mu^*(m, \epsilon) := m$ *and for all* $\rho \cdot k \in \mathsf{K}^+$*,* $\mu^*(m, \rho \cdot k) := \mu(\mu^*(m, \rho), k)$*.*

Remark 1.5. *When considering the memory skeleton, one can see that the update of the memory is done when a color is seen — it could be called a "chromatic" memory skeleton. Finite-memory strategies could be alternatively defined with an update with states seen in the game. The benefit of this definition of finite memory is that it does not depend on the underlying arena, only on the colors. Hence, we can talk about a memory skeleton that can be used with a winning objective in all arenas.*

To implement a strategy from a memory skeleton, we need a function mapping a state of the game along with a memory state into a GF-strategy. Such a function is called an action map, and it is defined below.

Definition 1.35 (Action map)**.** *Consider a concurrent arena* $\mathcal{C}$ *and a set of memory states* M*. An* action map *on* M *is a function* $\lambda : M \times Q \to \cup_{q \in Q} \Sigma_{\mathsf{A}}(\mathsf{F}(q))$ *such that for all* $q \in Q$ *and* $m \in M$ *we have* $\lambda(m, q) \in \Sigma_{\mathsf{A}}(\mathsf{F}(q))$*.*

With a memory skeleton and an action map, we can now implement a strategy. This defines finite-memory strategies. In the remainder of this dissertation we will be particularly interested in positional strategies i.e. strategies that can be implemented from a memory skeleton with only one state. In other words, what the strategy plays only depends on the current state (or position) of the game, hence the terminology. This is formally defined below.

Definition 1.36 (Positional, finite-memory strategies)**.** *Consider a concurrent arena* $\mathcal{C}$*. A memory skeleton* $\mathsf{M} = \langle M, m_{\text{init}}, \mu \rangle$ *on* K *and an action map* $\lambda : M \times Q \to \cup_{q \in Q} \Sigma_{\mathsf{A}}(\mathsf{F}(q))$ *implement the strategy* $\mathsf{s_A} : Q^+ \to \cup_{q \in Q} \Sigma_{\mathsf{A}}(\mathsf{F}(q))$ *such that, for all* $\rho \in Q^+$*,* $\mathsf{s_A}(\rho) := \lambda(\mu^*(m_{\text{init}}, \mathsf{col}^*(\mathsf{tl}(\rho))), \rho_{\mathsf{lt}}) \in \Sigma_{\mathsf{A}}(\mathsf{F}(q))$*.*

Given a memory skeleton M*, a strategy* $\mathsf{s_A}$ *is* implementable *by* M *if there is an action map* λ *such that* M *and* λ *implement* $\mathsf{s_A}$*. A strategy* $\mathsf{s_A}$ *is* finite memory *if there exists a finite memory skeleton* M *by which* $\mathsf{s_A}$ *is implemented. If* M *is a singleton, the strategy* $\mathsf{s_A}$ *is said to be* positional*. It can be seen as* *a function* $s_A : Q \to \cup_{q \in Q} \Sigma_A(F(q))$*. If a strategy is not finite memory, then it is* infinite memory.

Observation 1.3. *Consider a concurrent arena* $\mathcal{C}$*. Given a positional strategy per player* $s_A \in S_{\mathcal{C}}^{A}$ *and* $s_B \in S_{\mathcal{C}}^{B}$*, the stochastic tree* $\mathcal{T}_{\mathcal{C}}^{s_A,s_B}$ *induced by* s_A *and* s_B *is in fact a Markov chain.*

1.4.5 . Markov decision process

Finally, we present Markov decision processes. They can be seen as one-player games since, from every state of the game, exactly one player decides the next state. However, just like for turn-based game (recall Remark 1.4), we will only consider standard interactions.

Definition 1.37 (Markov decision process)**.** *A* Markov decision process *(MDP for short)* Γ *is a concurrent arena* $\Gamma := \langle Q, F, K, col \rangle$ *where all local interactions are turn-based for the same player. For all* $q \in Q$*, we denote by* Act_q *the set of actions available at state* q *and by* $\varrho_q : \mathsf{Act}_q \to \mathcal{D}(Q)$ *the function mapping each action to a distribution over successor states. (Both* Act_q *and* ϱ_q *are given by the game form* $F(q)$*.)*

The useful objects in MDPs are the end components [52], informally sub-MDPs that are strongly connected, similar to BSCC in Markov chains (recall Definition 1.9).

Definition 1.38 (End component)**.** *Consider an MDP* Γ*. An* end component *(EC for short)* H *in* Γ *is a pair* (Q_H, β_H) *such that* $Q_H \subseteq Q$ *is a subset of states and, for all* $q \in Q_H$*, we have* $\beta_H(q) \subseteq \mathsf{Act}_q$ *the subset of actions compatible with the EC* H *such that:*

- *for all* $q \in Q_H$ *and* $c \in \beta_H(q)$*, we have* $\mathsf{Sp}(\varrho_q(c)) \subseteq Q_H$*;*
- *the underlying graph* (Q_H, E) *is strongly connected, where* $(q, q') \in E$ *if and only if there is some* $c \in \beta_H(q)$ *such that* $q' \in \mathsf{Sp}(\varrho_q(c))$*.*

An end component H *can be seen as a concurrent arena. In that case, it is denoted* $\mathcal{C}_H$*. We denote by* $\mathcal{E}_\Gamma$ *the set of all ECs in the MDP* Γ*.*

In fact, similarly to what happens in Markov chains (recall Theorem 1.8): almost-surely the set of states seen infinitely often form a BSCC. Here, for all deterministic strategies, the set of states seen infinitely often form an EC. This is a well-known result, see for instance [43, Theorem 10.120], that we recall below in Theorem 1.17.

Theorem 1.17. *Consider a finite-state MDP* Γ *where Player* B *plays. Then, for all deterministic strategies* $s_B \in S_B^{\mathcal{C}}$*, we have:*

$$\mathbb{P}_q \left[\bigcup_{H \in \mathcal{E}_\Gamma} \left((Q^* \cdot Q_H^\omega) \cap \left(\bigcap_{q \in Q_H} (Q^* \cdot \{q\})^\omega \right) \right) \right] = 1$$

In any standard game, once Player A chooses a strategy, we obtain a Markov decision process where Player B plays alone.

Definition 1.39 (Induced Markov decision process). *Consider a standard arena $\mathcal{G}$. Let $\mathsf{s_A} \in \mathsf{PS}^{\mathsf{A}}_{\mathcal{C}}$ be a positional strategy. The Markov decision process $\Gamma^{\mathsf{s_A}}_{\mathcal{C}}$ (MDP for short)* induced *by the strategy $\mathsf{s_A}$ is the equal to $\Gamma^{\mathsf{s_A}}_{\mathcal{C}} := \langle Q, \mathsf{F}^{\mathsf{s_A}}, \mathsf{K}, \mathsf{col}\rangle$ where, for all $q \in Q$, we have $\mathsf{F}^{\mathsf{s_A}}(q) := \langle *, \mathsf{Act}^q_{\mathsf{B}}, Q, \varrho_q(\mathsf{s_A}(q), \cdot)\rangle$.*

1.5 Appendix

1.5.1 . Proof of Proposition 1.4

Proof. Consider an open set $B = \cup_{\rho \in A}\mathsf{Cyl}(\rho) \in \mathsf{Borel}(Q)$ for some $A \subseteq Q^*$. Then, we have $\pi \cdot B = \cup_{\rho \in \pi \cdot A}\mathsf{Cyl}(\rho) \in \mathsf{Borel}(Q)$. Furthermore, if there is a finite path $\rho \in A$ such that $\rho \subseteq \pi$ then $\pi^{-1} \cdot B = Q^\omega \in \mathsf{Borel}(Q)$. Otherwise, we let $A_\pi := \{\rho \in Q^+ \mid \pi \cdot \rho \in A\}$. Then, $\pi^{-1} \cdot B = \cup_{\rho \in A_\pi}\mathsf{Cyl}(\rho) \in \mathsf{Borel}(Q)$.

Consider now any Borel set $B \in \mathsf{Borel}(Q)$. We have $Q^\omega \setminus (\pi \cdot B) = \cup_{\rho \in Q^{|\pi|} \setminus \{\pi\}}\mathsf{Cyl}(\rho) \cup \pi \cdot (Q^\omega \setminus B)$. In addition, $Q^\omega \setminus (\pi^{-1} \cdot B) = \pi^{-1} \cdot (Q^\omega \setminus B)$. Finally, for all $(B_n)_{n \in \mathbb{N}} \in (\mathsf{Borel}(Q))^{\mathbb{N}}$, we have $\pi \cdot (\cup_{n \in \mathbb{N}} B_n) = \cup_{n \in \mathbb{N}}(\pi \cdot B_n)$ and $\pi^{-1} \cdot (\cup_{n \in \mathbb{N}} B_n) = (\cup_{n \in \mathbb{N}} \pi^{-1} \cdot B_n)$. By definition of the set $\mathsf{Borel}(Q)$, the property is ensured for all Borel sets $B \in \mathsf{Borel}(Q)$.

Furthermore, for all $\alpha \in [0, 1]$, we have:

$$(f^\pi)^{-1}([0, \alpha]) = \pi^{-1} \cdot f^{-1}([0, \alpha]) \in \mathsf{Borel}(Q)$$

Hence, the residual function f^π is measurable. □

1.5.2 . Proof of Lemma 1.7

Proof. We let $\Omega := Q \cup Q'$, $\Omega^*_{\mathsf{sq}} := (Q \cdot Q')^* \cup (Q \cdot Q')^+ \cdot Q$ and $\Omega^\omega_{\mathsf{sq}} := (Q \cdot Q')^\omega$. Note that $(Q \cdot Q')^\uparrow = \Omega^*_{\mathsf{sq}} \cup \Omega^\omega_{\mathsf{sq}}$. Let us first show that the function $f_{Q,Q'} : \Omega^\omega \to [0, 1]$ is measurable. First, the set $\Omega^\omega_{\mathsf{sq}} \subseteq \Omega^\omega$ is Borel since it is closed. Indeed we have

$$\Omega^\omega \setminus \Omega^\omega_{\mathsf{sq}} = \bigcup_{\pi \in \Omega^* \setminus ((Q \cdot Q')^* \cup (Q \cdot Q')^+ \cdot Q)} \mathsf{Cyl}(\pi)$$

Now, let us show that for all Borel sets $B \in \mathsf{Borel}(Q)$, we have $(\phi_{Q,Q'})^{-1}[B] \in \mathsf{Borel}(\Omega)$. We proceed similarly to what we did in the proof of Proposition 1.4. Consider first an open set $B = \cup_{\pi \in A}\mathsf{Cyl}(\pi)$ for some $A \subseteq Q^+$. For all $\pi = \pi_0 \ldots \pi_n \in A$, we let:

$$\mathsf{PrIm}(\pi) := \cup_{\rho = \rho_0 \cdots \rho_{n-1} \in (Q')^n}\mathsf{Cyl}(\pi_0 \cdot \rho_0 \cdots \rho_{n-1} \cdot \pi_n)$$

Then, we have:

$$(\phi_{Q,Q'})^{-1}[B] = \bigcup_{\pi \in A} \mathsf{PrIm}(\pi) \cap \Omega^\omega_{\mathsf{sq}}$$

Hence, $(\phi_{Q,Q'})^{-1}[B]$ is Borel. Furthermore, for all $B \in \mathsf{Borel}(Q)$, we have:

$$(\phi_{Q,Q'})^{-1}[Q^\omega \setminus B] = \Omega^\omega \setminus (\phi_{Q,Q'})^{-1}[B]$$

and for all $(B_n)_{n\in\mathbb{N}} \in (\mathsf{Borel}(Q))^{\mathbb{N}}$, we have:

$$(\phi_{Q,Q'})^{-1}[\cup_{n\in\mathbb{N}} B_n] = \cup_{n\in\mathbb{N}} (\phi_{Q,Q'})^{-1}[B_n]$$

It follows that $(\phi_{Q,Q'})^{-1}[B] \subseteq \Omega^\omega$ is Borel, for all $B \in \mathsf{Borel}(Q)$.

Now, consider some $\alpha \in [0,1]$. We have:

$$f_{Q,Q'}^{-1}[[0,\alpha]] = \Omega^\omega \setminus \Omega^\omega_{\mathsf{sq}} \cup (\phi_{Q,Q'})^{-1}[f^{-1}[[0,\alpha]]] \in \mathsf{Borel}(\Omega)$$

Therefore, the function $f_{Q,Q'}$ is measurable.

Let us now show the equality of the expected values. We let $\tilde{\mathbb{P}_q} : \mathsf{Borel}(Q) \to [0,1]$ be the function such that, for all Borel sets $B \in \mathsf{Borel}(Q)$, $\tilde{\mathbb{P}_q}[B] := \mathbb{P}'_q[\mathsf{Ant}_q[B]] \in [0,1]$ where $\mathsf{Ant}_q[B] := q^{-1} \cdot (\phi_{Q,Q'})^{-1}[q \cdot B] \in \mathsf{Borel}(\Omega)$. Then, consider countably many disjoint Borel sets $(B_n)_{n\in\mathbb{N}} \in (\mathsf{Borel}(Q))^{\mathbb{N}}$. We have:

$$\begin{aligned}
\mathsf{Ant}_q[\uplus_{n\in\mathbb{N}} B_n] &= \{\rho \in \Omega^\omega \mid q \cdot \rho \in (\phi_{Q,Q'})^{-1}[q \cdot (\uplus_{n\in\mathbb{N}} B_n)]\} \\
&= \{\rho \in \Omega^\omega \mid q \cdot \rho \in \uplus_{n\in\mathbb{N}} (\phi_{Q,Q'})^{-1}[q \cdot (B_n)]\} \\
&= \uplus_{n\in\mathbb{N}} \{\rho \in \Omega^\omega \mid q \cdot \rho \in (\phi_{Q,Q'})^{-1}[q \cdot (B_n)]\} \\
&= \uplus_{n\in\mathbb{N}} \mathsf{Ant}_q[B_n]
\end{aligned}$$

Therefore, since $\mathbb{P}'_q$ is a probability measure on Ω^ω and all sets $(\mathsf{Ant}_q[\uplus_{n\in\mathbb{N}} B_n])_{n\in\mathbb{N}}$ are disjoint, we have $\tilde{\mathbb{P}_q}[\uplus_{n\in\mathbb{N}} B_n] = \sum_{n\in\mathbb{N}} \tilde{\mathbb{P}_q}[B_n]$. In addition, $\mathsf{Ant}_q[\emptyset] = \emptyset$ and $\mathsf{Ant}_q[Q^\omega] = Q' \cdot \Omega^\omega_{\mathsf{sq}}$. Hence $\mathbb{P}'_q[\mathsf{Ant}_q[Q^\omega]] = 1$ since the stochastic tree $\mathcal{T}$ is (Q,Q')-alternating. Therefore, the function $\tilde{\mathbb{P}_q}$ is a probability measure over Q^ω.

Furthermore, for all $\pi \in Q^*$, we have $\mathsf{Ant}_q[\mathsf{Cyl}(\pi)] = \cup_{\pi' \in \mathsf{T}(\pi)} \mathsf{Cyl}(\pi') \cap \Omega^\omega_{\mathsf{sq}}$. Hence, we have, by assumption of the lemma and since the stochastic tree $\mathcal{T}$ is (Q,Q')-alternating:

$$\mathbb{P}_q[\mathsf{Cyl}(\pi)] = \mathbb{P}'_q[\cup_{\pi' \in \mathsf{T}(\pi)} \mathsf{Cyl}(\pi')] = \tilde{\mathbb{P}_q}[\mathsf{Cyl}(\pi)]$$

Hence, by Lemma 1.2, we have that, for all Borel sets $B \in \mathsf{Borel}(Q)$, $\mathbb{P}_q[B] = \tilde{\mathbb{P}_q}[B] = \mathbb{P}'_q[\mathsf{Ant}_q[B]]$. Now, consider any step function $g = \sum_{i=1}^n \alpha_i \cdot \mathbb{1}_{B_i} : Q^\omega \to [0,1]$, where for all $i \in [\![1,n]\!]$, we have $B_i \in \mathsf{Borel}(Q)$. Since $\mathcal{T}'$ is (Q,Q')-alternating, we have:

$$\mathbb{E}'_q[g^q_{Q,Q'}] = \mathbb{E}'_q[g^q_{Q,Q'} \cap Q' \cdot \Omega^\omega_{\mathsf{sq}}] = \sum_{i=1}^n \alpha_i \cdot \mathbb{P}'_q[q^{-1} \cdot (\phi_{Q,Q'})^{-1}[B_i]]$$

Furthermore:

$$\mathbb{E}_q[g^q] = \sum_{i=1}^n \alpha_i \cdot \mathbb{P}_q[q^{-1} \cdot B_i] = \sum_{i=1}^n \alpha_i \cdot \mathbb{P}'_q[q^{-1} \cdot (\phi_{Q,Q'})^{-1}[B_i]] = \mathbb{E}'_q[g^q_{Q,Q'}]$$

For all measurable functions $f : Q^\omega \to [0,1]$ and all step functions $f_n : Q^\omega \to [0,1]$ such that $f_n \leq f$, we have $(f_n)_{Q,Q'} \leq f_{Q,Q'}$. Hence, we can conclude that, for all measurable functions $f : Q^\omega \to [0,1]$, we have $\mathbb{E}_q[f^q] \leq \mathbb{E}'_q[(f_{Q,Q'})^q]$ by definition of these expected values. Furthermore, for $f : Q^\omega \to [0,1]$ a measurable function, $1-f$ is also a measurable function. Hence, $1 - \mathbb{E}_q[f^q] = \mathbb{E}_q[(1-f)^q] \leq \mathbb{E}'_q[((1-f)_{Q,Q'})^q]$. In addition, we have $\mathbb{E}'_q[((1-f)_{Q,Q'})^q] = \mathbb{E}'_q[1-(f_{Q,Q'})^q] = 1 - \mathbb{E}'_q[(f_{Q,Q'})^q]$ since the stochastic tree $\mathcal{T}'$ is (Q,Q')-alternating and the functions $((1-f)_{Q,Q'})^q$ and $1-(f_{Q,Q'})^q$ coincide on $Q' \cdot \Omega^\omega_{\mathsf{sq}}$. We obtain: $\mathbb{E}_q[f^q] = \mathbb{E}'_q[(f_{Q,Q'})^q]$. □

1.5.3 . Proof of Lemma 1.9

Proof. Consider a game in normal form $\mathcal{F}$. Let $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathcal{F})$. We have:

$$\mathsf{val}[\mathcal{F}](\sigma_\mathsf{A}) = \inf_{\sigma_\mathsf{B} \in \Sigma_\mathsf{B}(\mathcal{F})} \mathsf{out}[\mathcal{F}](\sigma_\mathsf{A}, \sigma_\mathsf{B}) \leq \inf_{\sigma_\mathsf{B} \in \Sigma_\mathsf{B}(\mathcal{F})} \sup_{\sigma'_\mathsf{A} \in \Sigma_\mathsf{A}(\mathcal{F})} \mathsf{out}[\mathcal{F}](\sigma'_\mathsf{A}, \sigma_\mathsf{B}) = \mathsf{val}[\mathcal{F}](\mathsf{B})$$

As this holds for all $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathcal{F})$, it follows that:

$$\mathsf{val}[\mathcal{F}](\mathsf{A}) = \sup_{\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathcal{F})} \mathsf{val}[\mathcal{F}](\sigma_\mathsf{A}) \leq \mathsf{val}[\mathcal{F}](\mathsf{B})$$

□

1.5.4 . Proof of Lemma 1.10

First, we prove a straightforward lemma that gives explicitely what is the outcome of a game in normal form.

Lemma 1.18. *Consider a game form $\mathcal{F}$, a valuation $v : \mathsf{O} \to [0,1]$ and a strategy per player $(\sigma_\mathsf{A}, \sigma_\mathsf{B}) \in \Sigma_\mathsf{A} \times \Sigma_\mathsf{B}$. We have:*

$$\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) = \sum_{o \in \mathsf{Sp}(\delta(\sigma_\mathsf{A}, \sigma_\mathsf{B}))} v(o) \cdot \varrho(\sigma_\mathsf{A}, \sigma_\mathsf{B})(o)$$

Proof. By definition of the expected value, we have:

$$\begin{aligned}
\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) &= \sum_{x \in [0,1]} \sum_{\substack{o \in \mathsf{Sp}(\delta(\sigma_\mathsf{A}, \sigma_\mathsf{B})) \\ v(o) = x}} x \cdot \varrho(\sigma_\mathsf{A}, \sigma_\mathsf{B})(o) \\
&= \sum_{x \in [0,1]} \sum_{o \in \mathsf{Sp}(\delta(\sigma_\mathsf{A}, \sigma_\mathsf{B}))} x \cdot \mathbb{1}_{v^{-1}[x]}(o) \cdot \varrho(\sigma_\mathsf{A}, \sigma_\mathsf{B})(o) \\
&= \sum_{o \in \mathsf{Sp}(\delta(\sigma_\mathsf{A}, \sigma_\mathsf{B}))} \sum_{x \in [0,1]} x \cdot \mathbb{1}_{v^{-1}[x]}(o) \cdot \varrho(\sigma_\mathsf{A}, \sigma_\mathsf{B})(o) \\
&= \sum_{o \in \mathsf{Sp}(\delta(\sigma_\mathsf{A}, \sigma_\mathsf{B}))} v(o) \cdot \varrho(\sigma_\mathsf{A}, \sigma_\mathsf{B})(o)
\end{aligned}$$

□

The above lemma will be implicitly used in this dissertation when manipulating outcomes of games in normal form.

We can now proceed to the proof of Lemma 1.10.

Proof. We have:

$$\begin{aligned}\mathsf{out}[\langle \mathcal{F}, \sum_{n \in \mathbb{N}} \lambda_n \cdot v_n \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) &= \sum_{o \in \mathsf{Sp}(\delta(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}))} (\sum_{n \in \mathbb{N}} \lambda_n \cdot v_n)(o) \cdot \varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})(o) \\ &= \sum_{n \in \mathbb{N}} \lambda_n \cdot \sum_{o \in \mathsf{Sp}(\delta(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}))} v_n(o) \cdot \varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})(o) \\ &= \sum_{n \in \mathbb{N}} \lambda_n \cdot \mathsf{out}[\langle \mathcal{F}, v_n \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})\end{aligned}$$

Consider any two valuations $v, v' : \mathsf{O} \to [0, 1]$, $\lambda \geq 0$ and $x \in \mathbb{R}$ such that $\lambda \cdot v + x \leq v'$. We have:

$$\begin{aligned}\mathsf{out}[\langle \mathcal{F}, v' \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) &= \sum_{o \in \mathsf{Sp}(\delta(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}))} v'(o) \cdot \varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})(o) \\ &\geq \sum_{o \in \mathsf{Sp}(\delta(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}))} (\lambda \cdot v + x)(o) \cdot \varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})(o) \\ &= \lambda \cdot \sum_{o \in \mathsf{Sp}(\delta(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}))} v(o) \cdot \varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})(o) + x \\ &= \lambda \cdot \mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) + x \\ &\geq \lambda \cdot \mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}) + x\end{aligned}$$

Since this holds for all $\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}$, it follows that:

$$\lambda \cdot \mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}) + x \leq \mathsf{val}[\langle \mathcal{F}, v' \rangle](\sigma_{\mathsf{A}}) \leq \mathsf{val}[\langle \mathcal{F}, v' \rangle](\mathsf{A})$$

Therefore:

$$\mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}) \leq \frac{1}{\lambda} \cdot (\mathsf{val}[\langle \mathcal{F}, v' \rangle](\mathsf{A}) - x)$$

Since this holds for all $\sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}}$, it follows that:

$$\mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{A}) \leq \frac{1}{\lambda} \cdot (\mathsf{val}[\langle \mathcal{F}, v' \rangle](\mathsf{A}) - x)$$

Thus:

$$\lambda \cdot \mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{A}) + x \leq \mathsf{val}[\langle \mathcal{F}, v' \rangle](\mathsf{A})$$

Finally, assume that $\lambda \cdot v + x : \mathsf{O} \to [0, 1]$. Then, denoting $v' := \lambda \cdot v + x : \mathsf{O} \to [0, 1]$ we have shown that:

$$\lambda \cdot \mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{A}) + x \leq \mathsf{val}[\langle \mathcal{F}, v' \rangle](\mathsf{A})$$

Furthermore, we have $\frac{1}{\lambda} v' - \frac{x}{\lambda} = v$. Therefore, we have:

$$\frac{1}{\lambda} \mathsf{val}[\langle \mathcal{F}, v' \rangle](\mathsf{A}) - \frac{x}{\lambda} \leq \mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{A})$$

That is:

$$\mathsf{val}[\langle \mathcal{F}, \lambda \cdot v \cdot x \rangle](\mathsf{A}) = \mathsf{val}[\langle \mathcal{F}, v' \rangle](\mathsf{A}) \leq \lambda \cdot \mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{A}) + x$$

We obtain:

$$\mathsf{val}[\langle \mathcal{F}, \lambda \cdot v \cdot x \rangle](\mathsf{A}) = \lambda \cdot \mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{A}) + x$$

This is analogous if $s = \sigma_{\mathsf{B}}$ or $s = \mathsf{B}$.

□

1.5.5 . Proof of Observation 1.1

Proof. Consider such a finite set $\mathsf{S_C} \subseteq \Sigma_{\mathsf{C}}(\mathcal{F})$ of Player-C strategies and assume that it supremizes $\mathcal{F}$. Let $v : \mathsf{O} \to [0,1]$ be a valuation of the outcomes. Since $\mathsf{S_C}$ supremizes the game form $\mathcal{F}$, it follows that for all $n \in \mathbb{N}$, there is a Player-C strategy $\sigma_n \in \Sigma_{\mathsf{C}}(\mathcal{F})$ that is $\frac{1}{n+1}$-optimal in the game in normal form $\langle \mathcal{F}, v \rangle$. Since the set $\mathsf{S_C}$ is finite, there must be a Player-C strategy $\sigma_{\mathsf{C}} \in \Sigma_{\mathsf{C}}(\mathcal{F})$ such that $\sigma_{\mathsf{C}} = \sigma_n$ for infinitely many $n \in \mathbb{N}$. Therefore, we have $\mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{C}}) = \mathsf{val}[\langle \mathcal{F}, v \rangle]$ since, for infinitely many $n \in \mathbb{N}$, we have $|\mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{C}}) - \mathsf{val}[\langle \mathcal{F}, v \rangle]| \leq \frac{1}{n+1}$. That is, the strategy $\sigma_{\mathsf{C}} \in \mathsf{S_C}$ is optimal in the game in normal form $\langle \mathcal{F}, v \rangle$. As this holds for all valuations of the outcomes $v : \mathsf{O} \to [0,1]$, it follows that the set $\mathsf{S_C}$ maximizes the game form $\mathcal{F}$ w.r.t. Player C. □

1.5.6 . Proof of Proposition 1.12

Before considering the proof of Proposition 1.12, let us state and prove the lemma below:

Lemma 1.19. *Consider a standard game form $\mathcal{F}$. Consider two valuations of the outcomes $v, v' : \mathsf{O} \to [0,1]$, two Player-A strategies $\sigma_{\mathsf{A}}, \sigma'_{\mathsf{A}} \in \Sigma_{\mathsf{A}}(\mathcal{F})$ and two Player-B strategies $\sigma_{\mathsf{B}}, \sigma'_{\mathsf{B}} \in \Sigma_{\mathsf{B}}(\mathcal{F})$. Then:*

$$|\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) - \mathsf{out}[\langle \mathcal{F}, v' \rangle](\sigma'_{\mathsf{A}}, \sigma'_{\mathsf{B}})| \leq \|v - v'\|_\infty + \|\sigma_{\mathsf{A}} - \sigma'_{\mathsf{A}}\|_1 + \|\sigma_{\mathsf{B}} - \sigma'_{\mathsf{B}}\|_1$$

Note that $\|\sigma_{\mathsf{A}} - \sigma'_{\mathsf{A}}\|_1$ is well defined since both σ_{A} and σ'_{A} have countable support, and similarly for σ_{B} and σ'_{B}. Hence:

$$|\mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}) - \mathsf{val}[\langle \mathcal{F}, v' \rangle](\sigma'_{\mathsf{A}})| \leq \|v - v'\|_\infty + \|\sigma_{\mathsf{A}} - \sigma'_{\mathsf{A}}\|_1$$

and

$$|\mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{A}) - \mathsf{val}[\langle \mathcal{F}, v' \rangle](\mathsf{A})| \leq \|v - v'\|_\infty$$

These inequalities also hold for Player B.

Proof. The proof is quite long while the idea is very simple: it just amounts to expand Definition 1.13 in the context of standard game forms. However, it involves nested sums, hence the length of the proof. We use several times that $\sum_{a \in \mathsf{Act}_\mathsf{A}} \sigma_\mathsf{A}(a) = 1 = \sum_{b \in \mathsf{Act}_\mathsf{B}} \sigma_\mathsf{B}(b)$ and that $v(o) \leq 1$ for all $o \in \mathsf{O}$. We have:

$$\begin{aligned}
&|\mathsf{out}[\langle \mathcal{F}, v\rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) - \mathsf{out}[\langle \mathcal{F}, v'\rangle](\sigma'_\mathsf{A}, \sigma'_\mathsf{B})| \\
&= |\sum_{o \in \mathsf{O}} \sum_{a \in \mathsf{Act}_\mathsf{A}} \sum_{b \in \mathsf{Act}_\mathsf{B}} (\sigma_\mathsf{A}(a) \cdot \sigma_\mathsf{B}(b) \cdot v(o) \cdot \varrho(a,b)(o) - \sigma'_\mathsf{A}(a) \cdot \sigma'_\mathsf{B}(b) \cdot v'(o) \cdot \varrho(a,b)(o))| \\
&\leq \sum_{o \in \mathsf{O}} \sum_{a \in \mathsf{Act}_\mathsf{A}} \sum_{b \in \mathsf{Act}_\mathsf{B}} |(\sigma_\mathsf{A}(a) \cdot \sigma_\mathsf{B}(b) \cdot v(o) \cdot \varrho(a,b)(o) - \sigma'_\mathsf{A}(a) \cdot \sigma'_\mathsf{B}(b) \cdot v'(o) \cdot \varrho(a,b)(o))| \\
&\leq \sum_{o \in \mathsf{O}} \sum_{a \in \mathsf{Act}_\mathsf{A}} \sum_{b \in \mathsf{Act}_\mathsf{B}} \sigma_\mathsf{A}(a) \cdot |\sigma_\mathsf{B}(b) \cdot v(o) \cdot \varrho(a,b)(o) - \sigma'_\mathsf{B}(b) \cdot v'(o) \cdot \varrho(a,b)(o)| \\
&+ \sum_{o \in \mathsf{O}} \sum_{a \in \mathsf{Act}_\mathsf{A}} \sum_{b \in \mathsf{Act}_\mathsf{B}} |\sigma_\mathsf{A}(a)' - \sigma_\mathsf{A}(a)| \cdot |\sigma'_\mathsf{B}(b) \cdot v'(o) \cdot \varrho(a,b)(o)|
\end{aligned}$$

We let:

$$x := \sum_{o \in \mathsf{O}} \sum_{a \in \mathsf{Act}_\mathsf{A}} \sum_{b \in \mathsf{Act}_\mathsf{B}} \sigma_\mathsf{A}(a) \cdot |\sigma_\mathsf{B}(b) \cdot v(o) \cdot \varrho(a,b)(o) - \sigma'_\mathsf{B}(b) \cdot v'(o) \cdot \varrho(a,b)(o)|$$

and

$$y := \sum_{o \in \mathsf{O}} \sum_{a \in \mathsf{Act}_\mathsf{A}} \sum_{b \in \mathsf{Act}_\mathsf{B}} |\sigma_\mathsf{A}(a)' - \sigma_\mathsf{A}(a)| \cdot |\sigma'_\mathsf{B}(b) \cdot v'(o) \cdot \varrho(a,b)(o)|$$

Thus, we have $|\mathsf{out}[\langle \mathcal{F}, v\rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) - \mathsf{out}[\langle \mathcal{F}, v'\rangle](\sigma'_\mathsf{A}, \sigma'_\mathsf{B})| \leq x + y$. Let us first deal with y:

$$\begin{aligned}
y &= \sum_{o \in \mathsf{O}} \sum_{a \in \mathsf{Act}_\mathsf{A}} \sum_{b \in \mathsf{Act}_\mathsf{B}} |\sigma_\mathsf{A}(a)' - \sigma_\mathsf{A}(a)| \cdot |\sigma'_\mathsf{B}(b) \cdot v'(o) \cdot \varrho(a,b)(o)| \\
&= \sum_{a \in \mathsf{Act}_\mathsf{A}} |(\sigma_\mathsf{A}(a)' - \sigma_\mathsf{A}(a))| \cdot \left(\sum_{b \in \mathsf{Act}_\mathsf{B}} \sigma'_\mathsf{B}(b) \cdot \sum_{o \in \mathsf{O}} |v'(o) \cdot \varrho(a,b)(o)| \right) \\
&\leq \sum_{a \in \mathsf{Act}_\mathsf{A}} |(\sigma_\mathsf{A}(a)' - \sigma_\mathsf{A}(a))| \cdot \left(\sum_{b \in \mathsf{Act}_\mathsf{B}} \sigma'_\mathsf{B}(b) \right) \\
&= \sum_{a \in \mathsf{Act}_\mathsf{A}} |(\sigma_\mathsf{A}(a)' - \sigma_\mathsf{A}(a))| = \|\sigma'_\mathsf{A} - \sigma_\mathsf{A}\|_1
\end{aligned}$$

Let us now deal with x:

$$x = \sum_{o \in \mathsf{O}} \sum_{a \in \mathsf{Act_A}} \sum_{b \in \mathsf{Act_B}} \sigma_{\mathsf{A}}(a) \cdot |\sigma_{\mathsf{B}}(b) \cdot v(o) \cdot \varrho(a,b)(o) - \sigma'_{\mathsf{B}}(b) \cdot v'(o) \cdot \varrho(a,b)(o)|$$
$$\leq \sum_{a \in \mathsf{Act_A}} \sigma_{\mathsf{A}}(a) \cdot \left(\sum_{b \in \mathsf{Act_B}} \sigma_{\mathsf{B}}(b) \cdot \sum_{o \in \mathsf{O}} \varrho(a,b)(o) \cdot |v(o) - v'(o)| \right.$$
$$\left. + \sum_{b \in \mathsf{Act_B}} |\sigma_{\mathsf{B}}(b) - \sigma'_{\mathsf{B}}(b)| \cdot \sum_{o \in \mathsf{O}} \varrho(a,b)(o) \cdot |v(o) - v'(o)| \right)$$
$$\leq \sum_{a \in \mathsf{Act_A}} \sigma_{\mathsf{A}}(a) \cdot \left(\sum_{b \in \mathsf{Act_B}} \sigma_{\mathsf{B}}(b) \cdot \|v - v'\|_\infty + \sum_{b \in \mathsf{Act_B}} |\sigma_{\mathsf{B}}(b) - \sigma'_{\mathsf{B}}(b)| \cdot \|v - v'\|_\infty \right)$$
$$= \sum_{a \in \mathsf{Act_A}} \sigma_{\mathsf{A}}(a) \cdot (\|v - v'\|_\infty + \|\sigma_{\mathsf{B}} - \sigma'_{\mathsf{B}}\|_1) = \|v - v'\| + \|\sigma_{\mathsf{B}} - \sigma'_{\mathsf{B}}\|_1$$

Overall, we do obtain:

$$|\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) - \mathsf{out}[\langle \mathcal{F}, v' \rangle](\sigma'_{\mathsf{A}}, \sigma'_{\mathsf{B}})| \leq \|\sigma'_{\mathsf{A}} - \sigma_{\mathsf{A}}\|_1 + \|\sigma_{\mathsf{B}} - \sigma'_{\mathsf{B}}\|_1 + \|v - v'\|_\infty$$

This proves the first inequality of Lemma 1.19. Let us consider the second one. For all positive $\varepsilon > 0$, let $\sigma^\varepsilon_{\mathsf{B}} \in \Sigma_{\mathsf{B}}(\mathcal{F})$ be such that

$$\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}, \sigma^\varepsilon_{\mathsf{B}}) \leq \mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}) + \varepsilon$$

Then:

$$\begin{aligned}
\mathsf{val}[\langle \mathcal{F}, v' \rangle](\sigma'_{\mathsf{A}}) &\leq \mathsf{out}[\langle \mathcal{F}, v' \rangle](\sigma'_{\mathsf{A}}, \sigma^\varepsilon_{\mathsf{B}}) \\
&\leq \mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}, \sigma^\varepsilon_{\mathsf{B}}) + \|v - v'\|_\infty + \|\sigma_{\mathsf{A}} - \sigma'_{\mathsf{A}}\|_1 \\
&\leq \mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}) + \|v - v'\|_\infty + \|\sigma_{\mathsf{A}} - \sigma'_{\mathsf{A}}\|_1 + \varepsilon
\end{aligned}$$

As this holds for all $\varepsilon > 0$, it follows that:

$$\mathsf{val}[\langle \mathcal{F}, v' \rangle](\sigma'_{\mathsf{A}}) \leq \mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}) + \|v - v'\|_\infty + \|\sigma_{\mathsf{A}} - \sigma'_{\mathsf{A}}\|_1$$

By symmetry, we obtain that:

$$|\mathsf{val}[\langle \mathcal{F}, v' \rangle](\sigma'_{\mathsf{A}}) - \mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}})| \leq +\|v - v'\|_\infty + \|\sigma_{\mathsf{A}} - \sigma'_{\mathsf{A}}\|_1$$

Let us now consider to the third inequality. We proceed similarly to the previous one: For all positive $\varepsilon > 0$, let $\sigma^\varepsilon_{\mathsf{A}} \in \Sigma_{\mathsf{A}}(\mathcal{F})$ be such that

$$\mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma^\varepsilon_{\mathsf{A}}) \geq \mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{A}) - \varepsilon$$

Then:

$$\begin{aligned}\mathsf{val}[\langle \mathcal{F}, v' \rangle](\mathsf{A}) &\geq \mathsf{val}[\langle \mathcal{F}, v' \rangle](\sigma_{\mathsf{A}}^{\varepsilon}) \\ &\geq \mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}^{\varepsilon}) - \|v - v'\|_{\infty} \\ &\geq \mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{A}) - \|v - v'\|_{\infty} - \varepsilon\end{aligned}$$

As this holds for all $\varepsilon > 0$, it follows that:

$$\mathsf{val}[\langle \mathcal{F}, v' \rangle](\mathsf{A}) \geq \mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{A}) - \|v - v'\|_{\infty}$$

By symmetry, we obtain that:

$$|\mathsf{val}[\langle \mathcal{F}, v' \rangle](\mathsf{A}) - \mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{A})| \leq \|v - v'\|_{\infty}$$

□

Now, we state and show below that there is a value and optimal strategies for both players in the case where the valuation takes finitely many values.

Lemma 1.20. *Consider a standard deterministic game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ on a set of outcomes O where either of the players has finitely many actions. Consider a valuation $v : \mathsf{O} \to [0,1]$ such that $v[\mathsf{O}] \subseteq [0,1]$ is finite. Then, there is a value in the game in normal form $\langle \mathcal{F}, v \rangle$ and both players have optimal strategies.*

Proof. We prove the lemma in the case where $\mathsf{Act}_{\mathsf{A}}$ is finite. The other case is analogous. For all $b, b' \in \mathsf{Act}_{\mathsf{B}}$, we say that b is equivalent to b', denoted $b \sim b'$ if, $v \circ \varrho(\cdot, b) : \mathsf{Act}_{\mathsf{A}} \to [0,1] = v \circ \varrho(\cdot, b') : \mathsf{Act}_{\mathsf{A}} \to [0,1]$. That is, the columns corresponding to actions b and b' are identical. Clearly, $\sim$ is an equivalence relation over $\mathsf{Act}_{\mathsf{B}} \times \mathsf{Act}_{\mathsf{B}}$. Let $R_{\mathsf{B}} \subseteq \mathcal{P}(\mathsf{Act}_{\mathsf{B}})$ be the (non-empty) set of equivalence classes of the equivalence relation $\sim$ where $\mathcal{P}(\mathsf{Act}_{\mathsf{B}})$ refers to the set of subsets of $\mathsf{Act}_{\mathsf{B}}$. In fact, R_{B} is finite. Indeed, for all Player-B actions $b \in \mathsf{Act}_{\mathsf{B}}$, we have $v \circ \varrho(\cdot, b) : \mathsf{Act}_{\mathsf{A}} \to v[\mathsf{O}]$. Since $\mathsf{Act}_{\mathsf{A}}$ and $v[\mathsf{O}]$ are finite, there are finitely many functions $\mathsf{Act}_{\mathsf{A}} \to v[\mathsf{O}]$. Since an element of R_{B} corresponds to a function $\mathsf{Act}_{\mathsf{A}} \to v[\mathsf{O}]$, it follows that R_{B} is finite. For all $T \in R_{\mathsf{B}}$, we let $b_T \in T$ be a representative of the equivalence class T. Consider now the game form $\mathcal{F}' := \langle \mathsf{Act}_{\mathsf{A}}, R_{\mathsf{B}}, \mathsf{O}, \varrho' \rangle$, where, for all $a \in \mathsf{Act}_{\mathsf{A}}$ and $T \in R_{\mathsf{B}}$, we have $\varrho'(a, T) := \varrho(a, b_T)$. This game in normal form is finite and deterministic and therefore has a value $u := \mathsf{val}[\langle \mathcal{F}', v \rangle] \in [0,1]$ by Theorem 1.11. Let us show that u is in fact the value of the game in normal form $\langle \mathcal{F}, v \rangle$.

We let $g : \Sigma_{\mathsf{B}}(\mathcal{F}) \to \Sigma_{\mathsf{B}}(\mathcal{F}')$ be such that, for all $\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}(\mathcal{F})$, for all $T \in R_{\mathsf{B}}$, we have $g(\sigma_{\mathsf{B}})(T) := \sum_{b \in T} \sigma_{\mathsf{B}}(b)$. We claim that:

$$\forall (\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \in \Sigma_{\mathsf{A}}(\mathcal{F}) \times \Sigma_{\mathsf{B}}(\mathcal{F}),\ \mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) = \mathsf{out}[\langle \mathcal{F}', v \rangle](\sigma_{\mathsf{A}}, g(\sigma_{\mathsf{B}})) \tag{1.1}$$

Indeed, let $(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \in \Sigma_{\mathsf{A}}(\mathcal{F}) \times \Sigma_{\mathsf{B}}(\mathcal{F})$. We have (recall that the game form $\mathcal{F}$ is deterministic):

$$\begin{aligned}
\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) &= \sum_{a \in \mathsf{Act}_{\mathsf{A}}} \sum_{b \in \mathsf{Act}_{\mathsf{B}}} \sigma_{\mathsf{A}}(a) \cdot \sigma_{\mathsf{B}}(b) \cdot v \circ \varrho(a, b) \\
&= \sum_{a \in \mathsf{Act}_{\mathsf{A}}} \sigma_{\mathsf{A}}(a) \cdot \left(\sum_{T \in R_{\mathsf{B}}} \sum_{b \in T} \sigma_{\mathsf{B}}(b) \cdot v \circ \varrho(a, b) \right) \\
&= \sum_{a \in \mathsf{Act}_{\mathsf{A}}} \sigma_{\mathsf{A}}(a) \cdot \left(\sum_{T \in R_{\mathsf{B}}} g(\sigma_{\mathsf{B}})(T) \cdot v \circ \varrho(a, b_T) \right) \\
&= \mathsf{out}[\langle \mathcal{F}', v \rangle](\sigma_{\mathsf{A}}, g(\sigma_{\mathsf{B}}))
\end{aligned}$$

We also let $g' : \Sigma_{\mathsf{B}}(\mathcal{F}') \to \Sigma_{\mathsf{B}}(\mathcal{F})$ be such that, for all $\sigma'_{\mathsf{B}} \in \Sigma_{\mathsf{B}}(\mathcal{F}')$, for all $b \in \mathsf{Act}_{\mathsf{B}}$, we have:

$$g'(\sigma'_{\mathsf{B}})(b) := \begin{cases} \sigma'_{\mathsf{B}}(T) & \text{if } b = b_T \text{ for some } T \in R_{\mathsf{B}} \\ 0 & \text{otherwise} \end{cases}$$

One can see that for all $\sigma'_{\mathsf{B}} \in \Sigma_{\mathsf{B}}(\mathcal{F}')$, we have $g \circ g'(\sigma'_{\mathsf{B}}) = \sigma'_{\mathsf{B}}$.

Consider optimal strategies σ'_{A} and σ'_{B} for both players in the game in normal form $\langle \mathcal{F}', v \rangle$. We claim that $\sigma'_{\mathsf{A}} \in \Sigma_{\mathsf{A}}(\mathcal{F})$ and $g'(\sigma'_{\mathsf{B}}) \in \Sigma_{\mathsf{B}}(\mathcal{F})$ have value u in the game in normal form $\langle \mathcal{F}, v \rangle$.

Consider any Player-B strategy σ_{B} in the game form $\mathcal{F}$. Then, by Equation 1.1:

$$\begin{aligned}
\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma'_{\mathsf{A}}, \sigma_{\mathsf{B}}) &= \mathsf{out}[\langle \mathcal{F}', v \rangle](\sigma'_{\mathsf{A}}, g(\sigma_{\mathsf{B}})) \\
&\geq \mathsf{val}[\langle \mathcal{F}', v \rangle](\sigma'_{\mathsf{A}}) = u
\end{aligned}$$

Hence, $\mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma'_{\mathsf{A}}) \geq u$. Furthermore, for any Player-A strategy σ_{A} in the game form $\mathcal{F}$, by Equation 1.1:

$$\begin{aligned}
\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}, g'(\sigma'_{\mathsf{B}})) &= \mathsf{out}[\langle \mathcal{F}', v \rangle](\sigma_{\mathsf{A}}, g \circ g'(\sigma'_{\mathsf{B}})) \\
&= \mathsf{out}[\langle \mathcal{F}', v \rangle](\sigma_{\mathsf{A}}, \sigma'_{\mathsf{B}}) \\
&\leq \mathsf{val}[\langle \mathcal{F}', v \rangle](\sigma'_{\mathsf{B}}) = u
\end{aligned}$$

Hence, $\mathsf{val}[\langle \mathcal{F}, v \rangle](g(\sigma'_{\mathsf{B}})) \leq u$. Overall, we obtain $\mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}) = u = \mathsf{val}[\langle \mathcal{F}, v \rangle](g'(\sigma'_{\mathsf{B}}))$. ☐

We can proceed to the proof of Proposition 1.12.

Proof. Consider a valuation of the outcomes $v : \mathsf{O} \to [0, 1]$. For all $n \in \mathbb{N}$, we let $v_n : \mathsf{O} \to [0, 1]$ be the valuation of the outcomes such that, for all outcomes $o \in \mathsf{O}$:

$$v_n(o) := \frac{\lfloor 2^n \cdot v(o) \rfloor}{2^n} \in \left\{ \frac{i}{2^n} \mid i \in [\![0, 2^n]\!] \right\}$$

where $\lfloor \cdot \rfloor : \mathbb{R} \to \mathbb{Z}$ is the floor function, that is, for all $x \in \mathbb{R}$, we have $\lfloor x \rfloor \leq x < \lfloor x \rfloor + 1$. Therefore, we have $\|v_n - v\|_\infty \leq \frac{1}{2^n}$. Since v_n takes finitely many values, it follows, by Lemma 1.20, that the game in normal form $\langle \mathcal{F}, v_n \rangle$ has a value and there are optimal strategies for both players in that game. We denote by $u_n := \mathsf{val}[\langle \mathcal{F}, v_n \rangle]$ the value of the game in normal form $\langle \mathcal{F}, v_n \rangle$ and by $\sigma_{\mathsf{A}}^n \in \Sigma_{\mathsf{A}}(\mathcal{F})$ (resp. $\sigma_{\mathsf{B}}^n \in \Sigma_{\mathsf{B}}(\mathcal{F})$) a Player-$\mathsf{A}$ (resp. Player-B) optimal strategy in the game in normal form $\langle \mathcal{F}, v_n \rangle$.

Let $u \in [0,1]$ be subsequential limit (w.r.t. $\| \cdot \|_\infty$) of $(u_n)_{n \in \mathbb{N}}$, that is $u = \lim_{n \to \infty} u_{\varphi(n)}$ for some increasing $\varphi : \mathbb{N} \to \mathbb{N}$ (that exists by Theorem 1.1). Let also σ_{A} be a subsequential limit (w.r.t. $\| \cdot \|_\infty$) of $(\sigma_{\mathsf{A}}^{\varphi(n)})_{n \in \mathbb{N}}$, that is $\sigma_{\mathsf{A}} = \lim_{n \to \infty} \sigma_{\mathsf{A}}^{\psi \circ \varphi(n)}$ for some increasing $\psi : \mathbb{N} \to \mathbb{N}$ (which is possible by Theorem 1.1 since $\mathsf{Act}_{\mathsf{A}}$ is finite). We let $h := \psi \circ \varphi : \mathbb{N} \to \mathbb{N}$. First, let us show that $\sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}}(\mathcal{F})$. Since, for all $a \in \mathsf{Act}_{\mathsf{A}}$ and $n \in \mathbb{N}$, we have $\sigma_{\mathsf{A}}^{h(n)}(a) \in [0,1]$, it follows that we also have $\sigma_{\mathsf{A}}(a) \in [0,1]$. Furthermore:

$$\sum_{a \in \mathsf{Act}_{\mathsf{A}}} \sigma_{\mathsf{A}}(a) = \sum_{a \in \mathsf{Act}_{\mathsf{A}}} \lim_{n \to \infty} \sigma_{\mathsf{A}}^{h(n)}(a) = \lim_{n \to \infty} \sum_{a \in \mathsf{Act}_{\mathsf{A}}} \sigma_{\mathsf{A}}^{h(n)}(a) = \lim_{n \to \infty} 1 = 1$$

Hence, we do have $\sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}}(\mathcal{F})$.

Let us now show that $\mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}) \geq u$. Let $\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}(\mathcal{F})$. For all $n \in \mathbb{N}$, we have, by Lemma 1.19:

$$\begin{aligned}
\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) &\geq \mathsf{out}[\langle \mathcal{F}, v_{h(n)} \rangle](\sigma_{\mathsf{A}}^{h(n)}, \sigma_{\mathsf{B}}) - \|v_{h(n)} - v\|_\infty - \|\sigma_{\mathsf{A}} - \sigma_{\mathsf{A}}^{h(n)}\|_1 \\
&\geq \mathsf{val}[\langle \mathcal{F}, v_{h(n)} \rangle](\sigma_{\mathsf{A}}^{h(n)}) - \frac{1}{2^{h(n)}} - \|\sigma_{\mathsf{A}} - \sigma_{\mathsf{A}}^{h(n)}\|_1 \\
&= u_{h(n)} - \frac{1}{2^{h(n)}} - \|\sigma_{\mathsf{A}} - \sigma_{\mathsf{A}}^{h(n)}\|_1
\end{aligned}$$

Thus:

$$\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \geq u_{h(n)} - \frac{1}{2^{h(n)}} - \|\sigma_{\mathsf{A}} - \sigma_{\mathsf{A}}^{h(n)}\|_1 \tag{1.2}$$

Furthermore, since $\sigma_{\mathsf{A}} = \lim_{n \to \infty} \sigma_{\mathsf{A}}^{h(n)}$, it follows that

$$\lim_{n \to \infty} \|\sigma_{\mathsf{A}} - \sigma_{\mathsf{A}}^{h(n)}\|_1 = 0$$

Since $(u_{h(n)})_{n \in \mathbb{N}}$ is a subsequence of $(u_{\varphi(n)})_{n \in \mathbb{N}}$ and $\lim_{n \to \infty} (u_{\varphi(n)})_{n \in \mathbb{N}} = u$, it follows that:

$$\lim_{n \to \infty} u_{h(n)} = u$$

Since Equation 1.2 holds for all $n \in \mathbb{N}$, we obtain $\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \geq u$. As this holds for all $\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}(\mathcal{F})$, it follows that $\mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}) \geq u$.

Let us now show that $\mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{B}) \leq u$. Let $n \in \mathbb{N}$. We consider the Player-B strategy $\sigma_{\mathsf{B}}^{h(n)} \in \Sigma_{\mathsf{B}}(\mathcal{F})$. Consider any Player-$\mathsf{A}$ strategy $\sigma'_{\mathsf{A}} \in \Sigma_{\mathsf{A}}(\mathcal{F})$.

By Lemma 1.19, we have:

$$\begin{aligned}\mathsf{out}[\langle \mathcal{F}, v\rangle](\sigma'_{\mathsf{A}}, \sigma_{\mathsf{B}}^{h(n)}) &\leq \mathsf{out}[\langle \mathcal{F}, v_{h(n)}\rangle](\sigma'_{\mathsf{A}}, \sigma_{\mathsf{B}}^{h(n)}) + \|v_{h(n)} - v\|_\infty \\ &\leq \mathsf{val}[\langle \mathcal{F}, v_{h(n)}\rangle](\sigma_{\mathsf{B}}^{h(n)}) + \frac{1}{2^{h(n)}} \\ &= u_{h(n)} + \frac{1}{2^{h(n)}}\end{aligned}$$

Thus, $\mathsf{val}[\langle \mathcal{F}, v\rangle](\sigma_{\mathsf{B}}^{h(n)}) \leq u_{h(n)} - \frac{1}{2^{h(n)}}$. Hence, for all $n \in \mathbb{N}$, we have:

$$\mathsf{val}[\langle \mathcal{F}, v\rangle](\mathsf{B}) \leq \mathsf{val}[\langle \mathcal{F}, v\rangle](\sigma_{\mathsf{B}}^{h(n)}) \leq u_{h(n)} + \frac{1}{2^{h(n)}}$$

Again, since $\lim_{n\to\infty} u_{h(n)} = u$, we obtain $\mathsf{val}[\langle \mathcal{F}, v\rangle](\mathsf{B}) \leq u$. Overall, we obtain, $\mathsf{val}[\langle \mathcal{F}, v\rangle] = u = \mathsf{val}[\langle \mathcal{F}, v\rangle](\sigma_{\mathsf{A}})$. □

1.5.7 . Proof of Lemma 1.13

Proof. Consider a non-empty set of outcomes O' and a map $d : \mathsf{O} \to \mathcal{D}(\mathsf{O}')$. For all valuations $v : \mathsf{O}' \to [0,1]$, we let $e_v : \mathsf{O} \to [0,1]$ be equal to $e_v := \mathbb{E}_d(v)$. For all valuations $v : \mathsf{O}' \to [0,1]$ and $(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \in \Sigma_{\mathsf{A}} \times \Sigma_{\mathsf{B}}$, we have:

$$\mathsf{out}[\langle \mathcal{F}, e_v\rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) = \mathsf{out}[\langle \mathcal{F}^d, v\rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \tag{1.3}$$

Indeed, we have:

$$\begin{aligned}\mathsf{out}[\langle \mathcal{F}, e_v\rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) = \mathbb{E}_{\varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})}(e_v) &= \mathbb{E}_{\varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})}(\mathbb{E}_d(v)) \\ &= \sum_{o \in \mathsf{O}} \varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})(o) \cdot \mathbb{E}_d(v)(o) \\ &= \sum_{o \in \mathsf{O}} \varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})(o) \cdot \sum_{o' \in \mathsf{O}'} d(o)(o') \cdot v(o') \\ &= \sum_{o' \in \mathsf{O}'} \sum_{o \in \mathsf{O}} \varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})(o) \cdot d(o)(o') \cdot v(o') \\ &= \sum_{o' \in \mathsf{O}'} \mathbb{E}_{\varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})}(d)(o') \cdot v(o') \\ &= \mathbb{E}_{\mathbb{E}_{\varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})}(d)}(v) = \mathsf{out}[\langle \mathcal{F}^d, v\rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})\end{aligned}$$

In fact, Equation 1.3 gives that the game forms $\langle \mathcal{F}, e_v\rangle$ and $\langle \mathcal{F}^d, v\rangle$ are the same. The lemma follows. □

1.5.8 . Proof of Corollary 1.14

Proof. We let $D := \{\varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \mid \sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}},\ \sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}\} \subseteq \mathcal{D}(\mathsf{O})$. Consider the standard deterministic game form $\mathcal{F}'$ defined by $\mathcal{F}' := \langle \Sigma_{\mathsf{A}}, \Sigma_{\mathsf{B}}, D, \varrho'\rangle$ where, for all $\sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}}$ and $\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}$, we have $\varrho'(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) := \varrho(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \in D$. Furthermore, considering $d : D \to \mathcal{D}(\mathsf{O})$ the identity function, one can realize that $\mathcal{F} = (\mathcal{F}')^d$. The result then follows from Proposition 1.12 applied to $\mathcal{F}'$ and Lemma 1.13 to transfer the result from $\mathcal{F}'$ to $\mathcal{F} = (\mathcal{F}')^d$. □

1.5.9 . Proof of Proposition 1.15

Proof. First, let us show that, for all Borel sets $B \in \mathsf{Borel}(\mathsf{K})$, we have $(\mathsf{col}^\omega)^{-1}[B] \in \mathsf{Borel}(Q)$. Consider any open set $B = \cup_{\pi \in A} \mathsf{Cyl}(\pi)$ for some $A \subseteq \mathsf{K}^*$. We have:

$$\mathsf{col}^{-1}[B] = \bigcup_{\pi \in A} \bigcup_{\rho \in (\mathsf{col}^+)^{-1}[\pi]} \mathsf{Cyl}(\rho) \in \mathsf{Borel}(Q)$$

Furthermore, for all Borel sets $B \in \mathsf{Borel}(\mathsf{K})$, we have $(\mathsf{col}^\omega)^{-1}[\mathsf{K}^\omega \setminus B] = Q^\omega \setminus (\mathsf{col}^\omega)^{-1}[B]$. In addition, for all $(B_n)_{n\in\mathbb{N}} \in (\mathsf{Borel}(\mathsf{K}))^{\mathbb{N}}$, we have $(\mathsf{col}^\omega)^{-1}[\cup_{n\in\mathbb{N}} B_n] = \cup_{n\in\mathbb{N}} (\mathsf{col}^\omega)^{-1}[B_n]$. It follows that, for all Borel sets $B \in \mathsf{Borel}(\mathsf{K})$, we have $(\mathsf{col}^\omega)^{-1}[B] \in \mathsf{Borel}$.

Then, for all $\alpha \in [0,1]$, we have:

$$(f_{\mathcal{C}})^{-1}[[0,\alpha]] = \bigcup_{q \in Q_{\mathsf{s}}, \mathsf{val}(q) \leq \alpha} \bigcup_{\pi \in Q_{\mathsf{ns}}^*} \mathsf{Cyl}(\pi \cdot q) \cup (\mathsf{col}^\omega)^{-1}[f^{-1}[0,\alpha]] \in \mathsf{Borel}(Q)$$

since f is measurable. Therefore, the function $f_{\mathcal{C}}$ is measurable. In addition, by Lemma 1.4, for all $\rho \in Q^*$, we have $(f_{\mathcal{C}})^\rho$ measurable. □

1.5.10 . Proof of Lemma 1.16

Proof. We prove the result when $\mathsf{C} = \mathsf{A}$, the other case being analogous. We denote by $\mathcal{C}^{S,\mathsf{A}}$ the arena underlying the game $\mathcal{G}^{S,\mathsf{A}}$. Clearly, for all states $q \in Q_{\mathsf{s}} \cup S$, we have $\chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{A}](q) = \chi_{\mathcal{G}}[\mathsf{A}]$. Consider some state $q \in Q_{\mathsf{ns}} \setminus S$.

Consider a Player-A strategy $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$. This strategy can be seen as a strategy in the game $\mathcal{G}^{S,\mathsf{A}}$, the game ending as soon as the set S is reached. Consider now any Player-B strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}^{S,\mathsf{A}}}$ in the arena $\mathcal{C}^{S,\mathsf{A}}$. The games $\mathcal{G}$ and $\mathcal{G}^{S,\mathsf{A}}$ coincide on $(Q \setminus S)^*$ and $(Q \setminus S)^\omega$. Let $\varepsilon > 0$ and let us define a Player-B strategy $\mathsf{s}_{\mathsf{B}}^{\varepsilon} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$ in the game $\mathcal{G}$ that coincides with the strategy $\mathsf{s_B}$ on $(Q \setminus S)^*$ and that plays a ε-optimal strategy — against the strategy $\mathsf{s_A}$ —- as soon as a state in S is reached. That is, the expected value of f given that the state q' is eventually reached is at most $\chi_{\mathcal{G}}[\mathsf{A}](q') + \varepsilon$. Note that this holds because the objective is prefix-independent: it does not matter the sequence of states seen before reaching q'. Formallu, we have the following inequality:

$$\mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s}_{\mathsf{B}}^{\varepsilon}}[f_{\mathcal{C}} \cdot \mathbb{1}_{(Q\setminus S)^* \cdot q'}] \leq (\chi_{\mathcal{G}}[\mathsf{A}](q') + \varepsilon) \cdot \mathbb{P}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s}_{\mathsf{B}}^{\varepsilon}}[(Q \setminus S)^* \cdot q']$$

As this holds for all $q' \in S$, we obtain:

$$\sum_{q' \in S} \mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s}_{\mathsf{B}}^{\varepsilon}}[f_{\mathcal{C}} \cdot \mathbb{1}_{(Q\setminus S)^* \cdot q'}] \leq \varepsilon + \sum_{q' \in S} \chi_{\mathcal{G}}[\mathsf{A}](q') \cdot \mathbb{P}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s}_{\mathsf{B}}^{\varepsilon}}[(Q \setminus S)^* \cdot q']$$

Hence, we have:

$$\begin{aligned}
\mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^{S,\mathsf{A}},q}[f_{\mathcal{C}^{S,\mathsf{A}}}] &= \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^{S,\mathsf{A}},q}[f_{\mathcal{C}^{S,\mathsf{A}}} \cdot \mathbb{1}_{(Q\setminus S)^\omega}] + \sum_{q' \in S} \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^{S,\mathsf{A}},q'}[f_{\mathcal{C}^{S,\mathsf{A}}} \cdot \mathbb{1}_{(Q\setminus S)^* \cdot q'}] \\
&= \mathbb{E}^{\mathsf{s_A},\mathsf{s}^\varepsilon_\mathsf{B}}_{\mathcal{C},q}[f_\mathcal{C} \cdot \mathbb{1}_{(Q\setminus S)^\omega}] + \sum_{q' \in S} \mathbb{P}^{\mathsf{s_A},\mathsf{s}^\varepsilon_\mathsf{B}}_{\mathcal{C},q}[(Q \setminus S)^* \cdot q'] \cdot \chi_\mathcal{G}[\mathsf{A}](q') \\
&\geq \mathbb{E}^{\mathsf{s_A},\mathsf{s}^\varepsilon_\mathsf{B}}_{\mathcal{C},q}[f_\mathcal{C} \cdot \mathbb{1}_{(Q\setminus S)^\omega}] + \sum_{q' \in S} \mathbb{E}^{\mathsf{s_A},\mathsf{s}^\varepsilon_\mathsf{B}}_{\mathcal{C},q}[f_\mathcal{C} \cap \mathbb{1}_{(Q\setminus S)^* \cdot q'}] - \varepsilon \\
&\geq \mathbb{E}^{\mathsf{s_A},\mathsf{s}^\varepsilon_\mathsf{B}}_{\mathcal{C},q}[f_\mathcal{C}] - \varepsilon \geq \chi_\mathcal{G}[\mathsf{s_A}](q') - \varepsilon
\end{aligned}$$

As this holds for all Player-B strategies $\mathsf{s_B} \in \mathsf{S}^{\mathcal{C}^{S,\mathsf{A}}}_\mathsf{B}$, it follows that $\chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{s_A}](q) \geq \chi_\mathcal{G}[\mathsf{s_A}](q) - \varepsilon$. As this holds for all $\varepsilon > 0$, we have $\chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{A}](q) \geq \chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{s_A}](q) \geq \chi_\mathcal{G}[\mathsf{s_A}](q)$. Since this holds for all Player-A strategies $\mathsf{s_A} \in \mathsf{S}^\mathcal{C}_\mathsf{A}$, we have $\chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{A}](q) \geq \chi_\mathcal{G}[\mathsf{A}](q)$. Furthermore, for an optimal Player-A strategy $\mathsf{s_A}$ in $\mathcal{G}$ from q, we have $\chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{s_A}](q) \geq \chi_\mathcal{G}[\mathsf{A}](q)$.

Let us now show the other inequality: $\chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{A}](q) \leq \chi_\mathcal{G}[\mathsf{A}](q)$. With what we have shown above, this will prove that the values of all states are the same and that if there is an optimal Player-A strategy in $\mathcal{G}$, then there is also one in $\mathcal{G}^{S,\mathsf{A}}$. We proceed very similarly than for the other inequality. Let $\varepsilon > 0$ and $\mathsf{s_A} \in \mathsf{S}^{\mathcal{C}^{S,\mathsf{A}}}_\mathsf{A}$ be a Player-A strategy such that $\chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{s_A}](q) \geq \chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{A}](q) - \varepsilon$. Consider a Player-$\mathsf{A}$ strategy $\mathsf{s}^\varepsilon_\mathsf{A} \in \mathsf{S}^\mathcal{C}_\mathsf{A}$ in the game $\mathcal{G}$ that coincides with the strategy $\mathsf{s_A}$ on $(Q \setminus S)^*$ and that plays a ε-optimal strategy as soon as a state in S is reached. Consider then any Player-B strategy $\mathsf{s_B} \in \mathsf{S}^\mathcal{C}_\mathsf{B}$, that can also be seen as a strategy in $\mathcal{C}^{S,\mathsf{A}}$. For all states $q' \in S$, with strategy $\mathsf{s}^\varepsilon_\mathsf{A}$ and $\mathsf{s_B}$, from q, we have that the expected value of f given that q' is eventually reached is at least $\chi_\mathcal{G}[\mathsf{A}](q') - \varepsilon$, that is:

$$\mathbb{E}^{\mathsf{s}^\varepsilon_\mathsf{A},\mathsf{s_B}}_{\mathcal{C},q}[f_\mathcal{C} \cdot \mathbb{1}_{(Q\setminus S)^* \cdot q'}] \geq (\chi_\mathcal{G}[\mathsf{A}](q') - \varepsilon) \cdot \mathbb{P}^{\mathsf{s}^\varepsilon_\mathsf{A},\mathsf{s_B}}_{\mathcal{C},q}[(Q \setminus S)^* \cdot q']$$

As this holds for all states $q' \in S$, we have :

$$\sum_{q' \in S} \mathbb{E}^{\mathsf{s}^\varepsilon_\mathsf{A},\mathsf{s_B}}_{\mathcal{C},q}[f_\mathcal{C} \cdot \mathbb{1}_{(Q\setminus S)^* \cdot q'}] \geq \sum_{q' \in S} \chi_\mathcal{G}[\mathsf{A}](q') \cdot \mathbb{P}^{\mathsf{s}^\varepsilon_\mathsf{A},\mathsf{s_B}}_{\mathcal{C},q}[(Q \setminus S)^* \cdot q'] - \varepsilon$$

Hence, we have:

$$\begin{aligned}
\mathbb{E}^{\mathsf{s}^\varepsilon_\mathsf{A},\mathsf{s_B}}_{\mathcal{C},q}[f_\mathcal{C}] &= \mathbb{E}^{\mathsf{s}^\varepsilon_\mathsf{A},\mathsf{s_B}}_{\mathcal{C},q}[f_\mathcal{C} \cdot \mathbb{1}_{(Q\setminus S)^\omega}] + \sum_{q' \in S} \mathbb{E}^{\mathsf{s}^\varepsilon_\mathsf{A},\mathsf{s_B}}_{\mathcal{C},q}[f_\mathcal{C} \cdot \mathbb{1}_{(Q\setminus S)^* \cdot q'}] \\
&\geq \mathbb{E}^{\mathsf{s}^\varepsilon_\mathsf{A},\mathsf{s_B}}_{\mathcal{C},q}[f_\mathcal{C} \cdot \mathbb{1}_{(Q\setminus S)^\omega}] + \sum_{q' \in S} \mathbb{P}^{\mathsf{s}^\varepsilon_\mathsf{A},\mathsf{s_B}}_{\mathcal{C},q}[(Q \setminus S)^* \cdot q'] \cdot \chi_\mathcal{G}[\mathsf{A}](q') - \varepsilon \\
&= \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^{S,\mathsf{A}},q}[f_{\mathcal{C}^{S,\mathsf{A}}} \cdot \mathbb{1}_{(Q\setminus S)^\omega}] + \sum_{q' \in S} \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^{S,\mathsf{A}},q}[f_{\mathcal{C}^{S,\mathsf{A}}} \cdot \mathbb{1}_{(Q\setminus S)^* \cdot q'}] - \varepsilon \\
&= \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^{S,\mathsf{A}},q}[f_{\mathcal{C}^{S,\mathsf{A}}}] - \varepsilon \geq \chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{A}](q') - 2\varepsilon
\end{aligned}$$

As this holds for all Player-B strategies $\mathsf{s}_\mathsf{B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$, it follows that $\chi_\mathcal{G}[\mathsf{s}_\mathsf{A}](q) \geq \chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{A}](q) - 2\varepsilon$. As this holds for all $\varepsilon > 0$, we have $\chi_\mathcal{G}[\mathsf{A}](q) \geq \chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{A}](q)$.

This also holds for $\mathsf{C} = \mathsf{B}$. Hence, if the game $\mathcal{G}$ has a value, we have $\chi_{\mathcal{G}^{S,\mathsf{A}}}[\mathsf{A}] = \chi_\mathcal{G}[\mathsf{A}] = \chi_\mathcal{G}[\mathsf{B}] = \chi_{\mathcal{G}^{S,\mathsf{B}}}[\mathsf{B}]$ where the games $\mathcal{G}^{S,\mathsf{A}}$ and $\mathcal{G}^{S,\mathsf{B}}$ are the same. Therefore, the game $\mathcal{G}^{S,\mathsf{A}}$ has also a value. $\square$

Part I

General results with arbitrary bounded payoff functions

In this first part, we study concurrent games while making minor assumptions on the local interactions and payoff functions involved. This is in sharp contrast with what we do in Part II, where we consider only concurrent parity games; and with what we do in Part III where we define restrictions on the local interactions occurring in concurrent games.

Among the results that we show in this part, we would like to mention here two of them that are essential to this dissertation and that we believe are important results on concurrent games.

The first one deals with, arguably, the most important result on concurrent two-player antagonistic games: Martin's determinacy results of Blackwell games. The original version of this theorem [12] states that all standard concurrent games with finite local interactions have a value. The main focus of Chapter 2 is the proof of a generalization of this result to arbitrary (non necessarily standard) games. The idea is that, by von Neuman's minimax theorem [4], all standard finite game forms are valuable. This is actually the assumption that Martin uses in his proof[2]. We show that all (arbitrary) concurrent games with valuable local interactions have a value. This is stated as part of Theorem 2.3 in Chapter 2. Note that Theorem 2.3 states other results than the one described above, without assuming that the local interactions are valuable.

The second result that we would like to highlight is novel and essential to this dissertation, though it is much easier to prove than the previous one. It is stated as Theorem 3.12 in Chapter 3, with Corollaries 3.14 and 3.16 being relevant special cases. Informally, this theorem states a sufficient condition for the value of a Player-A strategy to be greater than or equal to some threshold, and symmetrically for Player B. This theorem or one of its corollaries are used several times in this dissertation, in Chapters 4, 5, 7 and 8.

As mentioned above, in Chapter 2 we focus on the well-known determinacy result for Blackwell games by Martin. Then, in Chapter 3, we focus on subgame (ε-)optimal strategies (notion to be defined).

[2] In fact, Martin states that the result still holds even if the local interactions are such that one set of actions is finite, while the other is countable. Even with this weaker assumption, the local interactions are still valuable.

2 - Blackwell determinacy

In 1975, Martin showed the determinacy of Borel games: in all deterministic turn-based games with win/lose objectives, one of the players has a winning strategy [8]. Then, in 1998, Martin used this result to show that all standard concurrent games where each (standard) local interaction is finite[1] have a value [12]. Note that this holds for arbitrary payoff functions (recall that all the payoff functions we consider are measurable and into $[0, 1]$). This is a central result in the theory of standard concurrent games since the assumptions are quite mild while the conclusion, i.e. that all games have values, is very useful when studying concurrent games. In this chapter, we extend Martin's result and obtain a (slightly) more general one. The additional power is invoked several times in this dissertation, in places where Martin's original result would not suffice. Informally, this extension follows two directions:

- (1): By closely examining the construction that Martin uses to prove the result, we show that almost-optimal strategies (i.e. ε-optimal strategies, for all $\varepsilon > 0$) can be found among specific subsets of strategies. This is proved regardless of the local interactions involved (i.e. they need not be valuable).
- (2): We show that as soon as all local interactions are valuable, the game has a value.

More formally, we show the following. In an arbitrary game $\mathcal{G}$:

- (1): almost-optimal strategies (i.e. ε-optimal for all $\varepsilon > 0$) can be chosen among specific subsets of strategies, namely:
  - (1.a): first, without any additional assumption, they can be found among strategies generated by subsets of GF-strategies that supremize the corresponding local interactions;
  - (1.b): second, if $\mathcal{G}$ is win/lose, under a specific condition on the coloring function, it holds that we can further reduce the subset of strategies to consider only the ones that depend on the sequence of colors seen and on the current state of the game, not on the exact sequence of states seen. (This amounts to some kind of uniformization of strategies, see for instance [53] in the context of turn-based games and winning strategies.)

[1]In fact, Martin mentioned that this also holds if either of the player action set is finite while the other one is countable.

- (2): in addition, if $\mathcal{G}$ is valuable, i.e. if all local interactions in $\mathcal{G}$ are valuable, then $\mathcal{G}$ has a value.

This is stated formally in Section 2.2 as Theorem 2.3.

The proof. Result (2) generalizes Martin's result in the context of arbitrary concurrent games, and not just standard games. That is, we identify a sufficient condition on local interactions, namely being valuable, for concurrent games whose local interactions satisfy this condition to have a value. As discussed in Chapter 6, this condition on local interactions is also somehow necessary, see Proposition 6.1.

One of the key ideas of Martin's proof is to derive a spoiler/verifier style turn-based game $\mathcal{G}_{\mathsf{tb}}$ from a concurrent game $\mathcal{G}$. By Borel determinacy [8], from all states of the game $\mathcal{G}_{\mathsf{tb}}$, either of the players has a winning strategy. These winning strategies are translated into almost-optimal strategies in $\mathcal{G}$. If we only wanted to prove result (2), we could use this idea as is and adapt it straightforwardly to our framework. However, because we also want to prove result $(1.a)$, even when the local interactions are not valuable, more work is required: instead of only defining a unique spoiler/verifier turn-based game, we define infinitely many. Up to that change, our proof of result $(1.a)$ follows the footsteps of Martin's proof. In particular, we also use Borel determinacy. Moreover, result (2) is a direct consequence of lemmas dedicated to prove result $(1.a)$. Finally, to prove result $(1.b)$, we need to extract additional properties from the turn-based games mentioned above.

Furthermore, we prove this new version of Blackwell determinacy with elementary arguments. It is in particular the case for the intermediate results that we show on stochastic trees. These intermediate results, that we prove from elementary definitions in probability theory, are existing results on martingales. We discuss them in Section 2.3. More generally, we have also added intermediate lemmas and examples to explain and illustrate the ideas behind the proofs.

Consequences. As mentioned above, Theorem 2.3 extends Martin's result in two directions. The benefit of the second direction (stated as result (2)) is rather straightforward since it extends the set of games to which Martin's result can be applied.

Let us now consider the first extension. It contains two results: $(1.a)$ and $(1.b)$. We will give several applications of these results in this dissertation.

Result $(1.a)$ is used in Chapter 3. In this chapter, we show that subgame almost-optimal strategies (notion defined in Definition 3.3) exists and, with result $(1.a)$, we show that they can be found among a specific subsets of strategies (see Theorem 3.1). Furthermore, Chapter 6 is entirely dedicated to applications of result $(1.a)$. These consist in showing that, if the local interactions occurring in a concurrent game $\mathcal{G}$ belong to a specific set of game forms, then the whole concurrent game $\mathcal{G}$ enjoys nice properties (see Theorems 6.4, 6.11, 6.15). In particular, Borel determinacy that Martin proved in [8] is a logical consequence of Theorem 6.4. However, as stated above, we use Borel determinacy to prove Theorem 2.3, hence we do **not** provide a new proof of Borel determinacy. However, note that Borel determinacy is not a logical consequence of Martin's original result [12]. All these applications do not use result (1.b), they sometimes use result (2).

Consider now result (1.b). In Section 2.5, we use this result to prove the following. Consider a standard win/lose game with finitely many actions at each state for both players. Consider also action strategies, i.e. strategies that may also depend on the actions played by the players, which are formally defined in that section. Then, the values that can be achieved with action strategies is the same as the value with the strategies we have considered so far. The latter are called state strategies: they depend only on the history of states. This application does not use result (1.a), however it uses result (2). Finally, in this section, we also exhibit a standard finite game satisfying the following properties. The values achieved by action and state strategies are the same, which is a consequence of the above-mentioned result. However, there is an optimal strategy among action strategies, while there is none among state strategies. This is stated in Proposition 2.21.

The work presented in this chapter is not published yet.

2.1 Martin's results

In this section, we recall two of Martin's theorems. In the original papers [8, 12], the formalism which is used is quite different from the one used in this dissertation. Specifically, Martin uses the notion of game trees without considering an underlying graph (i.e. with an explicit set of states). However, the theorems we state here are equivalent to the ones showed in [8, 12].

First, Martin proved the existence of winning strategies in deterministic turn-based win/lose games [8]. This is also known as Borel determinacy, and it is stated below in Theorem 2.1.

Theorem 2.1 (Borel determinacy [8]). *Consider a turn-based deterministic win/lose game $\mathcal{G}$ without stopping states. For all $q \in Q$, either Player A or Player B has a winning (deterministic) strategy from state q. This holds even if the set of states is not countable.*

Furthermore, standard concurrent games with specific local interactions have a value [12]. This is also known as Blackwell determinacy, and it is stated below in Theorem 2.2.

Theorem 2.2 (Blackwell determinacy [12]). *Consider a standard concurrent game $\mathcal{G}$. Assume that, for all $q \in Q$, in the game form $\mathsf{F}(q)$ both action sets are countable and at least one of them is finite. Then, the game $\mathcal{G}$ has a value.*

2.2 Blackwell determinacy: a new version

Before formally stating another version of Theorem 2.2, we need to introduce two definitions. Below in Definition 2.1, we consider the notion of strategies that only depend on the colors seen. This is also defined for valuations of finite sequences of states.

Definition 2.1 (Color-Uniform strategies and valuations). *Consider an arbitrary concurrent arena $\mathcal{C}$. A* uniformizing pair *is a pair (U, m) where U is a non-empty set and $m : Q \to \mathsf{U}$ maps each state to an element in U. (In particular, the pair $(\mathsf{K}, \mathsf{col})$ is a uniformizing pair[2].) For such a pair, we say that two finite sequences of states $\rho, \rho' \in Q^+$ are* (U, m)-equivalent *if $\rho_{\mathsf{lt}} = \rho'_{\mathsf{lt}}$ and $m^+(\rho) = m^+(\rho')$.*

A function $g : Q^+ \to X$ mapping finite sequences of states to any non-empty set X is said to be (U, m)-uniform *if, for all pairs $\rho, \rho' \in Q^+$ of* (U, m)-equivalent *paths, we have $g(\rho) = g(\rho')$. They can be seen as maps $\mathsf{U}^* \times Q \to X$.*

In Definition 2.2 below, we define the notion of coloring function with a finite representative: the functions such that each color has only finitely many preimages w.r.t. that function.

Definition 2.2 (Coloring function with a finite representative). *Consider an arbitrary concurrent arena $\mathcal{C}$. We say that $(\mathsf{K}, \mathsf{col})$ has a* finite representative *in Q if, for all $k \in \mathsf{K}$, the set $\mathsf{col}^{-1}[\{k\}] \subseteq Q$ is finite[3].*

We can now state our main theorem of this chapter: the new version of Blackwell determinacy.

Theorem 2.3. *Consider an arbitrary concurrent game $\mathcal{G}$. Let $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$ be a Player. Consider a collection $(S^q_{\mathsf{C}})_{q \in Q}$ of sets of Player-C GF-strategies that supremize the game $\mathcal{G}$ w.r.t. Player C. Then:*

- *(1) For all $\varepsilon > 0$:*
  - *(1.a) There is a Player-C strategy $\mathsf{s}^{\varepsilon}_{\mathsf{C}} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{C}}$ generated by $(S^q_{\mathsf{C}})_{q \in Q}$ that is ε-optimal.*
  - *(1.b) If we additionally assume that $(\mathsf{K}, \mathsf{col})$ has a finite representative in Q and that $\mathcal{G}$ is win/lose, then the strategy $\mathsf{s}^{\varepsilon}_{\mathsf{C}}$ above can be chosen $(\mathsf{K}, \mathsf{col})$-uniform.*
- *(2) If the game $\mathcal{G}$ is valuable, then it has a value: $\chi_{\mathcal{G}}[\mathsf{A}] = \chi_{\mathcal{G}}[\mathsf{B}] : Q \to [0, 1]$.*

Remark 2.1. *It is quite straightforward that Theorem 2.3 implies Theorem 2.2 since, by Proposition 1.12, any standard game form with at least one*

[2]To properly fit this definition, the function col needs to be defined on Q_{ns}.

[3]This does not imply that the set Q is finite if the set of colors K is infinite.

action set that is finite is valuable. In fact, Theorem 2.3 also implies Theorem 2.1, see Statement 6.1. However, note that Theorem 2.1 is used to prove Theorem 2.3 (but Theorem 2.2 is not).

Before considering the proof of Theorem 2.3, we need to make a detour via stochastic trees.

2.3 A result on stochastic trees

In this section, we establish results in stochastic trees that will be used in this chapter to prove Theorem 2.3. We will also use them in Chapter 3. We would like to mention that the main results stated in this section (that is Proposition 2.6 and Proposition 2.9) already exist (see for instance [54, Thm. 3.17]) as the underlying object we consider – namely, non-decreasing valuation — in fact correspond to sub-martingales. However, the arguments we use in this section are elementary in the sense that we do not use at all results on martingales. This has the benefit of being readable by someone who is not familiar with this notion.

Before considering the few definitions we need to properly state and prove the results of interest of this section, we first consider superior and inferior limit functions (given a valuation of finite sequences of states) in stochastic trees and realize that they are measurable.

Proposition 2.4 (Proof 2.7.1)**.** *Consider a stochastic tree $\mathcal{T}$ and a valuation $v : Q^+ \to [0,1]$ of the finite sequences of states of $\mathcal{T}$. Consider the superior (resp. inferior) limit function $\mathsf{limsup}_v : Q^\omega \to [0,1]$ defined by, for all $\rho \in Q^\omega$: $\mathsf{limsup}_v(\rho) := \mathsf{limsup}\ (v(\rho_{\leq n}))_{n \in \mathbb{N}} \in [0,1]$ (resp. $\mathsf{liminf}_v(\rho) := \liminf\ (v(\rho_{\leq n}))_{n \in \mathbb{N}} \in [0,1]$). Then, this superior (resp. inferior) limit function is measurable.*

Let us also recall that comparing two measurable functions yields a measurable set.

Proposition 2.5 (Proof 2.7.2)**.** *Consider a non-empty set Q and two measurable functions $f, g : Q^\omega \to [0,1]$. For all $\bowtie \in \{\leq, <, \geq, >, =, \neq\}$, the event $\{f \bowtie g\} := \{\rho \in Q^\omega \mid f(\rho) \bowtie g(\rho)\} \subseteq Q^\omega$ is Borel: $\{f \bowtie g\} \in \mathsf{Borel}(Q)$.*

2.3.1 . Comparing superior and inferior limits

Consider a stochastic tree $\mathcal{T}$ and a valuation of finite sequences of states $v : Q^+ \to [0,1]$. By definition, for all infinite paths $\rho \in Q^\omega$, the superior limit w.r.t. v of ρ is greater then or equal to the inferior limit w.r.t. v of ρ: $\mathsf{limsup}_v(\rho) \geq \mathsf{liminf}_v(\rho)$. Hence, the expected value of the superior limit limsup_v is greater than or equal to the the expected value of the inferior limit liminf_v. Without any assumption on v, the difference between these expected values may be equal to 1: that is it could be that almost-surely, the superior limit of an infinite path is equal to 1 while the inferior limit is almost-surely equal to 0. The goal of this subsection is to show that, under a specific condition on v, the difference between the superior and inferior limit is null. More specifically, we show that if the expected value of the valuation v does not decrease in any single step, then the superior and inferior limits are equal almost-surely.

Let us define formally the notion of non-decreasing valuation.

Definition 2.3 (Non-decreasing valuation in stochastic trees)**.** *Consider a stochastic tree $\mathcal{T} = \langle Q, \mathbb{P} \rangle$, a finite path $\pi \in Q^+$ and a valuation $v : Q^* \to [0,1]$. It is* non-decreasing *from π if, for all $\rho \in Q^*$, we have:*

$$v(\rho) \leq \sum_{q \in Q} \mathbb{P}_{\pi \cdot \rho}(q) \cdot v(\rho \cdot q)$$

Moreover, a valuation $v : Q^+ \to [0,1]$ is said to be non-decreasing *if, for all $q \in Q$, the valuation $v^q : Q^* \to [0,1]$ is non-decreasing from q.*

Remark 2.2. *By definition, if a valuation is non-decreasing from some finite path $\pi \in Q^+$, then for all $\rho \in Q^*$, the valuation v^ρ is also non-decreasing from $\pi \cdot \rho$.*

With a non-decreasing valuation, infinite paths have a limit (i.e. the inferior equals the superior limit) almost-surely, as stated below.

Proposition 2.6. *Consider a stochastic tree $\mathcal{T}$ and a valuation $v : Q^* \to [0,1]$ non-decreasing from some $\pi \in Q^+$. Then, we have $\mathbb{P}_\pi(\mathsf{liminf}_v < \mathsf{limsup}_v) = 0$*

Before proving this proposition, let us first show two intermediate results.

Lemma 2.7. *Consider a stochastic tree $\mathcal{T}$ and a valuation $v : Q^* \to [0,1]$ non-decreasing from some $\pi \in Q^+$. Then, for all $i \in \mathbb{N}$, we have: $v(\epsilon) \leq \sum_{\rho \in Q^i} \mathbb{P}_\pi(\rho) \cdot v(\rho)$.*

Proof. We show this property by induction on i. This straightforwardly holds for $i = 0$. Assume now that this property holds for some $i \in \mathbb{N}$. We have:

$$\begin{aligned}
\sum_{\rho \in Q^{i+1}} \mathbb{P}_\pi(\rho) \cdot v(\rho) &= \sum_{\rho \in Q^i} \sum_{q \in Q} \mathbb{P}_\pi(\rho) \cdot \mathbb{P}_{\pi \cdot \rho}(q) \cdot v(\cdot \rho \cdot q) \\
&= \sum_{\rho \in Q^i} \mathbb{P}_\pi(\rho) \cdot \sum_{q \in Q} \mathbb{P}_{\pi \cdot \rho}(q) \cdot v(\rho \cdot q) \\
&= \sum_{\rho \in Q^i} \mathbb{P}_\pi(\rho) \cdot v(\rho) \geq v(\epsilon)
\end{aligned}$$

Hence, the property holds for all $i \in \mathbb{N}$. □

Lemma 2.8. *Consider a stochastic tree $\mathcal{T}$, a valuation $v : Q^* \to [0,1]$ non-decreasing from some $\pi \in Q^+$. Let $u := v(\epsilon)$ and let $0 \leq u' < u$ be a value less* *than u and $E := \{\rho \in Q^* \mid v(\rho) \leq u'\}$ be the set of finite paths whose values w.r.t. the valuation v are less than or equal to u'. Then, $\mathbb{P}_\pi[E] \leq \frac{1-u}{1-u'} < 1$.*

Proof. First, note that $\epsilon \notin E$. Let $A := \cup_{\rho \in E} \rho \cdot Q^*$ and $B := Q^* \setminus A$. Consider the set $D \subseteq E$ of finite sequences of states in E with no strict prefix in E. This set ensures that:

- $A = \cup_{\rho \in D} \rho \cdot Q^*$;
- for all $\rho, \rho' \in D$, we have $\rho \cdot Q^* \cap \rho' \cdot Q^* = \emptyset$.

Now, let us define a new valuation $w : Q^* \to [0,1]$ such that, for all $\rho \in Q^*$:

$$w(\rho) := \begin{cases} v(\rho) & \text{if } \rho \in B \\ v(\rho') & \text{if } \rho \in \rho' \cdot Q^* \text{ for some } \rho' \in D \end{cases}$$

Let us show that the valuation w is non-decreasing from π. Consider $\rho \in Q^*$. If $\rho \in B$, then we have $w(\rho) = v(\rho)$ and for all $q \in Q$, $w(\rho \cdot q) = v(\rho \cdot q)$. Hence, we do have $w(\rho) \leq \sum_{q \in Q} \mathbb{P}_{\pi \cdot \rho}(q) \cdot w(\rho \cdot q)$ since v is non-decreasing from π. Furthermore, if $\rho \in \rho' \cdot Q^*$ for some $\rho' \in D$, then $w(\rho) = v(\rho')$ and for all $q \in Q$, we have $w(\rho \cdot q) = v(\rho') = w(\rho)$. Hence, we have $w(\rho) = \sum_{q \in Q} \mathbb{P}_{\pi \cdot \rho}(q) \cdot w(\rho \cdot q)$. Overall, the valuation w is non-decreasing from π.

Consider now some $n \in \mathbb{N}$. For all $\rho' \in D \cap Q^{\leq n}$, we have $\mathbb{P}_\pi(\rho') = \sum_{\rho \in \rho' \cdot Q^* \cap Q^n} \mathbb{P}_\pi(\rho)$. Hence, by applying Lemma 2.7, since w is non-decreasing from π:

$$\begin{aligned} u = w(\epsilon) &\leq \sum_{\rho \in Q^n} \mathbb{P}_\pi(\rho) \cdot w(\rho) = \sum_{\rho \in A \cap Q^n} \mathbb{P}_\pi(\rho) \cdot w(\rho) + \sum_{\rho \in B \cap Q^n} \mathbb{P}_\pi(\rho) \cdot w(\rho) \\ &= \sum_{\rho' \in D \cap Q^{\leq n}} \sum_{\rho \in \rho' \cdot Q^* \cap Q^n} \mathbb{P}_\pi(\rho) \cdot w(\rho) + \sum_{\rho \in B \cap Q^n} \mathbb{P}_\pi(\rho) \cdot w(\rho) \\ &= \sum_{\rho' \in D \cap Q^{\leq n}} \sum_{\rho \in \rho' \cdot Q^* \cap Q^n} \mathbb{P}_\pi(\rho) \cdot v(\rho') + \sum_{\rho \in B \cap Q^n} \mathbb{P}_\pi(\rho) \cdot v(\rho) \\ &= \sum_{\rho' \in D \cap Q^{\leq n}} \mathbb{P}_\pi(\rho') \cdot v(\rho') + \sum_{\rho \in B \cap Q^n} \mathbb{P}_\pi(\rho) \cdot v(\rho) \\ &\leq \sum_{\rho' \in D \cap Q^{\leq n}} \mathbb{P}_\pi(\rho') \cdot u' + \sum_{\pi' \in B \cap Q^n} \mathbb{P}_\pi(\pi') \\ &= \mathbb{P}_\pi[D \cap Q^{\leq n}] \cdot u' + 1 - \mathbb{P}_\pi[D \cap Q^{\leq n}] \end{aligned}$$

Hence, denoting $p_n := \mathbb{P}_\pi[Q^{\leq n} \cap D]$, we obtain: $u \leq p_n \cdot u' + 1 - p_n$. That is, $p_n \leq \frac{1-u}{1-u'} < 1$. Since this holds for all $n \in \mathbb{N}$ and $\lim_{n \to \infty} p_n = \mathbb{P}_\pi[Q^* \cap D] = \mathbb{P}_\pi[D] = \mathbb{P}_\pi[E]$ by continuity of $\mathbb{P}$ and since $\mathbb{P}_\pi[E] = \mathbb{P}_\pi[\cup_{\rho \in E} \mathsf{Cyl}(\rho)]$, we get $\mathbb{P}_\pi[E] \leq \frac{1-u}{1-u'}$. $\square$

We can now proceed to the proof of Proposition 2.6.

Proof. Let $p, q \in \mathbb{Q} \cap [0,1]$ be such that $q < p$. Let $d := p - q > 0$ and $p' := p - \frac{d}{4}$ and $q' := q + \frac{d}{4}$. Finally, let $x := \frac{1-p'}{1-q'} < 1$. Let $V_{\geq p'} := \{\rho \in Q^* \mid v(\rho) \geq p'\}$ and $V_{\leq q'} := \{\rho \in Q^* \mid v(\rho) \leq q'\}$. Now, we have:

$$\{p \leq \mathsf{limsup}_v\} \cap \{\mathsf{liminf}_v \leq q\} \subseteq \bigcap_{k \in \mathbb{N}} (V_{\geq p'} \cdot V_{\leq q'})^k$$

By Lemma 2.8 and Remark 2.2, for all finite paths $\rho \in V_{\geq p'}$, we have $\mathbb{P}_\rho[V_{\leq q'}] \leq \frac{1-v(\rho)}{1-q'} \leq x$. Hence, $\mathbb{P}_\pi[V_{\geq p'} \cdot V_{\leq q'}] \leq x$. It follows that, for all $k \in \mathbb{N}$, we have $\mathbb{P}_\pi[(V_{\geq p'} \cdot V_{\leq q'})^k] \leq x^k$. Since $x < 1$, we have $\mathbb{P}_\pi[\cap_{k \in \mathbb{N}} (V_{\geq p'} \cdot V_{\leq q'})^k] = \lim_{k \in \mathbb{N}} x^k = 0$. It follows that $\mathbb{P}[\{p \leq \mathsf{limsup}_v\} \cap \{\mathsf{liminf}_v \leq q\}] = 0$. As this holds for all $p, q \in \mathbb{Q} \cap [0,1]$ such that $q < p$, it follows that $\mathbb{P}_\pi[\mathsf{liminf}_v < \mathsf{limsup}_v] = 0$. □

2.3.2 . Expected value of the superior limit

In this subsection, we focus on non-decreasing valuations. We show that for such valuations, the expected value of the superior limit (which is almost-surely equal to inferior limit, recall Proposition 2.6 above) is at least the value of the starting state.

Proposition 2.9. *Consider a stochastic tree $\mathcal{T}$ and a valuation $v : Q^* \to [0,1]$ non-decreasing from some $\pi \in Q^+$. Then:*

$$v(\epsilon) \leq \mathbb{E}_\pi[\mathsf{limsup}_v] = \mathbb{E}_\pi[\mathsf{liminf}_v]$$

In fact, to prove Theorem 2.3, we will only use the inequality. However, we have also stated the equality so that Proposition 2.9 implies straightforwardly Proposition 2.6 (though we use Proposition 2.6 to prove Proposition 2.9) since $\mathsf{liminf}_v \leq \mathsf{limsup}_v$.

Proof. First, the equality is a direct consequence of Proposition 2.6. Then, we let $\mathbb{P} := \mathbb{P}_\pi$, $\mathbb{E} := \mathbb{E}_\pi$, $\mathsf{l_{sup}} := \mathsf{limsup}_v$ and $\mathsf{l_{inf}} := \mathsf{liminf}_v$. For all $j \in \mathbb{N}$ and subsets $I \subseteq \mathbb{R}$, we denote by $V(j, I)$ the open set

$$V(j, I) := \bigcup_{\substack{\rho \in Q^j \\ v(\rho) \in I}} \mathsf{Cyl}(\rho)$$

of paths whose j-th value is in the interval I.

For all $n \in \mathbb{N}$, we consider the function $f_n : Q^\omega \to [0,1]$ such that, for all $\rho \in Q^\omega$, we have:

$$f_n(\rho) := \frac{\lfloor 2^n \cdot \mathsf{l_{sup}}(\rho) \rfloor}{2^n} \in \left\{ \frac{i}{2^n} \mid 0 \leq i \leq 2^n \right\}$$

where $\lfloor \cdot \rfloor : \mathbb{R} \to \mathbb{Z}$ is the floor function, that is, for all $x \in \mathbb{R}$, we have $\lfloor x \rfloor \leq x < \lfloor x \rfloor + 1$. For all $n \in \mathbb{N}$, we have that f_n is a step function (recall Definition 1.2), $f_n \leq \mathsf{l_{sup}}$ and therefore $\mathbb{E}[f_n] \leq \mathbb{E}[\mathsf{l_{sup}}]$.

Now, let $\varepsilon > 0$. Consider some $n \in \mathbb{N}$ such that

$$\frac{2}{2^n} \leq \varepsilon \tag{2.1}$$

For all $a \leq b$, let $L_{a,b} := \{a \leq \mathsf{l_{inf}}\} \cap \{\mathsf{l_{sup}} < b\}$. Consider some $0 \leq i \leq 2^n - 1$. We have $\mathbb{P}(L_{\frac{i+1}{2^n}-\delta, \frac{i+1}{2^n}}) \to_{\delta \to 0} 0$. This allows us to consider a $0 < \delta_i < \frac{1}{4 \cdot 2^n}$ such that $\mathbb{P}(L_{\frac{i+1}{2^n}-\delta_i, \frac{i+1}{2^n}}) < \frac{1}{2^{n+1}(2^n+1)^2}$. We set $\delta := \min_{0 \leq i \leq 2^n-1} \delta_i > 0$ and for all $0 \leq i \leq 2^n - 1$, we let $L_i := L_{\frac{i}{2^n}, \frac{i+1}{2^n}}$ and $L'_i := L_{\frac{i}{2^n}, \frac{i+1}{2^n}-\delta}$. In addition, $L'_{2^n} := L_{2^n} := \{\mathsf{l_{sup}} = 1\}$. These definitions are illustrated in Figure 2.1. With these choices, we have, for all $0 \leq i \leq 2^n$:

$$\mathbb{P}(L_i) \leq \mathbb{P}(L'_i) + \frac{1}{2^{n+1}(2^n+1)^2} \tag{2.2}$$

Furthermore:

$$Q^\omega = \bigcup_{i=0}^{2^n-1} \left(\left\{ \frac{i}{2^n} \leq \mathsf{l_{sup}} \right\} \cap \left\{ \mathsf{l_{sup}} < \frac{i+1}{2^n} \right\} \right) \cup \{\mathsf{l_{sup}} = 1\}$$

Hence, since by Proposition 2.6, almost-surely the superior and inferior limits coincide, we have:

$$1 = \mathbb{P}(Q^\omega) = \sum_{i=0}^{2^n-1} \mathbb{P}\left(\left\{ \frac{i}{2^n} \leq \mathsf{l_{inf}} \right\} \cap \left\{ \mathsf{l_{sup}} < \frac{i+1}{2^n} \right\} \right) + \mathbb{P}\left(\{\mathsf{l_{sup}} = 1\}\right) = \sum_{i=0}^{2^n} \mathbb{P}(L_i)$$

We obtain:

$$\sum_{i=0}^{2^n} \mathbb{P}(L_i) = 1 \tag{2.3}$$

Finally, for all $1 \leq i \leq 2^n - 1$, we consider the subsets $J_i := [\frac{i}{2^n} - \frac{\delta}{2}, \frac{i+1}{2^n} - \frac{\delta}{2}[$, $J_0 := [0, \frac{1}{2^n} - \frac{\delta}{2}[$, and $J_{2^n} := [1 - \frac{\delta}{2}, 1]$. These definitions are also illustrated in Figure 2.1. With these choices, the J_i form a partition of the set $[0,1]$, i.e:

$$[0,1] = \biguplus_{i=0}^{2^n} J_i \tag{2.4}$$

Let $0 \leq i \leq 2^n$. By definition of the inferior and superior limits, we have:

$$L'_i \subseteq \bigcup_{l \in \mathbb{N}} \bigcap_{k \geq l} V(k, J_i)$$

Hence: $\mathbb{P}(L'_i) = \mathbb{P}(L'_i \cap \bigcup_{l \in \mathbb{N}} \bigcap_{k \geq l} V(k, J_i)) = \lim\limits_{l \to \infty} \mathbb{P}(L'_i \cap \bigcap_{k \geq l} V(k, J_i))$ by monotone continuity of the probability. Let us consider some $l_i \in \mathbb{N}$ such that: $\mathbb{P}(L'_i \cap \bigcap_{k \geq l_i} V(k, J_i)) \geq \mathbb{P}(L'_i) - \frac{1}{2^{2n+1}(2^n+1)}$. Now, let $l := \max_{0 \leq i \leq 2^n} l_i$. It follows that, for all $0 \leq i \leq 2^n$:

$$\mathbb{P}(V(l, J_i)) \geq \mathbb{P}\left(L'_i \cap \bigcap_{k \geq l_i} V(k, J_i) \right) \geq \mathbb{P}(L'_i) - \frac{1}{2^{2n+1}(2^n+1)} \tag{2.5}$$

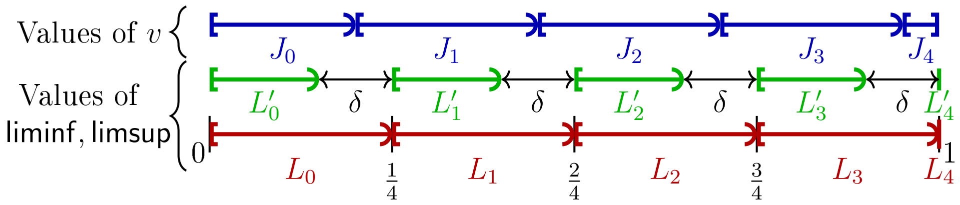

Figure 2.1: An illustration of the definitions of L_i, L'_i and J_i from the proof of Proposition 2.9 in the case where $n = 2$.

Let $0 \leq i \leq 2^n$. We want to establish the following upper bound on $\mathbb{P}(V(l, J_i))$:

$$\mathbb{P}(V(l, J_i)) \leq \mathbb{P}(L'_i) + \frac{1}{2^n(2^n+1)} \tag{2.6}$$

We have:

$$\begin{aligned}
1 &= \sum_{j=0}^{2^n} \mathbb{P}(L_j) && \text{by Equation (2.3)} \\
&\leq \sum_{j=0}^{2^n} \left(\mathbb{P}(L'_j) + \frac{1}{2^{n+1}(2^n+1)^2} \right) = \sum_{j=0}^{2^n} \mathbb{P}(L'_j) + \frac{1}{2^{n+1}(2^n+1)} && \text{by Equation (2.2)} \\
&\leq \mathbb{P}(L'_i) + \sum_{j=0, j\neq i}^{2^n} \left(\mathbb{P}(V(l, J_j)) + \frac{1}{2^{2n+1}(2^n+1)} \right) + \frac{1}{2^{n+1}(2^n+1)} && \text{by Equation (2.5)} \\
&\leq \mathbb{P}(L'_i) + \sum_{j=0, j\neq i}^{2^n} \mathbb{P}(V(l, J_j)) + \frac{1}{2^n(2^n+1)} \\
&= \mathbb{P}(L'_i) + (1 - \mathbb{P}(V(l, J_i))) + \frac{1}{2^n(2^n+1)} && \text{by Equation (2.4)}
\end{aligned}$$

Overall, we do obtain Equation 2.6. Hence, denoting $V_l := v[Q^l] \subseteq [0, 1]$, by

Lemma 2.7 for the first equality, we have:

$$
\begin{aligned}
v(\epsilon) &\leq \sum_{\rho \in Q^l} \mathbb{P}(\rho) \cdot v(\rho) = \sum_{i=0}^{2^n} \sum_{x \in J_i \cap V_l} \mathbb{P}(V(l,\{x\})) \cdot x && \text{by Equation (2.4)} \\
&\leq \sum_{i=0}^{2^n} \sum_{x \in J_i \cap V_l} \mathbb{P}(V(l,\{x\})) \cdot \left(\frac{i+1}{2^n}\right) && \text{by definition of } J_i \\
&= \sum_{i=0}^{2^n} \mathbb{P}(V(l,J_i)) \cdot \left(\frac{i+1}{2^n}\right) && \text{by definition of } V(l,J_i) \\
&= \sum_{i=0}^{2^n} \mathbb{P}(V(l,J_i)) \cdot \frac{i}{2^n} + \frac{1}{2^n} && \text{by Equation (2.4)} \\
&\leq \sum_{i=0}^{2^n} \left(\mathbb{P}(L'_i) + \frac{1}{2^n(2^n+1)}\right) \cdot \frac{i}{2^n} + \frac{1}{2^n} && \text{by Equation (2.6)} \\
&\leq \sum_{i=0}^{2^n} \mathbb{P}(L'_i) \cdot \frac{i}{2^n} + \sum_{i=0}^{2^n} \frac{1}{2^n(2^n+1)} + \frac{1}{2^n} \\
&\leq \sum_{i=0}^{2^n} \mathbb{P}(L_i) \cdot \frac{i}{2^n} + \frac{2}{2^n} && \text{since } L'_i \subseteq L_i \\
&= \sum_{i=0}^{2^n} \mathbb{P}\left(f_n^{-1}\left[\left\{\frac{i}{2^n}\right\}\right]\right) \cdot \frac{i}{2^n} + \frac{2}{2^n} && \text{by Proposition 2.6} \\
&= \mathbb{E}[f_n] + \frac{2}{2^n} \leq \mathbb{E}[\mathsf{l_{sup}}] + \varepsilon && \text{by Equation (2.1)}
\end{aligned}
$$

As this holds for all positive $\varepsilon > 0$, it follows that $v(\epsilon) \leq \mathbb{E}[\mathsf{l_{sup}}]$. □

2.4 The proof

This section is devoted to the proof of Theorem 2.3. The first step we take is to define non-decreasing valuations in concurrent games and to link them to non-decreasing valuations in stochastic trees so that we can use the results of the previous Section 2.3.

In concurrent arenas, we consider valuations of the finite sequences of states. Such valuations induce games in normal forms after each finite sequence of states. The notion of being non-decreasing (or non-increasing) can be defined with respect to different conditions. Specifically, a valuation is non-decreasing w.r.t. Player A if, after each finite sequence of states $\rho \in Q^*$, the value w.r.t. Player A of the game in normal form induced by the valuation after ρ is at least $v(\rho)$. We could also define non-decreasing valuation w.r.t. to a Player-B strategy. Symmetrically, we define the notion of non-increasing valuation w.r.t. Player B or a Player-A strategy. Before considering the formal definitions of non-decreasing and non-increasing valuations, let us define below the notion of guard, that is w.r.t. what we consider the values of local interactions.

Definition 2.4 (Guard). *Consider an arbitrary concurrent arena $\mathcal{C}$. We let* $\mathsf{Guard}_{\mathcal{C}} := \mathsf{Guard}^{\mathsf{A}}_{\mathcal{C}} \uplus \mathsf{Guard}^{\mathsf{B}}_{\mathcal{C}}$ *where* $\mathsf{Guard}^{\mathsf{A}}_{\mathcal{C}}$ *is the set of Player-*A *guards with* $\mathsf{Guard}^{\mathsf{A}}_{\mathcal{C}} := \{\mathsf{A}\} \cup \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$ *and* $\mathsf{Guard}^{\mathsf{B}}_{\mathcal{C}}$ *is the set of Player-*B *guards with* $\mathsf{Guard}^{\mathsf{B}}_{\mathcal{C}} := \{\mathsf{B}\} \cup \mathsf{S}^{\mathcal{C}}_{\mathsf{A}}$*. Furthermore, for all* $\mathsf{gd} \in \mathsf{Guard}_{\mathcal{C}}$*, we let* $\mathsf{Opnt}(\mathsf{gd})$ *be the set of strategies to be considered for the opponent with guard* gd*. That is:*

$$\mathsf{Opnt}(\mathsf{gd}) := \begin{cases} \mathsf{S}^{\mathcal{C}}_{\mathsf{B}} & \text{if } \mathsf{gd} = \mathsf{A} \\ \mathsf{S}^{\mathcal{C}}_{\mathsf{A}} & \text{if } \mathsf{gd} = \mathsf{B} \\ \{\mathsf{s}\} & \text{if } \mathsf{gd} = \mathsf{s} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{A}} \cup \mathsf{S}^{\mathcal{C}}_{\mathsf{B}} \end{cases}$$

Finally, for all $\mathsf{gd} \in \mathsf{Guard}_{\mathcal{C}}$ *and* $\pi \in Q^+$*, we let* $\mathsf{gd}(\pi) := \mathsf{gd}$ *if* $\mathsf{gd} \in \{\mathsf{A}, \mathsf{B}\}$ *and* $\mathsf{gd}(\pi) := \mathsf{s}_{\mathsf{A}}(\pi)$ *if* $\mathsf{gd} = \mathsf{s}_{\mathsf{A}} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{A}}$ *and* $\mathsf{gd}(\pi) := \mathsf{s}_{\mathsf{B}}(\pi)$ *if* $\mathsf{gd} = \mathsf{s}_{\mathsf{B}} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$*.*

Remark 2.3. *A quick remark on the terminologies Player-*A *guard and Player-*B *guard. A Player-*A *guard* $\mathsf{gd} \in \mathsf{Guard}_{\mathcal{C}}$ *tells what is the game that Player-*A *is playing. If* $\mathsf{gd} = \mathsf{A}$*, then Player* A *tries to maximize the expected value of the payoff function* f *against all Player-*B *strategies (that is against all strategies in* $\mathsf{Opnt}(\mathsf{A}) = \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$*), whereas if* $\mathsf{gd} = \mathsf{s}_{\mathsf{B}} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$*, Player* A *tries to maximize the expected value of the payoff function against the Player-*B *strategy* s_{B} *(that is against all strategies in* $\mathsf{Opnt}(\mathsf{s}_{\mathsf{B}}) = \{\mathsf{s}_{\mathsf{B}}\}$*). The situation is similar at the local level, with game forms.*

In the following, in Page 84, we will informally motivate why we use guards, but we first need to define important objects.

We define formally the notions of non-decreasing and non-increasing valuations w.r.t. guards.

Definition 2.5 (Non-decreasing valuation in concurrent arenas). *Consider a concurrent arena* $\mathcal{C}$ *and a valuation* $v : Q^+ \to [0, 1]$*. For all* $\rho \in Q^+$*, we denote by* $v^\rho : Q \to [0, 1]$ *the valuation such that, for all* $q \in Q$*,* $v^\rho(q) := v(\rho \cdot q) \in [0, 1]$*.*

Consider a Player-A *guard* $\mathsf{gd} \in \mathsf{Guard}^{\mathsf{A}}_{\mathcal{C}}$*. We say that the valuation* $v : Q^+ \to [0, 1]$ *is* non-decreasing *w.r.t.* gd *if, for all* $\rho \in Q^+$*, we have* $v(\rho) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle][\mathsf{gd}(\rho)]$*.*

Symmetrically, for all Player-B *guards* $\mathsf{gd} \in \mathsf{Guard}^{\mathsf{B}}_{\mathcal{C}}$*, we say that the valuation* $v : Q^+ \to [0, 1]$ *is* non-increasing *w.r.t.* gd *if, for all* $\rho \in Q^+$*, we have* $v(\rho) \geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle][\mathsf{gd}(\rho)]$*.*

We can now define what it means for a strategy to be dominating a valuation w.r.t. a guard.

Definition 2.6 (Dominating a valuation). *Consider a concurrent arena* $\mathcal{C}$ *and a Player-*A *guard* $\mathsf{gd} \in \mathsf{Guard}^{\mathsf{A}}_{\mathcal{C}}$*. A Player-*$\mathsf{A}$ *strategy* s_{A} *dominates a valuation*

$v : Q^+ \to [0,1]$ *w.r.t.* gd *if, for all* $\rho \in Q^+$*, we have:*

$$\begin{aligned} v(\rho) &\leq \mathsf{val}[\langle \mathsf{F}(\pi_{\mathsf{lt}}), v^\rho \rangle](\mathsf{s_A}(\rho)) && \textit{if } \mathsf{gd} = \mathsf{A} \\ v(\rho) &\leq \mathsf{out}[\langle \mathsf{F}(\pi_{\mathsf{lt}}), v^\rho \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) && \textit{if } \mathsf{gd} = \mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C} \end{aligned}$$

In particular, the valuation v *is non-decreasing w.r.t. guard* gd*. This is defined symmetrically for Player* B*.*

It fact, in a concurrent arena, when a Player-A strategy dominates a valuation v, in stochastic trees that can be induced by this strategy, the valuation v is non-decreasing. This is stated in the lemma below.

Lemma 2.10. *Consider an arbitrary concurrent arena* $\mathcal{C}$*. Let* $\mathsf{gd} \in \mathsf{Guard}_\mathcal{C}^\mathsf{A}$ *be a Player-*A *guard and* $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ *be a Player-*A *strategy dominating a valuation* $v : Q^+ \to [0,1]$ *w.r.t.* gd*. For all Player-*B *strategies* $\mathsf{s_B} \in \mathsf{Opnt}(\mathsf{gd})$*, the valuation* v *is non-decreasing in the induced stochastic tree* $\mathcal{T}_\mathcal{C}^{\mathsf{s_A},\mathsf{s_B}}$*.*

Symmetrically, let $\mathsf{gd} \in \mathsf{Guard}_\mathcal{C}^\mathsf{B}$ *be a Player-*B *guard and* $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$ *be a Player-*B *strategy dominating a valuation* $v : Q^+ \to [0,1]$ *w.r.t.* gd*. For all Player-*A *strategies* $\mathsf{s_A} \in \mathsf{Opnt}(\mathsf{gd})$*, the valuation* v *is non-increasing in the induced stochastic tree* $\mathcal{T}_\mathcal{C}^{\mathsf{s_A},\mathsf{s_B}}$*.*

Proof. We prove the result for Player A, the case of Player B being symmetrical. Let us denote $\mathbb{P}_\mathcal{C}^{\mathsf{s_A},\mathsf{s_B}}$ by $\mathbb{P}$. Let $q \in Q$ and $\rho \in Q^*$. Then, we have, by Lemma 1.10 (linearity games in normal form) and by Definition 1.29 of $\mathbb{P}$:

$$\begin{aligned} \sum_{q' \in Q} \mathbb{P}_{q \cdot \rho}(q') \cdot v^q(\rho \cdot q') &= \sum_{q' \in Q} \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), q' \rangle](\mathsf{s_A}(q \cdot \rho), \mathsf{s_B}(q \cdot \rho)) \cdot v^q(\rho \cdot q') \\ &= \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^{q \cdot \rho} \rangle](\mathsf{s_A}(q \cdot \rho), \mathsf{s_B}(q \cdot \rho)) \\ &(\geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^{q \cdot \rho} \rangle](\mathsf{s_A}(q \cdot \rho))) \\ &\geq v^q(\rho) \end{aligned}$$

Recall that q' may be seen as a distribution in $\mathcal{D}(Q)$ that maps q' to 1. Furthermore, the last inequality comes from the fact that the strategy $\mathsf{s_A}$ dominates the valuation v w.r.t. gd (the inequality in parenthesis may be read if $\mathsf{gd} = \mathsf{A}$). Hence v is non-decreasing (recall Definition 2.3) in $\mathcal{T}_\mathcal{C}^{\mathsf{s_A},\mathsf{s_B}}$. □

However, given a Player-A guard $\mathsf{gd} \in \mathsf{Guard}_\mathcal{C}^\mathsf{A}$, for a valuation v non-decreasing w.r.t. gd in an arena, there does not always exist a Player-A strategy dominating v w.r.t. gd. This is due to the fact that the local interactions in that arena may not be maximizable. However, for all positive $\varepsilon > 0$, the valuation v can be modified into an "ε-close" valuation for which there is a Player-A dominating strategy $\mathsf{s_A}$. Furthermore, if the valuation v is (U, m)-uniform for a uniformizing pair (U, m), then the Player-A strategy $\mathsf{s_A}$ can also be chosen (U, m)-uniform. This is formally stated below.

Lemma 2.11 (Proof 2.7.3). *Consider an arbitrary concurrent arena $\mathcal{C}$ supremized w.r.t. Player* A *by a collection* $(S^q_{\mathsf{A}})_{q \in Q}$ *of sets of* GF*-strategies. Consider a Player-*A *guard* gd $\in \mathsf{Guard}^{\mathsf{A}}_{\mathcal{C}}$*. Let* $v : Q^+ \to [0,1]$ *be a valuation non-decreasing w.r.t.* gd *that is* (U, m)*-uniform for a uniformizing pair* (U, m)*. Let* $\varepsilon > 0$ *be a positive real. Then, the valuation* $v_\varepsilon : Q^+ \to [0,1]$ *such that, for all* $\rho \in Q^+$*, we have* $v_\varepsilon(\rho) := \max(v(\rho) - \frac{\varepsilon}{2^{|\rho|-1}}, 0)$ *is such that:*

1. $v - \varepsilon \leq v_\varepsilon \leq v$*;*
2. *it is* (U, m)*-uniform;*
3. *there is a Player-*A (U, m)*-uniform strategy* $\mathsf{s_A}$ *generated by* $(S^{\mathsf{A}}_q)_{q \in Q}$ *dominating w.r.t.* gd *the valuation* v_ε*.*

This is symmetrical for Player B *and a non-increasing valuation.*

2.4.1 . Winning valuations and ε-optimal strategies

There are two main ideas in the proof of Theorem 2.3. The first idea is the following. Consider Proposition 2.9. It states that in a stochastic tree with a non-decreasing valuation, the value of a finite path is less than or equal to the expected value of the (superior) limit of this valuation from that path. In a game $\mathcal{G}$ from a state q_0 and for a value $\alpha \in [0,1]$, consider a valuation $v : Q^+ \to [0,1]$ ensuring the following: 1) it is non-decreasing w.r.t the guard A in the arena $\mathcal{C}$, 2) $v(q_0) = \alpha$ and 3) the superior limit w.r.t. v of all infinite paths, from q_0, is less than or equal to their values w.r.t. the payoff function f. Then, assume that there is a Player-A strategy $\mathsf{s_A}$ dominating this valuation. For all Player-B strategies $\mathsf{s_B} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$, from q_0, in the stochastic tree $\mathcal{T}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}}$ induced by game $\mathcal{G}$ by $\mathsf{s_A}, \mathsf{s_B}$, the expected value of the (superior) limit of v is less than or equal to the expected value of f, by 3). Furthermore, Proposition 2.9 and 1) ensure that the value of q_0 w.r.t. the valuation v — i.e. $v(q_0) = \alpha$, by 2) — is less than or equal to expected value of the (superior) limit of v. Overall, the value of such a Player-A strategy $\mathsf{s_A}$ dominating the valuation v would be at least α (from q_0). We will call such a valuation v a winning valuation (for Player A) w.r.t. (q_0, α). Note that it can be defined symmetrically for Player B. Such winning valuations are defined below in Definition 2.7.

Definition 2.7 (Winning valuations). *Consider a concurrent game* $\mathcal{G} = \langle \mathcal{C}, f \rangle$*, a starting state* q_0 *and a value* $\alpha \in [0,1]$*. Let* gd $\in \mathsf{Guard}^{\mathsf{A}}_{\mathcal{C}}$ *be a Player-*A *guard (resp.* gd $\in \mathsf{Guard}^{\mathsf{B}}_{\mathcal{C}}$ *be a Player-*B *guard). A valuation* $v : Q^+ \to [0,1]$ *is* winning *w.r.t.* (q_0, α) *and* gd *for Player* A *(resp.* B*) if:*

- $v(q_0) = \alpha$ *and* v *is non-decreasing (resp. non-increasing) w.r.t.* gd*;*
- *for all paths* $\rho \in q_0 \cdot Q^\omega$*, we have* $\mathsf{limsup}_i v(\rho_{\leq i}) \leq f_{\mathcal{C}}(\rho)$
(resp. $\mathsf{limsup}_i v(\rho_{\leq i}) \geq f_{\mathcal{C}}(\rho)$*).*

Note that, in the above definition, it is important to use a superior limit for both players instead of a superior limit for Player A and an inferior limit for Player B. We could instead have used an inferior limit for both players, but this would modify the remainder of the chapter. Also, we could have only required that $\alpha \leq v(q_0)$, the interesting properties ensured by winning valuations would still hold. Furthermore, recall that the function $f_{\mathcal{C}} : Q^\omega \to [0,1]$ takes into account the stopping states of the game $\mathcal{G}$, see Definition 1.30.

Lemma 2.12. *Consider a concurrent game $\mathcal{G}$ supremized by a collection $(S^q_{\mathsf{A}})_{q \in Q}$ of* GF*-strategies w.r.t. Player* A*. Consider also a starting state $q_0 \in Q$ and a value $\alpha \in [0,1]$. Let $\mathsf{gd} \in \mathsf{Guard}^{\mathsf{A}}_{\mathcal{C}}$ be a Player-*A *guard. Assume that there is a winning valuation $v : Q^+ \to [0,1]$ w.r.t. (q_0, α) and* gd *for Player* A *that is (U, m)-uniform for a uniformizing pair (U, m). Then, for all $\varepsilon > 0$, Player* A *has a (U, m)-uniform strategy generated by $(S^q_{\mathsf{A}})_{q \in Q}$ whose value against any Player-*B *strategy in $\mathsf{Opnt}(\mathsf{gd})$ is at least $\alpha - \varepsilon$ from q_0.*

This is symmetrical for Player B*.*

Proof. Let $v : Q^+ \to [0,1]$ be such a winning (U, m)-uniform valuation w.r.t. (q_0, α) and gd for Player A. Let $\varepsilon > 0$. Consider the valuation v_ε from Lemma 2.11. It is non-decreasing w.r.t. gd and such that $v - \varepsilon \leq v_\varepsilon \leq v$, hence $\alpha - \varepsilon \leq v(q_0) - \varepsilon \leq v_\varepsilon(q_0)$. Furthermore, it ensures that there is a (U, m)-uniform Player-A strategy s_ε generated by $(S^q_{\mathsf{A}})_{q \in Q}$ dominating it w.r.t. gd. Consider now a Player B strategy $\mathsf{s}_{\mathsf{B}} \in \mathsf{Opnt}(\mathsf{gd})$, the stochastic tree $\mathcal{T}_{\mathcal{C}}^{\mathsf{s}_\varepsilon, \mathsf{s}_{\mathsf{B}}}$ induced by both strategies s_ε and s_{B} and the valuation $v_\varepsilon : Q^+ \to [0,1]$ in that stochastic tree. Since the strategy s_ε dominates v_ε w.r.t. gd, it follows that the valuation v_ε is non-decreasing in the stochastic tree $\mathcal{T}^{\mathcal{C}}_{\mathsf{s}_\varepsilon, \mathsf{s}_{\mathsf{B}}}$ by Lemma 2.10. In particular, it is non-decreasing from q_0. Hence, by Proposition 2.9, we have $v_\varepsilon(q_0) = (v_\varepsilon)^{q_0}(\epsilon) \leq \mathbb{E}^{\mathsf{s}_\varepsilon, \mathsf{s}_{\mathsf{B}}}_{\mathcal{C}, q_0}[\mathsf{limsup}_{(v_\varepsilon)^{q_0}}]$. Since $v_\varepsilon \leq v$, it follows that $\mathsf{limsup}_{(v_\varepsilon)^{q_0}} \leq \mathsf{limsup}_{v^{q_0}}$. Furthermore, by assumption, for all paths $\rho \in q_0 \cdot Q^\omega$, we have $\mathsf{limsup}_v(\rho) \leq f_{\mathcal{C}}(\rho)$. It follows that $\mathbb{E}^{\mathsf{s}_\varepsilon, \mathsf{s}_{\mathsf{B}}}_{\mathcal{C}, q_0}[\mathsf{limsup}_{(v_\varepsilon)^{q_0}}] \leq \mathbb{E}^{\mathsf{s}_\varepsilon, \mathsf{s}_{\mathsf{B}}}_{\mathcal{C}, q_0}[(f_{\mathcal{C}})^{q_0}]$. Overall: $\alpha - \varepsilon \leq \mathbb{E}^{\mathsf{s}_\varepsilon, \mathsf{s}_{\mathsf{B}}}_{\mathcal{C}, q_0}[(f_{\mathcal{C}})^{q_0}]$. As this holds for all Player-B strategies $\mathsf{s}_{\mathsf{B}} \in \mathsf{Opnt}(\mathsf{gd})$, the (U, m)-uniform Player-A strategy s_ε generated by $(S^q_{\mathsf{A}})_{q \in Q}$ has value at least $\alpha - \varepsilon$ from q_0. □

2.4.2 . Existence of winning valuations

High level explanations: with valuable local interactions. The question is now why should there exist such winning valuations. This is where the second idea comes into play. Let us first give the intuition in the case where all local interactions are valuable, where the notion of guards is not used. This is very close to the original idea by Martin. The idea is as follows: we are going to define a standard deterministic win/lose turn-based game $\mathcal{G}_{\mathsf{tb}}$ from the concurrent game $\mathcal{G}$ such that the existence of winning strategies in $\mathcal{G}_{\mathsf{tb}}$ relates to the existence of winning valuation in $\mathcal{G}$. The way this game is played is the following: the game starts at state (q_0, α_0) with $\alpha_0 \in [0,1]$, it is Player

A's turn. She has at her disposal any valuation $v_0 : Q \to [0,1]$ of the states in Q such that the value of the game in normal form $\langle \mathsf{F}(q_0), v_0 \rangle$ is at least α_0. The idea is that she promises that the value of these states is at least the value she gives to them via the valuation v_0. Then, Player B responds by choosing a state q_1. In fact, Player B tries to show that Player A's promise cannot be kept: she tries to reach a state whose value w.r.t. the valuation v_0 is higher than its actual value. We then visit a new Player-A state $(q_0 \cdot q_1, \alpha_1)$ where $\alpha_1 := v_0(q_1)$. The process then repeats itself indefinitely: Player A chooses a valuation of the next states, while Player B picks the next state to visit. This induces an infinite path $((\pi_{\leq i}, \alpha_i) \cdot (\pi_{\leq i}, v_i))_{i \in \mathbb{N}}$. Player A wins this win/lose game if the superior limit of the values $(\alpha_i)_{i \in \mathbb{N}}$ is at most the value of the infinite path $\pi \in q_0 \cdot Q^\omega$ w.r.t. the payoff function $f_{\mathcal{C}}$. In other words, Player B loses if she is not able to show that Player A broke her promise, i.e. Player A has not over-estimated the values of the states visited.

The turn-based game $\mathcal{G}_{\mathsf{tb}}$ we have described is deterministic (and the objective is Borel), hence Borel determinacy (i.e. Theorem 2.1) ensures that, for all $u \in [0,1]$, either of the players have a winning strategy in the game $\mathcal{G}_{\mathsf{tb}}$ from the state (q_0, u). Furthermore, the higher u is, the more difficult it is for Player A to win from the state (q_0, u). This means that there is a threshold $\alpha(q_0) \in [0,1]$ such that, for all $u < \alpha(q_0)$, Player A has a winning strategy from the state (q_0, u) whereas, for all $u > \alpha(q_0)$, Player B has a winning strategy from the state (q_0, u). Then, for all $\varepsilon > 0$, from a Player-A winning strategy in the game $\mathcal{G}_{\mathsf{tb}}$ from the state $(q_0, \alpha(q_0) - \varepsilon)$ we are able to build a Player-A winning valuation w.r.t. $(q_0, \alpha(q_0) - \varepsilon)$. This is in fact rather straightforward from the definition of the game $\mathcal{G}_{\mathsf{tb}}$, and also because a Player-A strategy in the game $\mathcal{G}_{\mathsf{tb}}$ chooses values for the states. Almost symmetrically, from a Player-B winning strategy in the game $\mathcal{G}_{\mathsf{tb}}$ from the state $(q_0, \alpha(q_0) + \varepsilon)$ we are able to build a Player-B winning valuation w.r.t. $(q_0, \alpha(q_0) + 2 \cdot \varepsilon)$. This, however is less direct. This is due to the fact that, contrary to Player-A strategies, Player-B strategies in the game $\mathcal{G}_{\mathsf{tb}}$ do not choose values for the states. This shows that the value of the state q_0 (in the game $\mathcal{G}$) is at least and at most $\alpha(q_0)$, it is therefore equal to the threshold $\alpha(q_0)$.

When local interactions are not valuable: the need for guards. Let us now consider the case where the local interactions are not necessarily valuable. We will now use guards, in particular when considering winning valuations. We are trying to show result $(1.a)$. In that case, the above-described turn-based game $\mathcal{G}_{\mathsf{tb}}$ is not well-defined anymore. Indeed, when considering the valuations $v_0 : Q \to [0,1]$ of the states allowed to Player A at a state (q_0, α_0), we can no longer talk about the value of the game in normal form $\langle \mathsf{F}(q_0), v_0 \rangle$ since it does not exist, a priori. A possible way to fix this is to consider a new turn-based game $\mathcal{G}^{\mathsf{A}}_{\mathsf{tb}}$ in which we consider the Player-A value of the games in normal form. In that new turn-based game, that mimics the game

$\mathcal{G}_{\mathsf{tb}}$, the valuations $v_0 : Q \to [0,1]$ of the states allowed to Player A at a state (q_0, α_0) are the ones such that α_0 is at least the Player-A value of the game in normal form $\langle \mathsf{F}(q_0), v_0 \rangle$. Otherwise, the game $\mathcal{G}_{\mathsf{tb}}$ is left unchanged. Then, as before, from all states q_0, there is a threshold $\alpha^{\mathsf{A}}(q_0) \in [0,1]$ such that for all $u < \alpha(q_0, \mathsf{A})$, Player A has a winning strategy from the state (q_0, u) in $\mathcal{G}^{\mathsf{A}}_{\mathsf{tb}}$ whereas, for all $u > \alpha(q_0, \mathsf{A})$, Player B has a winning strategy from the state (q_0, u) in $\mathcal{G}^{\mathsf{A}}_{\mathsf{tb}}$. Furthermore, as before, for all $\varepsilon > 0$, from a Player-A winning strategy in the game $\mathcal{G}^{\mathsf{A}}_{\mathsf{tb}}$ from the state $(q_0, \alpha(q_0, \mathsf{A}) - \varepsilon)$ we are able to build a Player-A winning valuation v_ε w.r.t. $(q_0, \alpha(q_0, \mathsf{A}) - \varepsilon)$ and A. Furthermore, we will then be able to build a Player-A strategy $\mathsf{s}^{2\varepsilon}_{\mathsf{A}}$ of value at least $\alpha(q_0, \mathsf{A}) - 2 \cdot \varepsilon$ from q_0 that ensures the properties stated in result $(1.a)$ by using Lemma 2.11. However, we are no longer able to do the (almost) symmetrical thing for Player B. This is in fact unsurprising since there is no reason for $\alpha(q_0, \mathsf{A})$ to be equal to the value of the state q_0[4]. However, we can hope (and it is in fact the case) that $\alpha(q_0, \mathsf{A})$ is equal to the Player-A value of the state q_0 in the game $\mathcal{G}$. This would show that the Player-A strategy $\mathsf{s}^{2\varepsilon}_{\mathsf{A}}$ mentioned above is indeed 2ε-optimal from the state q_0. We have already shown that the Player-A value of the state q_0 is at least $\alpha(q_0, \mathsf{A})$. However, the question is now: how do we prove that it is at most $\alpha(q_0, \mathsf{A})$?

We cannot exhibit Player-B strategies of value arbitrarily close to $\alpha(q_0, \mathsf{A})$, as we did with valuable local interactions, since $\alpha(q_0, \mathsf{A})$ is a priori not the value of the state q_0. The only way (that we can think of) is to show that for any Player-A strategy $\mathsf{s_A}$, the value of the strategy $\mathsf{s_A}$ from q_0 is at most $\alpha(q_0, \mathsf{A})$. However, we need to link the value of the strategy $\mathsf{s_A}$ from q_0 with $\alpha(q_0, \mathsf{A})$. Our idea is then to define yet another turn-based game $\mathcal{G}^{\mathsf{s_A}}_{\mathsf{tb}}$ that mimics the game $\mathcal{G}^{\mathsf{A}}_{\mathsf{tb}}$ but changes the local condition: in that game, the valuations $v_0 : Q \to [0,1]$ of the states allowed to Player A at a state (q_0, α_0) are the ones such that α_0 is at most the value of the GF-strategy $\mathsf{s_A}(q_0)$ in the game in normal form $\langle \mathsf{F}(q_0), v_0 \rangle$. Then, as before, from all states q_0, there is a threshold $\alpha^{\mathsf{s_A}}(q_0) \in [0,1]$ such that for all $u < \alpha(q_0, \mathsf{s_A})$, Player A has a winning strategy from the state (q_0, u) in $\mathcal{G}^{\mathsf{s_A}}_{\mathsf{tb}}$ whereas, for all $u > \alpha(q_0, \mathsf{s_A})$, Player B has a winning strategy from the state (q_0, u) in $\mathcal{G}^{\mathsf{s_A}}_{\mathsf{tb}}$. Furthermore, it now holds that, for all $\varepsilon > 0$, from a Player-B winning strategy in the game $\mathcal{G}^{\mathsf{s_A}}_{\mathsf{tb}}$ from the state $(q_0, \alpha(q_0, \mathsf{s_A}) + \varepsilon)$ we are able to build a Player-B winning valuation w.r.t. $(q_0, \alpha(q_0, \mathsf{s_A}) + 2 \cdot \varepsilon)$ and $\mathsf{s_A}$. This shows that the $\mathsf{s_A}$-value of the state q_0 is at most $\alpha(q_0, \mathsf{s_A})$. In addition, from any state (q, u), if Player A wins in the game $\mathcal{G}^{\mathsf{s_A}}_{\mathsf{tb}}$, then she also wins in the game $\mathcal{G}^{\mathsf{A}}_{\mathsf{tb}}$. This is due to the fact that the local condition in $\mathcal{G}^{\mathsf{s_A}}_{\mathsf{tb}}$ allows less choices for Player A than the local condition in $\mathcal{G}^{\mathsf{A}}_{\mathsf{tb}}$. Hence, we have $\alpha(q_0, \mathsf{s_A}) \leq \alpha(q_0, \mathsf{A})$. This shows what we wanted: all Player-A strategies have value at most $(q_0, \alpha(q_0, \mathsf{A}))$ from q_0.

[4]Otherwise, it would imply that any concurrent game has a value, regardless of the local interactions involved, which we know is not true

Overall, in order to show result (1.a) for Player A, we define several turn-based games: first, the game $\mathcal{G}_{\mathsf{tb}}^{\mathsf{A}}$, similar to the game $\mathcal{G}_{\mathsf{tb}}$, except that the valuations allowed to Player A are determined by the Player-A value of the game in normal forms. By exhibiting Player-A winning valuations w.r.t. A, we can then show that the Player-A value of any state q_0 is at least $\alpha(q_0, \mathsf{A})$. To show that $\alpha(q_0, \mathsf{A})$ is actually equal to the Player-A value of the state q_0, we define, for all Player-A strategies $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$, the turn-based game $\mathcal{G}_{\mathsf{tb}}^{\mathsf{s_A}}$. It is also similar to the game $\mathcal{G}_{\mathsf{tb}}$, except that the valuations allowed to Player A are determined by the value, in games in normal form, of the Player-A GF-strategies obtained from the strategy $\mathsf{s_A}$. By exhibiting Player-B winning valuations w.r.t. $\mathsf{s_A}$, we can then show that the $\mathsf{s_A}$-value of any state q_0 is at most $\alpha(q_0, \mathsf{s_A})$ which is less than or equal to $\alpha(q_0, \mathsf{A})$. In terms of guards, we have that for the Player-A guard $\mathsf{gd} = \mathsf{A}$, we exhibit Player-A strategies, whereas for the Player-B guard $\mathsf{gd} = \mathsf{s_A}$, we exhibit Player-$\mathsf{B}$ strategies.

In the following, we also define turn-based games for the guards $\mathsf{gd} = \mathsf{B}$ and $\mathsf{gd} = \mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$ to do the same for Player B. Furthermore, we will handle at the same time all Player-A guards, by exhibiting winning valuations for Player A and all Player-B guards, by exhibiting winning valuations for Player B.

Formal definitions and proofs. For all guards $\mathsf{gd} \in \mathsf{Guard}_{\mathcal{C}}$, we formally define the game $\mathcal{G}_{\mathsf{tb}}^{\mathsf{gd}}$.

Definition 2.8. *Consider an arbitrary concurrent game $\mathcal{G}$ and a guard $\mathsf{gd} \in \mathsf{Guard}_{\mathcal{C}}$. We build the following deterministic turn-based win/lose game $\mathcal{G}_{\mathsf{tb}}^{\mathsf{gd}} := \langle \mathcal{C}_{\mathsf{tb}}^{\mathsf{gd}}, W_{\mathsf{tb}}(f) \rangle$. Note that this arena $\mathcal{C}_{\mathsf{tb}}^{\mathsf{gd}}$ need not be colored*[5]*, and the winning objective $W_{\mathsf{tb}}(f) \subseteq (Q_{\mathsf{tb}})^\omega$ is directly given as a Borel subset of infinite paths. Recalling Definition 1.11 for the definition of turn-based game forms, we let $\mathcal{C}_{\mathsf{tb}}^{\mathsf{gd}} = \langle Q_{\mathsf{tb}}, \mathsf{F}^{\mathsf{gd}} \rangle$ be such that:*

- $Q_{\mathsf{tb}} := Q_{\mathsf{A}} \uplus Q_{\mathsf{B}}$;
- $Q_{\mathsf{A}} := \{(\pi, \alpha) \mid \pi \in Q^+,\ \alpha \in [0,1]\}$ *is the set of Player-A states;*
- $Q_{\mathsf{B}} := \{(\pi, h) \mid \pi \in Q^+,\ h : Q \to [0,1]\}$ *is the set of Player-B states;*
- *For all Player-A states $(\pi, \alpha) \in Q_{\mathsf{A}}$, $\mathsf{F}^{\mathsf{gd}}((\pi, \alpha)) := \langle \mathsf{Move}_{\mathsf{A}}^{\mathsf{gd}}(\pi, \alpha), \{*\}, Q_{\mathsf{B}}, \varrho_{\mathsf{A}}^{\pi} \rangle$ with*

  $$\mathsf{Move}_{\mathsf{A}}^{\mathsf{gd}}(\pi, \alpha) := \{h : Q \to [0,1] \mid \mathsf{val}[\langle \mathsf{F}(\pi_{\mathsf{lt}}), h \rangle][\mathsf{gd}(\pi)] \geq \alpha\}$$

  and $\varrho_{\mathsf{A}}^{\pi} : \mathsf{Move}_{\mathsf{A}}^{\mathsf{gd}}(\pi, \alpha) \to Q_{\mathsf{B}}$ such that, for all $h \in \mathsf{Move}_{\mathsf{A}}^{\mathsf{gd}}(\pi, \alpha)$, we have $\varrho_{\mathsf{A}}^{\pi}(h) := (\pi, h) \in Q_{\mathsf{B}}$.

[5] If we want the definition to exactly fit in the formalism of Definition 1.23, we could consider Q_{tb} itself as set of colors and the identity function as coloring function.

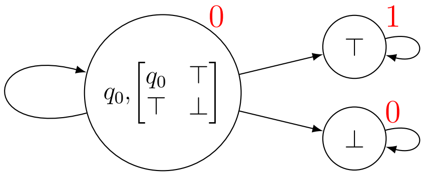

Figure 2.2: A deterministic concurrent reachability game $\mathcal{G} = \langle \mathcal{C}, W \rangle$ where Player A wants to reach the target $\{\top\}$.

$$\langle \mathsf{F}(q_0), h \rangle = \begin{bmatrix} x & 1 \\ 1 & 0 \end{bmatrix}$$

Figure 2.3: The local interaction at state q_0 in the game of Figure 2.2 valued with the valuation $h : \{q_0 \mapsto x, \top \mapsto 1, \bot \mapsto 0\}$.

- *For all Player-*B *states* $(\pi, h) \in Q_{\mathsf{B}}$, $\mathsf{F}^{\mathsf{gd}}((\pi, h)) := \langle \{*\}, \mathsf{Move}_{\mathsf{B}}, Q_{\mathsf{A}}, \varrho_{\mathsf{B}}^{\pi,h} \rangle$ *with*
$$\mathsf{Move}_{\mathsf{B}} := Q$$
and $\varrho_{\mathsf{B}}^{h} : \mathsf{Move}_{\mathsf{B}} \to Q_{\mathsf{A}}$ *such that, for all* $q \in \mathsf{Move}_{\mathsf{B}}$*, we have* $\varrho_{\mathsf{B}}^{\pi,h}(q) := (\pi \cdot q, h(q)) \in Q_{\mathsf{A}}$.

- $W_{\mathsf{tb}}(f) := \{(\pi_0, \alpha_0) \cdot (\pi_0, h_0) \cdot (\pi_{\leq 1}, \alpha_1) \cdots \in (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^\omega \mid \mathsf{limsup}(\alpha_i)_{i \in \mathbb{N}} \leq f_{\mathcal{C}}(\pi)\} \subseteq (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^\omega$.

Note that the winning set $W_{\mathsf{tb}}(f)$ *is Borel since the functions* $f_{\mathcal{C}}$ *and the superior limit* limsup *are measurable.*

Note that the different games $\mathcal{G}_{\mathsf{tb}}^{\mathsf{gd}}$ indexed by gd only differ by their sets of moves available to Player A at her states.

Example 2.1. *Consider the standard concurrent game* $\mathcal{G} = \langle \mathcal{C}, W \rangle$ *from Figure 2.2. Note that all local interactions are finite and therefore valuable. Hence,* $\mathcal{G}_{\mathsf{tb}}^{\mathsf{A}} = \mathcal{G}_{\mathsf{tb}}^{\mathsf{B}}$ *(we will discuss it again below in Observation 2.1). The game* $\mathcal{G}$ *is win/lose and Player* A *wins if and only if the state* $\top$ *is reached. (Recall the reachability objective from Definition 1.25.) Hence, Player* B *wants to either loop indefinitely on* q_0 *or reach the state* $\bot$*. Let us exemplify Definition 2.8 on this game. Part of the arena* $\mathcal{C}_{\mathsf{tb}}^{\mathsf{A}} = \mathcal{C}_{\mathsf{tb}}^{\mathsf{B}}$ *is represented in Figure 2.4. Player-*A *states are rectangle-shaped, whereas Player-*B *states are hexagon-shaped*[6]*.*

From the state (q_0, α)*, Player* A *chooses a valuation* h *of successor states. In what is depicted in Figure 2.4, we only drew valuations* h *such that* $h(\top) = 1$ *and* $h(\bot) = 0$*. The reason for that is the following: both* $\bot$ *and* $\top$ *are self-looping states. Furthermore, reaching and looping on* $\top$ *is winning for Player* A *while reaching and looping on* q_0 *is loosing for Player* A *. Hence, Player* A *can safely value* $\top$ *with 1 since that would not lead to her overestimating its value. However, she cannot*[7] *value* $\bot$ *with a positive value* $x > 0$*. Indeed,*

[6]Note that in Definition 2.8, Player-B states are pairs of state and valuation of successor states. We did not indicate the state (which is q_0) to simplify our drawing.

[7]More precisely, she can but she should not if she wants to win.

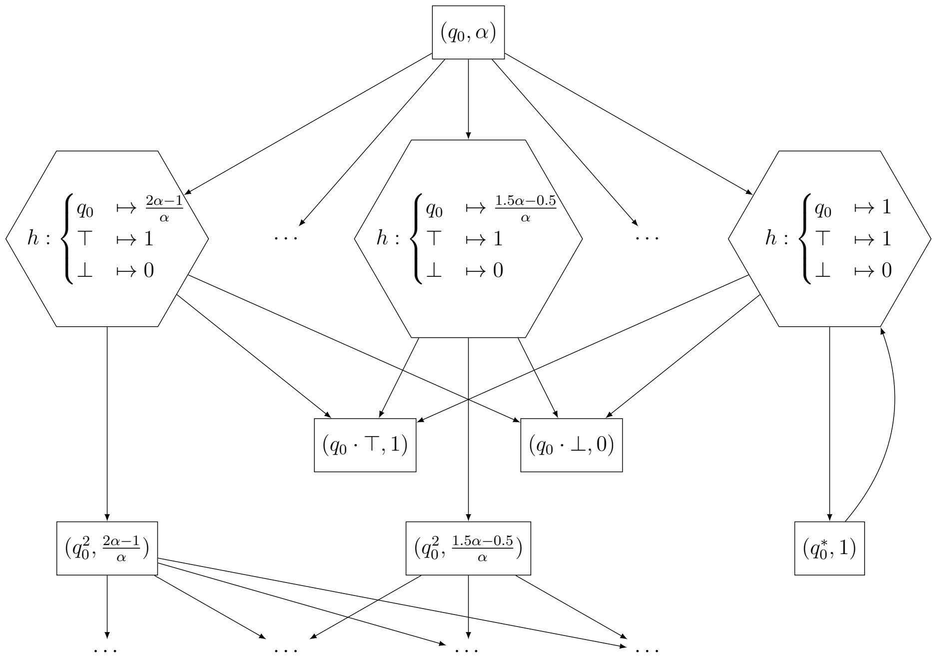

Figure 2.4: A part of the turn-based arena $\mathcal{C}_{\mathsf{tb}}$ from Definition 2.8 built from the concurrent reachability game of Figure 2.2.

in that case by definition of $\mathsf{Move}_{\mathsf{A}}^{\mathsf{gd}}(\bot, \cdot)$*, it would imply that the superior limit of the values seen is at least* x *whereas the target is not reached (i.e. the corresponding infinite path is losing, in terms of payoff functions, it has value 0). In other words the value of* $\bot$ *is 0, hence giving it a positive value would be overestimating this value. Alternatively,* $\top$ *could be stopping state of value 1 and* $\bot$ *could be a stopping state of value 0. Furthermore, note that if the game is at a state* $(q, \cdot) \in Q_{\mathsf{A}}$*, for any state* $q' \in Q$ *such that* q' *does not appear in* $\mathsf{F}(q)$*, then Player* A *should always choose a valuation* h *such that* $h(q') := 0$ *since it is always easier to win for her if the value of the state is smaller.*

The only relevant choice for Player A *remains in how to value* q_0*. We have depicted several possibilities (only for* $\alpha \geq \frac{1}{2}$*, otherwise* $\frac{2\alpha-1}{\alpha} < 0$*). In that case, the minimum that Player* A *can value* q_0 *is equal to* $\frac{2\alpha-1}{\alpha}$*. Indeed, one can check that in the game in normal form* $\langle \mathsf{F}(q_0), h \rangle$ *with* $h : \{q_0 \mapsto x, \top \mapsto 1, \bot \mapsto 0\}$ *— that is depicted in Figure 2.3 — we have* $\mathsf{val}[\langle \mathsf{F}(q_0), h \rangle] = \frac{1}{2-x}$*. Hence,* $\mathsf{val}[\langle \mathsf{F}(q_0), h \rangle] \geq \alpha$ *if and only if* $x \geq \frac{2\alpha-1}{\alpha}$*. In particular, if* $\alpha \leq \frac{1}{2}$*, then* $x = 0$ *works. Furthermore, the greater* α *is, the greater* x *must be. This can be informally seen at the bottom of Figure 2.4 between states* $(q_0, \frac{2\alpha-1}{\alpha})$ *and* $(q_0, \frac{1.5\alpha-0.5}{\alpha})$*. Since* $\frac{1.5\alpha-0.5}{\alpha} \geq \frac{2\alpha-1}{\alpha}$*, Player* A *has more possibilities from* $(q_0, \frac{2\alpha-1}{\alpha})$ *than from* $(q_0, \frac{1.5\alpha-0.5}{\alpha})$*, as represented in the arrows exiting from these states.*

We make a central observation below. Note that it is only in the proof of this observation that we use Theorem 2.1 (Borel determinacy: the existence of winning strategies in deterministic standard turn-based win/lose games).

Observation 2.1. *Consider the arbitrary concurrent game $\mathcal{G}$, a guard* $\mathsf{gd} \in \mathsf{Guard}_{\mathcal{C}}$ *and a starting state q_0. For all Borel winning conditions $W \subseteq (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^\omega$, there is a value $\alpha^{\mathsf{gd}}(q_0, W) \in [0,1]$, such that, for all $\alpha_{\mathsf{A}} < \alpha^{\mathsf{gd}}(q_0, W) < \alpha_{\mathsf{B}}$, Player* A *has a winning strategy from $(q_0, \alpha_{\mathsf{A}})$ and Player* B *has a winning strategy from $(q_0, \alpha_{\mathsf{B}})$ in the game $\langle \mathcal{C}^{\mathsf{gd}}_{\mathsf{tb}}, W \rangle$. This value $\alpha^{\mathsf{gd}}(q_0, W)$ is called the* threshold *of the parameterized game $\langle \mathcal{C}^{\mathsf{gd}}_{\mathsf{tb}}, W \rangle$ from q_0. Furthermore, for all $\mathsf{s}_{\mathsf{A}} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{A}}$ and $\mathsf{s}_{\mathsf{B}} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$:*

$$\alpha^{\mathsf{s}_{\mathsf{A}}}(q_0, W) \leq \alpha^{\mathsf{A}}(q_0, W) \leq \alpha^{\mathsf{B}}(q_0, W) \leq \alpha^{\mathsf{s}_{\mathsf{B}}}(q_0, W)$$

with $\alpha^{\mathsf{A}}(q_0, W) = \alpha^{\mathsf{B}}(q_0, W)$ as soon as all local interactions are valuable.

Finally, in the game $\mathcal{G}^{\mathsf{gd}}_{\mathsf{tb}} = \langle \mathcal{C}^{\mathsf{gd}}_{\mathsf{tb}}, W_{\mathsf{tb}}(f) \rangle$, Player A *has a winning strategy from $(q_0, 0)$.*

Proof. Consider two thresholds $\alpha < \alpha'$. Then, in the arena $\mathcal{C}^{\mathsf{gd}}_{\mathsf{tb}}$, from (q_0, α) Player A has less strategies available than from (q_0, α') while the strategies available to Player B are the same. Since from every state, either of the players has winning strategy — by Theorem 2.1, as we consider a deterministic win/lose turn-based game with a Borel objective — the first result follows. Furthermore, for all $\pi \in Q^+$ and $h : Q \to [0,1]$, letting $\mathcal{F} := \mathsf{F}(\pi_{\mathsf{lt}})$, we have:

$$\mathsf{val}[\langle \mathcal{F}, h \rangle][\mathsf{s}_{\mathsf{A}}(\pi)] \leq \mathsf{val}[\langle \mathcal{F}, h \rangle][\mathsf{A}] \leq \mathsf{val}[\langle \mathcal{F}, h \rangle][\mathsf{B}] \leq \mathsf{val}[\langle \mathcal{F}, h \rangle][\mathsf{s}_{\mathsf{B}}(\pi)]$$

In addition, if $\mathcal{F}$ is valuable, we have $\mathsf{val}[\langle \mathcal{F}, h \rangle][\mathsf{A}] = \mathsf{val}[\langle \mathcal{F}, h \rangle][\mathsf{B}]$ and $\mathsf{Move}^{\mathsf{A}}_{\mathsf{A}}(\pi, \alpha) = \mathsf{Move}^{\mathsf{B}}_{\mathsf{A}}(\pi, \alpha)$. The second result follows. Finally, in the game $\mathcal{G}^{\mathsf{gd}}_{\mathsf{tb}}$, from the state $(q_0, 0)$, Player A has the winning strategy consisting in always playing the valuation 0 mapping every state to 0, ensuring that the superior limit is less than or equal to f. □

Example 2.2. *Let us compute the threshold $\alpha^{\mathsf{A}}(q_0, W_{\mathsf{tb}}(W)) = \alpha^{\mathsf{B}}(q_0, W_{\mathsf{tb}}(W))$ in the game of Figure 2.4. Note that since W — the objective in the original concurrent game $\mathcal{G}$ — is win/lose, the objective $W_{\mathsf{tb}}(W)$ in the turn-based game $\mathcal{G}_{\mathsf{tb}}$ can be reformulated as follows: an infinite path $\rho \in (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^\omega$ is in $W_{\mathsf{tb}}(W)$ if and only if either a state $(\top, \cdot) \in Q_{\mathsf{A}}$ is seen or the limit of the values in states of the shape $(q_0^n, \alpha) \in Q_{\mathsf{A}}$ is 0. A similar reformulation will be used in Subsection 2.4.3 to show result $(1.b)$ of Theorem 2.3.*

As mentioned in Example 2.1, if $\alpha \leq \frac{1}{2}$, Player A *can value q_0 with 0 and ensure winning. Indeed, Player* B *may go to $(q_0, 0) \in Q_{\mathsf{A}}$ or $(\bot, 0) \in Q_{\mathsf{A}}$ and in both cases Player* A *will win as mentioned in Observation 2.1. Player* B *may also go to $(\top, 1)$ and in that case Player* A *will also win. Now, consider some $\frac{1}{2} < \alpha < 1$. Let $g : [\frac{1}{2}, 1) \to [0,1]$ be such that, for all $x \in [\frac{1}{2}, 1)$, we have $g(x) := \frac{2x-1}{x}$. From (q_0, α), Player* A *first chooses a valuation h_0 such* *that $h_0(q_0) := g(\alpha)$. Then if Player* B *has chosen to come back to q_0, Player* A *chooses a valuation h_1 such that $h_1(q_0) := g \circ g(\alpha)$. This is repeated until a state $(q_0^n, x) \in Q_\mathsf{A}$ is reached with $x \leq \frac{1}{2}$, after which Player* A *values q_0 with 0. (Note that for all $x \in [\frac{1}{2}, 1)$, there is some $n \in \mathbb{N}$ such that $g^{(n)}(x) \leq \frac{1}{2}$, where $g^{(n)}$ refers to the composition n times of g.) Hence, $\alpha^\mathsf{A}(q_0, W_\mathsf{tb}(W)) = \alpha^\mathsf{B}(q_0, W_\mathsf{tb}(W)) = 1$. However, note that when $\alpha = 1$, it is Player* B *who has a winning strategy from (q_0, α). Player* A *has to value q_0 with 1 indefinitely, as depicted in Figure 2.4. Hence, Player* B *can loop indefinitely on q_0 (the state $(q_0^*, 1)$ refers to the set of states $\{(q_0^n, 1) \mid n \geq 2\}$) while ensuring that the superior limit is positive (in fact, it is equal to 1).*

Crucially, the existence of winning strategies in $\mathcal{G}_\mathsf{tb}^\mathsf{gd}$ for a starting state (q_0, α) from $q_0 \in Q$ and $\alpha \in [0, 1]$ implies the existence of winning valuations w.r.t. (q_0, α) and gd in the game $\mathcal{G}$. Note that although this holds as is for Player A, we prove a slightly weaker statement for Player B. Furthermore, the proof for Player A is quite straightforward, while it is harder for Player B. In addition, with these lemmas, we are able to show results (1.a) and (2) from Theorem 2.3. We will use an additional lemma — that we discuss in the next subsection — to obtain result (1.b).

Lemma 2.13. *Consider an arbitrary concurrent game $\mathcal{G}$, a starting state q_0 and a value $\alpha \in [0, 1]$. Let $\mathsf{gd} \in \mathsf{Guard}_\mathcal{C}^\mathsf{A}$ be a Player-*A *guard. Assume that Player* A *has a winning (deterministic) strategy in the game $\mathcal{G}_\mathsf{tb}^\mathsf{gd}$ from the state (q_0, α). Then, there is a valuation $v : Q^+ \to [0, 1]$ that is winning w.r.t. (q_0, α) and* gd *for Player* A.

Before proving this lemma, we first introduce below in Definition 2.9 a useful function mapping finite paths $\mathcal{C}$ to finite paths in $\mathcal{C}_\mathsf{tb}$ that are compatible with a Player-A strategy in $\mathcal{C}_\mathsf{tb}$ from a given starting state.

Definition 2.9 (Map to finite paths compatible with a strategy)**.** *Consider an arbitrary concurrent arena $\mathcal{C}$, a Player-*A *guard $\mathsf{gd} \in \mathsf{Guard}_\mathcal{C}^\mathsf{A}$ and the turn-based arena $\mathcal{C}_\mathsf{tb}^\mathsf{gd}$ from Definition 2.8. For all Player-*A *strategies $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}_\mathsf{tb}^\mathsf{gd}}$ in the turn-based arena $\mathcal{C}_\mathsf{tb}^\mathsf{gd}$ and a starting state $(q_0, \alpha) \in Q_\mathsf{A}$. We let $p_{\mathsf{s_A}}^{(q_0,\alpha)} : q_0 \cdot Q^* \to (Q_\mathsf{A} \cdot Q_\mathsf{B})^* \cdot Q_\mathsf{A}$ be defined inductively by, for all $\pi \in q_0 \cdot Q^*$:*

$$p_{\mathsf{s_A}}^{(q_0,\alpha)}(\pi) := \begin{cases} (q_0, \alpha) & \text{if } \pi = q_0 \\ p_{\mathsf{s_A}}^{(q_0,\alpha)}(\rho) \cdot (\rho, \mathsf{s_A} \circ p_{\mathsf{s_A}}^{(q_0,\alpha)}(\rho)) \cdot (\pi, \mathsf{s_A} \circ p_{\mathsf{s_A}}^{(q_0,\alpha)}(\rho)(q)) & \text{if } \pi = \rho \cdot q \end{cases}$$

This function is then extended to infinite paths $(p_{\mathsf{s_A}}^{q_0,\alpha})^\omega : q_0 \cdot Q^\omega \to (Q_\mathsf{A} \cdot Q_\mathsf{B})^\omega$.

Observation 2.2. *For all concurrent arenas $\mathcal{C}$, Player-*A *guards $\mathsf{gd} \in \mathsf{Guard}_\mathcal{C}^\mathsf{A}$ and Player-*A *strategies $\mathsf{s_A}$ in the turn-based arena $\mathcal{C}_\mathsf{tb}^\mathsf{gd}$, for all infinite paths $\rho \in q_0 \cdot Q^\omega$, the infinite path $(p_{\mathsf{s_A}}^{q_0,\alpha})^\omega(\rho) \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^\omega$ is compatible with the strategy $\mathsf{s_A}$ from the state (q_0, α): $(p_{\mathsf{s_A}}^{q_0,\alpha})^\omega(\rho) \in \mathsf{CP}_{\mathcal{C}_\mathsf{tb},(q_0,\alpha)}(\mathsf{s_A})$. (Recall Definition 1.33.)*

Let us proceed to the proof of Lemma 2.13.

Proof. Consider a Player-A winning strategy $\mathsf{s_A}$ in the game $\mathcal{G}_{\mathsf{tb}}^{\mathsf{gd}}$ from the state (q_0, α). We let $p := p_{\mathsf{s_A}}^{(q_0,\alpha)}$. We define inductively a valuation $v : q_0 \cdot Q^* \to [0,1]$ in the following way:

- $v(q_0) := \alpha$;
- for all $\rho \in q_0 \cdot Q^*$, we let $v^\rho := \mathsf{s_A} \circ p(\rho) : Q \to [0,1]$.

Furthermore, for all $\rho \in Q^+ \setminus q_0 \cdot Q^*$, we set $v(\rho) := 0$. With this definition, this valuation v ensures the property:

$$\forall \rho \in q_0 \cdot Q^*,\ p(\rho)_{\mathsf{lt}} = (\rho, v(\rho))$$

Indeed, $p(q_0) = (q_0, \alpha) = (q_0, v(q_0))$. Furthermore, for all $\rho \cdot q \in q_0 \cdot Q^+$, we have $p(\rho \cdot q)_{\mathsf{lt}} = (\rho \cdot q, \mathsf{s_A} \circ p(\rho)(q)) = (\rho \cdot q, v(\rho \cdot q))$.

Let us show that this valuation is winning for Player A w.r.t. (q_0, α) and gd in the game $\mathcal{G}$. First, it is non-decreasing w.r.t. gd. Indeed, this holds for finite paths in $Q^+ \setminus q_0 \cdot Q^*$ and for all finite paths $\rho \in q_0 \cdot Q^*$, we have $\mathsf{s_A} \circ p(\rho) \in \mathsf{Move}_{\mathsf{A}}^{\mathsf{gd}}(p(\rho)_{\mathsf{lt}}) = \mathsf{Move}_{\mathsf{A}}^{\mathsf{gd}}((\rho, v(\rho)))$. That is, $v(\rho) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \mathsf{s_A} \circ p(\rho) \rangle][\mathsf{gd}(\rho)] = \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle][\mathsf{gd}(\rho)]$. In addition, the valuation v ensures that $v(q_0) = \alpha$. Furthermore, consider a path $\rho \in q_0 \cdot Q^\omega$. The infinite path $p(\rho)$ is compatible with the strategy $\mathsf{s_A}$ from (q_0, α) in $\mathcal{G}_{\mathsf{tb}}^{\mathsf{gd}}$ as mentioned in Observation 2.2. Since this strategy is winning, it follows that $p(\rho) \in W$. Since for all $i \in \mathbb{N}$, we have $p(\rho^{\leq i})_{\mathsf{lt}} = (\rho_{\leq i}, v(\rho^{\leq i}))$, it follows that $\mathsf{limsup}\ v(\rho^{\leq i}) \leq f_{\mathcal{C}}(\rho)$. As this holds for all paths $\rho \in q_0 \cdot Q^\omega$, it follows that the valuation v is winning w.r.t. (q_0, α) and gd for Player A. □

Lemma 2.14 below is the analogue of Lemma 2.13 for Player B.

Lemma 2.14. *Consider an arbitrary concurrent game $\mathcal{G}$, a starting state q_0 and a value $\alpha \in [0,1]$. Let $\mathsf{gd} \in \mathsf{Guard}_{\mathcal{C}}^{\mathsf{B}}$ be a Player-B guard. Assume that Player B has a winning (deterministic) strategy in the game $\mathcal{G}_{\mathsf{tb}}^{\mathsf{gd}}$ from the state (q_0, α). Then, for all $0 < \varepsilon < 1 - \alpha$, there is a valuation $v : Q^+ \to [0,1]$ that is winning for Player B w.r.t. $(q_0, \alpha + \varepsilon)$ and gd.*

The proof is not symmetric compared to the proof of Lemma 2.13. Here, it is harder to come up with the appropriate valuation since a Player-B winning strategy $\mathsf{s_B}$ in the game $\mathcal{G}_{\mathsf{tb}}^{\mathsf{gd}}$ does not choose values for the states but pick states once Player A has chosen a value for them. The idea to define a value for a state q is to consider the infimum over the values that Player A can choose for q that makes the Player-B winning strategy $\mathsf{s_B}$ go to q.

Proof. Consider a Player-B winning strategy $\mathsf{s_B}$ from the state (q_0, α) in the game $\mathcal{G}_{\mathsf{tb}}^{\mathsf{gd}}$. Let $0 < \varepsilon < 1 - \alpha$. We want to define a valuation $v : Q^+ \to [0,1]$. First, for all $\rho \in Q^+ \setminus q_0 \cdot Q^*$, we let $v(\rho) := 1$. Then, we define inductively in parallel the valuation v on $q_0 \cdot Q^*$ and a map $p : q_0 \cdot Q^* \to (Q_\mathsf{A} \cdot Q_\mathsf{B})^* \cdot Q_\mathsf{A}$ ensuring the following, for all $\rho \in q_0 \cdot Q^*$:

a. if $v(\rho) = 1$ then for all $\pi \in Q^*$, we have $v(\rho \cdot \pi) = 1$;

b. if $v(\rho) < 1$, then $p(\rho) \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^* \cdot Q_\mathsf{A}$ is compatible with the strategy $\mathsf{s_B}$ from (q_0, α).

We let $v(q_0) := \alpha + \varepsilon$ and $p(q_0) := (q_0, \alpha) \in Q_\mathsf{A}$. Now, assume that v and p are both defined on a path $\rho \in q_0 \cdot Q^*$. First, if $v(\rho) = 1$, then for all $q \in Q$, we set $v(\rho \cdot q) := 1$, thus ensuring property a. Assume now that $v(\rho) < 1$. For all states $q \in Q$, we define the set of valuations that Player A can choose that make the Player-B strategy $\mathsf{s_B}$ go to q:

$$H_\rho(q) := \{h \in \mathsf{Move}_\mathsf{A}^{\mathsf{gd}}(p(\rho)_{\mathsf{lt}}) \mid \mathsf{s_B}(p(\rho) \cdot (\rho, h)) = q\}$$

Then, we define a function $h_\rho : Q \to [0, 1]$ such that, for all states $q \in Q$:

$$h_\rho(q) := \inf\{h(q) \mid h \in H_\rho(q)\}$$

Furthermore $H_\rho(q) = \emptyset$ means that regardless of what Player A chooses as value for the state q, the Player-B winning strategy $\mathsf{s_B}$ never goes to q. Hence, we set $h_\rho(q) := 1$. Then, we set:

$$v^\rho := \min\left(h_\rho + \frac{\varepsilon}{2^{|\rho|}}, 1\right) : Q \to [0, 1] \tag{2.7}$$

Let us now define $p(\rho \cdot q) \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^* \cdot Q_\mathsf{A}$. If $v(\rho \cdot q) = 1$, then we let $p(\rho \cdot q) := p(\rho) \cdot (\rho, 1) \cdot (\rho \cdot q, 1)$. Assume now that $v(\rho \cdot q) < 1$. That is, $v(\rho \cdot q) = h_\rho(q) + \frac{\varepsilon}{2^{|\rho|}} < 1$ and $H_\rho(q) \neq \emptyset$. In that case, we consider some $h_\rho^q \in H_\rho^q$ such that $h_\rho^q(q) \leq h_\rho(q) + \frac{\varepsilon}{2^{|\rho|+1}} = v(\rho \cdot q) - \frac{\varepsilon}{2^{|\rho|+1}}$. Then, we define:

$$p(\rho \cdot q) := p(\rho) \cdot (\rho, h_\rho^q) \cdot (q, h_\rho^q(q))$$

Since we have $h_\rho^q \in H_\rho^q$, it follows that the finite path $p(\rho \cdot q)$ is compatible with the strategy $\mathsf{s_B}$. The valuation v and the map p are now entirely defined and satisfy properties a. and b. In fact, we obtain for all $\rho \in q_0 \cdot Q^\omega$, that either:

- there is some $i \in \mathbb{N}$ such that, for all $j \geq i$, we have $v(\rho_j) = 1$; or
- the infinite path $p^\omega(\rho) \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^\omega$ is compatible with the strategy $\mathsf{s_B}$ from the state (q_0, α).

Let us show that the valuation v is non-increasing w.r.t. gd. This holds for paths in $Q^+ \setminus q_0 \cdot Q^*$. Consider some $\rho \in q_0 \cdot Q^*$. If $v(\rho) = 1$, it is straightforward that $v(\rho) \geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle][\mathsf{gd}(\rho)]$. Assume now that $v(\rho) < 1$ and let us write

$\rho = \rho' \cdot q$. We have $p(\rho)_{\mathsf{lt}} = (\rho, h^q_{\rho'}(q))$ with $h^q_{\rho'}(q) \leq v(\rho) - \frac{\varepsilon}{2^{|\rho|}}$ by definition of $h^q_{\rho'}$.

Let us show that $h^q_{\rho'}(q) \geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), h_\rho \rangle][\mathsf{gd}(\rho)]$. Assume towards a contradiction that $h^q_{\rho'}(q) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), h_\rho \rangle][\mathsf{gd}(\rho)] - \delta$ for some positive $\delta > 0$. Then, let $h'_\rho : Q \to [0,1]$ be such that $h'_\rho := \max(0, h_\rho - \delta)$. In that case, $h'_\rho \geq h_\rho - \delta$, hence by Lemma 1.10, we have:

$$\mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), h'_\rho \rangle][\mathsf{gd}(\rho)] \geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), h_\rho \rangle][\mathsf{gd}(\rho)] - \delta \geq h^q_{\rho'}(q)$$

It follows that $h'_\rho \in \mathsf{Move}^{\mathsf{gd}}_{\mathsf{A}}(p(\rho)_{\mathsf{lt}})$. Consider now the state $q' \in Q$ equal to $q' := \mathsf{s_B}(p(\rho) \cdot (\rho, h'_\rho))$. In particular, we have $h'_\rho \in H_\rho(q')$. Furthermore, it is not possible that $h'_\rho(q') = 0$. Indeed, since $p(\rho)$ is compatible with $\mathsf{s_B}$ from (q_0, α), the path $p(\rho) \cdot (\rho, h'_\rho) \cdot (q', h'_\rho(q'))$ also is. However, Player A has a winning strategy from any state $(t, 0) \in Q_{\mathsf{A}}$ for $t \in Q$ (see Observation 2.1) and the Player-B strategy $\mathsf{s_B}$ is winning from (q_0, α). Hence, $h'_\rho(q') > 0$. That is, $h'_\rho(q') = h_\rho(q') - \delta < h_\rho(q')$. This is in contradiction with the definition of $h_\rho(q') = \inf\{h'(q') \mid h' \in H_\rho(q')\}$ since $h'_\rho \in H_\rho(q')$. In fact, we have $h^q_{\rho'}(q) \geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), h_\rho \rangle][\mathsf{gd}(\rho)]$.

Furthermore, $v^\rho - \frac{\varepsilon}{2^{|\rho|}} \leq h_\rho$ by Equation 2.7. Recall also that $v(\rho) - \frac{\varepsilon}{2^{|\rho|}} \geq h^q_{\rho'}(q)$. Hence, we obtain, with Lemma 1.10:

$$v(\rho) - \frac{\varepsilon}{2^{|\rho|}} \geq h^q_{\rho'}(q) \geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), h_\rho \rangle][\mathsf{gd}(\rho)] \geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle][\mathsf{gd}(\rho)] - \frac{\varepsilon}{2^{|\rho|}}$$

That is, $v(\rho) \geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v_\rho \rangle][\mathsf{gd}(\rho)]$. As this holds for all $\rho \in q_0 \cdot Q^*$, it follows that the valuation v is non-increasing w.r.t gd.

Consider now an infinite path $\rho \in q_0 \cdot Q^\omega$. If there is some $i \in \mathbb{N}$ such that for all $j \geq i$, we have $v(\rho^{\leq j}) = 1$, it follows that $\mathsf{limsup}\ v(\rho^{\leq i}) = 1 \geq f(\rho)$. Otherwise, the infinite path $p(\rho) \in (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^\omega$ is compatible with the winning Player-B strategy $\mathsf{s_B}$ from the state (q_0, α) with $p(\rho)$ equal to:

$$p(\rho) = (\rho_0, \alpha) \cdot (\rho_0, h^{\rho_1}_{\rho_0}) \cdot (\rho_{\leq 1}, h^{\rho_1}_{\rho_0}(\rho_1)) \cdot (\rho_{\leq 1}, h^{\rho_2}_{\rho_{\leq 1}}) \cdot (\rho_{\leq 2}, h^{\rho_2}_{\rho_{\leq 1}}(\rho_2)) \cdots$$

Therefore, since the Player-B strategy $\mathsf{s_B}$ is winning from (q_0, α), we have $\mathsf{limsup}\ h^{\rho_{i+1}}_{\rho_{\leq i}}(\rho_{i+1}) > f_{\mathcal{C}}(\rho)$. Furthermore, for all $i \in \mathbb{N}$, we have $v(\rho_{\leq i} \cdot \rho_{i+1}) \geq h^{\rho_{i+1}}_{\rho_{\leq i}}(\rho_{i+1})$ by definition of $h^{\rho_{i+1}}_{\rho_{\leq i}}$. Hence, $\mathsf{limsup}\ v(\rho_{\leq i+1}) \geq f_{\mathcal{C}}(\rho)$. As this holds for all paths $\rho \in q_0 \cdot Q^\omega$, it follows that the valuation v is winning for Player B w.r.t. $(q_0, \alpha + \varepsilon)$ and gd. □

We have now all the ingredients to prove results 1.a and 2 of Theorem 2.3. A reader who wants to see the proof of that part of the theorem now can skip the next subsection — which we need to prove result 1.b — go to Subsection 2.4.4 and read the corresponding part of the proof.

2.4.3 . Win/lose objectives with finite representatives for $(\mathsf{K}, \mathsf{col})$ in Q

In this subsection, we focus on how to prove result $1.b$. From a high level perspective, let us explain what we want to prove in this subsection. Consider an arbitrary concurrent game $\mathcal{G} = \langle \mathcal{C}, f \rangle$ without stopping states. Then, the payoff function f only depends on the colors seen in the game, not on the exact sequence of states visited. Hence, it is natural to assume, that to play almost-optimally, when they make a decision, the players only need to know their current position and the sequence of colors seen, not the exact sequence of states visited. This is what we show in this subsection. More precisely, we prove a lemma that can then be used to prove this result. However, we are not able to show it in all generality. We will use two additional assumptions: first, that the game is win/lose (i.e. $f[\mathsf{K}^\omega] \subseteq \{0, 1\}$), and second, that the pair $(\mathsf{K}, \mathsf{col})$ has a finite representative. Let us first describe how we use this first assumption. If the game is win/lose and without stopping states, the winning condition $W_{\mathsf{tb}}(f)$ can reformulated as follows: this consists in the set of infinite paths in $(Q_\mathsf{A} \cdot Q_\mathsf{B})^\omega$ such that either the infinite sequence of states (in Q) has value 1 w.r.t. f (i.e. it is winning for Player A) or the infimum of the values seen in states in Q_B is 0. Consider now another objective $W'_{\mathsf{tb}}(f)$ that is the set of infinite paths in $(Q_\mathsf{A} \cdot Q_\mathsf{B})^\omega$ such that either the infinite sequence of states (in Q) has value 1 w.r.t. f (i.e. it is winning for Player A) or ρ visits a state in Q_A of value 0. In fact, with this slight modification, we do not make it harder for Player A: $\alpha^\mathsf{A}(q, W_{\mathsf{tb}}(f)) \leq \alpha^\mathsf{A}(q, W'_{\mathsf{tb}}(f))$. This is formally proved in Lemma 2.15 below.

Let us first introduce some useful functions.

Definition 2.10. *Consider an arbitrary concurrent arena $\mathcal{C}$ and the turn-based arena $\mathcal{C}^\mathsf{A}_{\mathsf{tb}}$. We denote by $\mathsf{P}_0 \subseteq Q^\uparrow_{\mathsf{tb}}$ the set of finite or infinite paths visiting the value 0: $\mathsf{P}_0 := \{\rho \in Q^\uparrow_{\mathsf{tb}} \mid \exists i < |\rho|,\ \rho_i = (q, 0) \in Q_\mathsf{A}\}$.*

Furthermore, we let $\mathsf{val} : (Q_\mathsf{A} \cdot Q_\mathsf{B})^* \cdot Q_\mathsf{A} \to [0, 1]$ and $\mathsf{sta} : (Q_\mathsf{A} \cdot Q_\mathsf{B})^* \cdot Q_\mathsf{A} \to Q$ be such that:

$$\forall \rho = (q_0, \alpha_0) \cdot (q_0, h_0) \cdots (q_n, \alpha_n) \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^* \cdot Q_\mathsf{A}, \begin{cases} \mathsf{val}(\rho) := \alpha_n \\ \mathsf{sta}(\rho) := q_n \end{cases}$$

We let also $\phi_Q : (Q_\mathsf{A} \cdot Q_\mathsf{B})^\uparrow \cdot Q_\mathsf{A} \to Q^\uparrow$ and $\phi_{[0,1]} : (Q_\mathsf{A} \cdot Q_\mathsf{B})^\uparrow \cdot Q_\mathsf{A} \to [0, 1]^\uparrow$ be such that:

$$\forall \rho = (q_0, \alpha_0) \cdot (q_0, h_0) \cdots \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^\uparrow \cdot Q_\mathsf{A}, \begin{cases} \phi_Q(\rho) := q_0 \cdot q_1 \cdots \in Q^\uparrow \\ \phi_{[0,1]}(\rho) := \alpha_0 \cdot \alpha_1 \cdots \in [0, 1]^\uparrow \end{cases}$$

Let us now formally define this other winning condition for Player A in the game $\mathcal{G}_{\mathsf{tb}}$.

Definition 2.11 (Another winning condition)**.** *Consider an arbitrary concurrent game $\mathcal{G}$, the turn-based arena $\mathcal{C}_{\mathsf{tb}}$ from Definition 2.8. We define the* *winning objective:*

$$W'_{\mathsf{tb}}(f) := \{(\rho_0, \alpha_0) \cdot (\rho_0, h_0) \cdots \in (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^\omega \mid f_{\mathcal{C}}(\rho) = 1 \text{ or } \exists i \in \mathbb{N},\, \alpha_i = 0\}$$

We let $\mathcal{G}'_{\mathsf{tb}} := \langle \mathcal{C}^{\mathsf{A}}_{\mathsf{tb}}, W'_{\mathsf{tb}}(f)\rangle$ *be the corresponding turn-based game.*

With this definition, we have:

Lemma 2.15. *Consider an arbitrary win/lose concurrent game* $\mathcal{G} = \langle \mathcal{C}, f\rangle$ *without stopping states. Let* $q_0 \in Q$*. Then:* $\alpha^{\mathsf{A}}(q_0, W_{\mathsf{tb}}(f)) \leq \alpha^{\mathsf{A}}(q_0, W'_{\mathsf{tb}}(f))$*.*

Proof sketch. For $\alpha := \alpha^{\mathsf{A}}(q_0, W_{\mathsf{tb}}(f))$, we let $0 < \varepsilon < \alpha$ and we want to show that $\alpha - \varepsilon \leq \alpha^{\mathsf{A}}(q_0, W'_{\mathsf{tb}}(f))$. Consider a winning Player-A strategy s_{A} in the game $\mathcal{G}^{\mathsf{A}}_{\mathsf{tb}}$ from the staring state $(q_0, \alpha - \frac{\varepsilon}{2})$. Our goal is define a Player-A strategy s'_{A} that is winning from the state $(q_0, \alpha - \varepsilon)$ in the game $\mathcal{G}'_{\mathsf{tb}}$. To do so, we will define this strategy s'_{A} from s_{A} by subtracting (when possible) $\varepsilon/2$ to all values of the states that it chooses. That way, for all infinite paths ρ compatible with s'_{A} from $(q_0, \alpha - \varepsilon)$, there is going to be an infinite path $p(\rho)$ that is compatible with s_{A} from $(q_0, \alpha - \frac{\varepsilon}{2})$ such that the sequences of states (in Q) seen in ρ and $p(\rho)$ are the same. Furthermore, if ρ never visits a state in Q_{A} of value 0, then the infimum of the values seen in $p(\rho)$ is at least $\varepsilon/2$. Therefore, since s_{A} is winning $\mathcal{G}_{\mathsf{tb}}$ from $(q_0, \alpha - \frac{\varepsilon}{2})$, then s'_{A} is winning in $\mathcal{G}'_{\mathsf{tb}}$ from $(q_0, \alpha - \varepsilon)$. □

Let us now formally prove this lemma.

Proof. If $\alpha^{\mathsf{A}}(q_0, W_{\mathsf{tb}}(f)) = 0$, this straightforwardly holds. Assume now that $\alpha^{\mathsf{A}}(q_0, W_{\mathsf{tb}}(f)) \neq 0$, which we denote by $\alpha := \alpha^{\mathsf{A}}(q_0, W_{\mathsf{tb}}(f))$. Let $0 < \varepsilon < \alpha$. Let us show that $\alpha - \varepsilon \leq \alpha^{\mathsf{A}}(q_0, W'_{\mathsf{tb}}(f))$. Consider a winning Player-A strategy s_{A} in the game $\mathcal{G}^{\mathsf{A}}_{\mathsf{tb}}$ from the starting state $(q_0, \alpha - \frac{\varepsilon}{2})$.

Formally, we define inductively a Player-A strategy s'_{A} in the arena $\mathcal{C}^{\mathsf{A}}_{\mathsf{tb}}$ along with a map p taking a finite path ρ in $(Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^* \cdot Q_{\mathsf{A}} \setminus \mathsf{P}_0$ that is compatible with s'_{A} from $(q_0, \alpha - \varepsilon)$ and returning a finite path $p(\rho) \in (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^* \cdot Q_{\mathsf{A}}$ ensuring:

a. $\phi_Q(p(\rho)) = \phi_Q(\rho)$, i.e. the finite sequence of states in Q seen is the same in $p(\rho)$ and ρ;

b. $p(\rho) \notin \mathsf{P}_0$ is compatible with the strategy s_{A} from $(q_0, \alpha - \varepsilon/2)$;

c. $\mathsf{val}(p(\rho)) = \mathsf{val}(\rho) + \varepsilon/2$;

d. for all $\rho \in (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^* \cdot Q_{\mathsf{A}} \cap \mathsf{P}_0$ compatible with s'_{A}, we have $\mathsf{val}(\rho) = 0$;

e. for all $\rho \in (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^* \cdot Q_{\mathsf{A}} \setminus \mathsf{P}_0$ compatible with s'_{A}, we have $\mathsf{s}'_{\mathsf{A}}(\rho) := \max(\mathsf{s}_{\mathsf{A}} \circ p(\rho) - \frac{\varepsilon}{2}, 0) : Q \to [0,1]$.

Initially, we set $p((q_0, \alpha-\varepsilon)) := (q_0, \alpha-\varepsilon/2)$ and $\mathsf{s}'_\mathsf{A}((q_0, \alpha-\varepsilon)) := \max(\mathsf{s}_\mathsf{A} \circ p(q_0, \alpha - \varepsilon) - \frac{\varepsilon}{2}, 0) : Q \to [0,1]$ thus ensuring properties $a. - e.$

Assume now that s'_A and p, for some $n \in \mathbb{N}$, are defined on paths of length at most $2 \cdot n + 1$ while ensuring properties $a. - e.$. Consider a path

$$\rho = \rho' \cdot (\pi, \mathsf{s}'_\mathsf{A}(\rho')) \cdot (\pi \cdot q', \mathsf{s}'_\mathsf{A}(\rho')(q)) \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^{n+1} \cdot Q_\mathsf{A}$$

compatible with s'_A. If $\rho \in \mathsf{P}_0$, then by assumption $\mathsf{val}(\rho) = 0$. Hence, we let $\mathsf{s}_\mathsf{A}(\rho) := 0 : Q \to [0,1] \in \mathsf{Move}_\mathsf{A}(\rho_{\mathsf{lt}})$ ensuring property d. Now, assume that $\rho \notin \mathsf{P}_0$. In that case, we have $\mathsf{s}'_\mathsf{A}(\rho')(q) > 0$, that is $\mathsf{s}'_\mathsf{A}(\rho')(q) = \mathsf{s}_\mathsf{A} \circ p(\rho')(q) - \frac{\varepsilon}{2} > 0$. We let

$$p(\rho) := p(\rho') \cdot (\pi, \mathsf{s}_\mathsf{A} \circ p(\rho')) \cdot (\pi \cdot q, \mathsf{s}_\mathsf{A} \circ p(\rho')(q)) \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^{n+1} \cdot Q_\mathsf{A}$$

thus ensuring properties a.-c. Furthermore, defining $\mathsf{s}'_\mathsf{A}(\rho) := \max(\mathsf{s}_\mathsf{A} \circ p(\rho) - \frac{\varepsilon}{2}, 0) : Q \to [0,1]$ ensures property e. This concludes the definitions of the strategy s'_A and the map p.

Now, let us show that the strategy s'_A is winning in the game $\mathcal{G}'_{\mathsf{tb}}$ from the state $(q_0, \alpha - \varepsilon)$. Consider some infinite path $\rho \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^\omega$ that is compatible with s'_A in $\mathcal{C}^\mathsf{A}_{\mathsf{tb}}$ from $(q_0, \alpha - \varepsilon)$. If $\rho \in \mathsf{P}_0$, then $\rho \in W'_{\mathsf{tb}}(f)$. Now, assume that $\rho \notin \mathsf{P}_0$. It follows that, for all $i \in \mathbb{N}$, we have $\rho_{\leq i}$ compatible with s'_A and $\rho_{\leq i} \notin \mathsf{P}_0$. The map p is therefore defined on all prefixes $\rho_{\leq i}$ for $i \in \mathbb{N}$. Hence, we can consider the infinite path $p(\rho) \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^\omega$. By property b., we have $p(\rho) \notin \mathsf{P}_0$ and $p(\rho)$ compatible with the strategy s_A from the state $(q_0, \alpha - \varepsilon/2)$. Furthermore, by property c., for all $i \in \mathbb{N}$, we have $\mathsf{val}(p(\rho)_{\leq 2i+1}) = \mathsf{val}(p(\rho_{\leq 2i+1})) = \mathsf{val}(\rho_{\leq 2i+1}) + \varepsilon/2 \geq \varepsilon/2$. It follows that $\mathsf{limsup}\ \phi_{[0,1]}(p(\rho)) \geq \varepsilon/2 > 0$. Since the strategy s_A is winning from the state $(q_0, \alpha - \varepsilon/2)$ for the winning condition $W_{\mathsf{tb}}(f)$, we have $p(\rho) \in W_{\mathsf{tb}}(f)$. That is, $0 < \mathsf{limsup}\ \phi_{[0,1]}(p(\rho)) \leq f_\mathcal{C}(\phi_Q(p(\rho)))$. We can conclude that, since f is win/lose and there are no stopping states, $f_\mathcal{C}(\phi_Q(p(\rho))) = 1$. As $\phi_Q(p(\rho)) = \phi_Q(\rho)$ (by property $a.$), it follows that $\rho \in W'_{\mathsf{tb}}(f)$. As this holds for all infinite paths ρ compatible with s'_A from $(q_0, \alpha - \varepsilon)$, it follows that this strategy is winning from $(q_0, \alpha - \varepsilon)$. Hence, $\alpha^\mathsf{A}(q_0, W_{\mathsf{tb}}(f)) - \varepsilon = \alpha - \varepsilon \leq \alpha^\mathsf{A}(q_0, W'_{\mathsf{tb}}(f))$. Since this holds for all $\varepsilon > 0$, it follows that $\alpha^\mathsf{A}(q_0, W_{\mathsf{tb}}(f)) \leq \alpha^\mathsf{A}(q_0, W'_{\mathsf{tb}}(f))$. □

Lemma 2.16. *Consider an arbitrary win/lose concurrent game $\mathcal{G}$. Assume that $(\mathsf{K}, \mathsf{col})$ has finite representatives in Q. Then, for all $0 < \alpha < \alpha^\mathsf{A}(q_0, W_{\mathsf{tb}}(f))$, there is a $(\mathsf{K}, \mathsf{col})$-uniform valuation $v : Q^+ \to [0,1]$ that is winning for Player A w.r.t. (q_0, α) and A.*

Proof sketch. To build this valuation v that is winning for Player A w.r.t. (q_0, α), by Lemma 2.15, we can use a Player-A strategy that is winning in the game $\mathcal{G}'_{\mathsf{tb}}$ from the state (q_0, α). However, since the valuation v has to be

$(\mathsf{K}, \mathsf{col})$-uniform, we need to define it on $\mathsf{K}^* \times Q$. To do so, given any pair $(\gamma, q) \in \mathsf{K}^* \times Q$, we will consider all the possible paths $\pi \in Q^*$ whose sequence of colors may be equal to γ. Since $(\mathsf{K}, \mathsf{col})$ has finite representatives in Q, there are only finitely many of them. By appropriately choosing the finite path π to consider, it is possible to define v while ensuring that the valuation v that we define is non-decreasing w.r.t. A. In addition, for all $\rho \in q_0 \cdot Q^\omega$, we can build an infinite path in the arena $\mathcal{C}^{\mathsf{A}}_{\mathsf{tb}}$ whose sequence of colors corresponds to the sequence of colors ρ that is compatible with the strategy $\mathsf{s_A}$ from (q_0, α). We achieve this by using the finite representatives assumption (with König's lemma-like argument.). In addition, since we have changed the objective into $W_{\mathsf{tb}}(f)'$, we can show that if $f(\rho) < 1$, then there is some $i \in \mathbb{N}$ such that, for all $j \geq i$, we have $v(\rho_{\leq j}) = 0$. Therefore, the superior limit of $v(\rho_{\leq i})$ is equal to 0. This allows us to conclude that the valuation v that we have defined is winning for Player A w.r.t. (q_0, α). □

Let us now formally prove this lemma.

Proof. Consider such an $\alpha < \alpha^{\mathsf{A}}(q_0, W_{\mathsf{tb}}(w))$. By Lemma 2.15, we have $\alpha < \alpha^{\mathsf{A}}(q_0, W'_{\mathsf{tb}}(w))$. Hence, we can consider a Player-A strategy $\mathsf{s_A}$ that is winning from (q_0, α) in the game $\mathcal{G}'_{\mathsf{tb}}$. Consider the function $p = p^{q_0,\alpha}_{\mathsf{s_A}} : q_0 \cdot Q^* \to (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^* \cdot Q_{\mathsf{A}}$ from Definition 2.9. Then, for all $\gamma \in \mathsf{K}^*$ and $q \in Q$, we define the set:

$$C_{\gamma,q} := \{\pi = \pi' \cdot q \in q_0 \cdot Q^* \cdot q \mid \mathsf{col}^*(\pi') = \gamma \ \wedge \ p(\pi) \notin \mathsf{P}_0\}$$

Note that for all $\gamma \in \mathsf{K}^*$ and $q \in Q$, the set $C_{\gamma,q}$ is finite since $(\mathsf{K}, \mathsf{col})$ has finite representatives in Q. Hence, we can define the maximum value $\alpha_{\gamma,q} \in [0,1]$ achieved by paths in $C_{\gamma,q}$:

$$\alpha_{\gamma,q} := \max\{\alpha \in [0,1] \mid \exists \pi \in C_{\gamma,q},\ \mathsf{val} \circ p(\pi) = \alpha\}$$

Whenever $C_{\gamma,q}$ is the empty set, we set $\alpha_{\gamma,q} := 0$. Finally, consider a function $\iota : \mathsf{K}^* \times Q \to Q^+$ such that, for all $(\gamma, q) \in \mathsf{K}^* \times Q$, if $C_{\gamma,q} \neq \emptyset$, then $\iota(\gamma, q) \in C_{\gamma,q}$ and:

$$\mathsf{val} \circ p(\iota(\gamma, q)) = \alpha_{\gamma,q} \tag{2.8}$$

Note that, by definition of the set $C_{\gamma,q}$, we have:

$$\mathsf{col}^* \circ \iota(\gamma, q) = \gamma \cdot \mathsf{col}(q) \tag{2.9}$$

We can now define a valuation of finite sequences of states $v : Q^+ \to [0,1]$. First, we let $v(q_0) := \alpha$ and $v(q) := 0$ for all $q \in Q \setminus \{q_0\}$. Then, for all $(\gamma \cdot q) \in \mathsf{K}^* \times Q$, we define the valuation $v^{\gamma \cdot q} : Q \to [0,1]$ in the following way:

$$v^{\gamma,q} := \begin{cases} 0 & : Q \to [0,1] \text{ if } C_{\gamma,q} = \emptyset \\ \mathsf{s_A} \circ p(\iota(\gamma, q)) & : Q \to [0,1] \text{ otherwise} \end{cases}$$

Then, for all $\rho \cdot q \in q_0 \cdot Q^*$, we have $v^{\rho \cdot q} := v^{\mathsf{col}^*(\rho) \cdot q}$. Hence, by definition, v is $(\mathsf{K}, \mathsf{col})$-uniform. Let us show that v is a winning valuation for Player A w.r.t. (q_0, α) and A.

First, we show that it is non-decreasing w.r.t. A. We have $v^{q_0} = \mathsf{s_A}((q_0, \alpha)) \in \mathsf{Move}^{\mathsf{A}}_{\mathsf{A}}((q_0, \alpha))$. Hence, $v(q_0) = \alpha \leq \mathsf{out}[\langle \mathsf{F}(q_0), v^{q_0} \rangle]$. Furthermore, $v(q) = 0$ for all $q \in Q \setminus \{q_0\}$. Hence, the valuation v is non-decreasing at all states $q \in Q$. Now, let $\rho := \rho' \cdot q' \cdot q \in Q \cdot Q^+$ with $\rho' \in Q^*$, $\gamma := \mathsf{col}^*(\rho') \in \mathsf{K}^*$ and $\gamma' := \mathsf{col}^*(\rho' \cdot q') = \gamma \cdot \mathsf{col}(q') \in \mathsf{K}^+$. If $v(\rho) = 0$, then $v(\rho) \leq \mathsf{out}[\langle \mathsf{F}(\rho), v^{\rho} \rangle]$ holds straightforwardly. Assume now that $v(\rho) > 0$. That is, $v^{\rho' \cdot q'}(q) > 0$. It follows that $C_{\gamma, q'} \neq \emptyset$. Hence, we have $\iota(\gamma, q') \in C_{\gamma, q'}$ and $v^{\rho' \cdot q'} = \mathsf{s_A} \circ p(\iota(\gamma, q'))$. Furthermore,

$$\begin{aligned} p(\iota(\gamma, q') \cdot q) &= p(\iota(\gamma, q')) \cdot (\iota(\gamma, q'), \mathsf{s_A} \circ p(\iota(\gamma, q'))) \cdot (\iota(\gamma, q') \cdot q, \mathsf{s_A} \circ p(\iota(\gamma, q'))(q)) \\ &= p(\iota(\gamma, q')) \cdot (\iota(\gamma, q'), v^{\gamma, q'}) \cdot (\iota(\gamma, q') \cdot q, v^{\gamma, q'}(q)) \end{aligned}$$

with $v^{\gamma \cdot q'}(q) = v(\rho) > 0$. Since $p(\iota(\gamma, q')) \notin \mathsf{P}_0$, we have $p(\iota(\gamma, q') \cdot q) \notin \mathsf{P}_0$. In addition, since by Equation 2.9 we have $\mathsf{col}^* \circ \iota(\gamma, q') = \gamma \cdot \mathsf{col}(q') = \gamma'$, it follows that $\iota(\gamma, q') \cdot q \in C_{\gamma', q}$. Hence, by definition, $\alpha_{\gamma', q} \geq \mathsf{val} \circ p(\iota(\gamma, q') \cdot q) = v(\rho)$. Furthermore, $v^{\gamma', q} = \mathsf{s_A} \circ p(\iota(\gamma', q)) \in \mathsf{Move}^{\mathsf{A}}_{\mathsf{A}}(p(\iota(\gamma', q))_{\mathsf{lt}}) = \mathsf{Move}^{\mathsf{A}}_{\mathsf{A}}((\iota(\gamma, q') \cdot q, \alpha_{\gamma', q}))$ by Equation 2.8. In other words, $v(\rho) \leq \alpha_{\gamma', q} \leq \mathsf{val}[\langle \mathsf{F}(q), v^{\gamma', q} \rangle][\mathsf{A}] = \mathsf{val}[\langle \mathsf{F}(q), v^{\rho} \rangle][\mathsf{A}]$. That is, v is non-decreasing w.r.t. A at ρ.

Consider now some $\rho \in q_0 \cdot Q^\omega$. Let $N_{>0} := \{i \in \mathbb{N} \mid C_{\mathsf{col}^*(\rho_{\leq i}), \rho_{i+1}} \neq \emptyset\}$. For all $k \notin N_{>0}$, we have $v^{\mathsf{col}^*(\rho_{\leq k}), \rho_{k+1}}(\rho_{k+2}) = v(\rho_{\leq k+2}) = 0$. Hence, if $N_{>0}$ is finite, then there is some $k \in \mathbb{N}$ such that $v(\rho_{\leq i}) = 0$ for all $i \geq k$. Hence, $\mathsf{limsup}(v(\rho_{\leq i}))_{i \in \mathbb{N}} = 0$. Now, assume that $N_{>0}$ is infinite. Let us define inductively a sequence of states $\pi \in Q^\omega$ such that, for all $i \in \mathbb{N}$, we have:

a. $\mathsf{col}^*(\pi_{\leq i}) = \mathsf{col}^*(\rho_{\leq i})$;

b. $p(\pi_{\leq i}) \notin \mathsf{P}_0$;

c. the set $N^i_{>0} := \{j \geq i \mid C_{\mathsf{col}^*(\rho_{\leq j}), \rho_{j+1}} \cap \pi_{\leq i} \cdot Q^* \neq \emptyset\}$ is infinite.

Initially, $\pi_0 = q_0$ and all properties are ensured since $N_{>0}$ is infinite. Let us now assume that π is defined up to index $i \in \mathbb{N}$ and that all properties above are ensured up to that index. Let us define $\pi_{i+1} \in Q$. Let $\tilde{N}^i_{>0} := \{j \geq i+1 \mid C_{\mathsf{col}^*(\rho_{\leq j}), \rho_{j+1}} \cap \pi_{\leq i} \cdot Q^* \neq \emptyset\}$. By assumption, this set is infinite. Now, let $c := \mathsf{col}(\rho_{i+1}) \in \mathsf{K}$. For all $q \in \mathsf{col}^{-1}[c]$, we consider the set $N^i_{>0}(q) := \{j \geq i+1 \mid C_{\mathsf{col}^*(\rho_{\leq j}), \rho_{j+1}} \cap \pi_{\leq i} \cdot q \cdot Q^* \neq \emptyset\}$. We have:

$$\tilde{N}^i_{>0} = \bigcup_{q \in \mathsf{col}^{-1}[c]} N^i_{>0}(q)$$

Since $(\mathsf{K}, \mathsf{col})$ has finite representatives in Q, the set $\mathsf{col}^{-1}[c]$ is finite. Hence, there is some $q \in \mathsf{col}^{-1}[c]$ such that $N^i_{>0}(q)$ is infinite. Then, we set $\pi_{i+1} := q$.

With this choice, we have $N_{>0}^{i+1} = N_{>0}^{i}(q)$. Hence, it is infinite and property $c.$ is ensured. It also follows that there is some $j \geq i+1$, such that $C_{\mathsf{col}^*(\rho_{\leq j}),\rho_{j+1}} \cap \pi_{\leq i+1} \cdot Q^* \neq \emptyset$. Consider $\theta = \pi_{\leq i+1} \cdot \pi' \in C_{\mathsf{col}^*(\rho_{\leq j}),\rho_{j+1}} \cap \pi_{\leq i+1} \cdot Q^*$. By definition of $C_{\mathsf{col}^*(\rho_{\leq j}),\rho_{j+1}}$, we have $p(\theta) \notin \mathsf{P}_0$. Since $\pi_{\leq i+1}$ is a prefix of θ, it follows that $p(\pi_{\leq i+1}) \notin \mathsf{P}_0$. Hence, property $b.$ is ensured. Furthermore, we have $\mathsf{col}^*(\pi_{\leq i+1}) = \mathsf{col}^*(\pi_{\leq i}) \cdot \mathsf{col}(\pi_{i+1}) = \mathsf{col}^*(\pi_{\leq i}) \cdot c = \mathsf{col}^*(\rho_{\leq i}) \cdot \mathsf{col}(\rho_{i+1}) = \mathsf{col}^*(\rho_{\leq i+1})$ and property $a.$ is ensured. This concludes the definition of $\pi \in Q^\omega$.

The infinite path $p(\pi) \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^\omega$ is compatible with the strategy $\mathsf{s_A}$ from (q_0, α) (recall Observation 2.2). Hence, since the strategy $\mathsf{s_A}$ is winning from (q_0, α), we have $p(\pi) \in W'_{\mathsf{tb}}(f)$. However, because of property $b.$, we have $p(\pi) \notin \mathsf{P}_0$. Hence, $f(\pi) = f(\phi_Q(p(\pi))) = 1$. Since $\mathsf{col}^\omega(\pi) = \mathsf{col}^\omega(\rho)$, it follows that $f(\rho) = 1$. That is, we have $\mathsf{limsup}(v(\rho_{\leq i}))_{i \in \mathbb{N}} \leq f(\rho)$ and the valuation v is winning for Player A w.r.t. (q_0, α) and A. □

2.4.4 . Proof of Theorem 2.3

Proof. Consider an arbitrary concurrent game $\mathcal{G}$ and assume that it is supremized by a collection $(S_q^\mathsf{A})_{q \in Q}$ of GF-strategies w.r.t. Player A and supremized by a collection $(S_q^\mathsf{B})_{q \in Q}$ of GF-strategies w.r.t. Player B. Note that such collections of GF-strategies always exist since $(\Sigma_\mathsf{A}^q)_{q \in Q}$ (resp. $(\Sigma_\mathsf{B}^q)_{q \in Q}$) works for Player A (resp. B).

Let $q \in Q$ be a starting state. Let $\alpha_q := \alpha^\mathsf{A}(q, W_{\mathsf{tb}}(f))$. Let us show that $\alpha_q = \chi_\mathcal{G}[\mathsf{A}](q)$. Let $0 < \varepsilon < \alpha_q$. By definition of $\alpha^\mathsf{A}(q, W_{\mathsf{tb}}(f))$ (recall Observation 2.1), Player A is winning from $(q, \alpha_q - \frac{\varepsilon}{2})$ in $\mathcal{G}_{\mathsf{tb}}^\mathsf{A}$. Hence, by Lemma 2.13, she has a winning valuation w.r.t. $(q, \alpha_q - \frac{\varepsilon}{2})$ and A. It follows, from Lemma 2.12 that Player A has a strategy $\mathsf{s}_\mathsf{A}^{q,\varepsilon}$ generated by $(S_q^\mathsf{A})_{q \in Q}$ whose value is at least $\alpha_q - \varepsilon$ from q in $\mathcal{G}$ against any Player-B strategy in $\mathsf{Opnt}(\mathsf{A}) = \mathsf{S}_\mathsf{B}^\mathcal{C}$. Hence, $\alpha_q \leq \chi_\mathcal{G}[\mathsf{A}](q)$.

Now, consider any Player-A strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$. Let $\alpha_q^{\mathsf{s_A}} := \alpha^{\mathsf{s_A}}(q, W_{\mathsf{tb}}(f))$. Let us show that $\chi_\mathcal{G}[\mathsf{s_A}](q) \leq \alpha_q^{\mathsf{s_A}}$. Let $0 < \varepsilon < 1 - \alpha_q^{\mathsf{s_A}}$. By definition of $\alpha_q^{\mathsf{s_A}}$, Player B is winning from $(q, \alpha_q^{\mathsf{s_A}} + \frac{\varepsilon}{3})$ in $\mathcal{G}_{\mathsf{tb}}^{\mathsf{s_A}}$. Hence, by Lemma 2.14, she has a winning valuation w.r.t. $(q, \alpha_q^{\mathsf{s_A}} + \frac{2 \cdot \varepsilon}{3})$ and $\mathsf{s_A}$. It follows, from Lemma 2.12 that Player B has a strategy generated by $(S_q^\mathsf{B})_{q \in Q}$ whose value is at least $\alpha_q^{\mathsf{s_A}} + \varepsilon$ from q_0 against all strategies in $\mathsf{Opnt}(\mathsf{s_A}) = \mathsf{s_A}$. Hence, $\chi_\mathcal{G}[\mathsf{s_A}](q) \leq \alpha_q^{\mathsf{s_A}} + \varepsilon$. As this holds for all $\varepsilon > 0$, it follows that $\chi_\mathcal{G}[\mathsf{s_A}](q) \leq \alpha_q^{\mathsf{s_A}}$. Furthermore, by Observation 2.1, we have $\alpha_q^{\mathsf{s_A}} = \alpha^{\mathsf{s_A}}(q, W_{\mathsf{tb}}(w)) \leq \alpha^\mathsf{A}(q, W_{\mathsf{tb}}(w)) = \alpha_q$. That is, $\chi_\mathcal{G}[\mathsf{s_A}](q) \leq \alpha_q$. As this holds for all Player-A strategies $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$, it follows that $\chi_\mathcal{G}[\mathsf{A}](q) \leq \alpha_q$. Overall, we obtain $\chi_\mathcal{G}[\mathsf{A}](q) = \alpha_q = \alpha^\mathsf{A}(q, W_{\mathsf{tb}}(w))$. This proves result $1.a$ since, for all $\varepsilon > 0$, the Player-A strategy $\mathsf{s}_\mathsf{A}^\varepsilon \in \mathsf{S}_\mathcal{C}^A$ such that, for all $q \in Q$ and $\rho \in Q^*$, $\mathsf{s}_\mathsf{A}^\varepsilon(q \cdot \rho) := \mathsf{s}_\mathsf{A}^{q,\varepsilon}(q \cdot \rho)$ is generated by $(S_q^\mathsf{A})_{q \in Q}$ and is ε-optimal.

We can prove similarly — by using the counterparts of the lemma cited above for Player B — that $\chi_\mathcal{G}[\mathsf{B}](q) = \alpha^\mathsf{B}(q, W_{\mathsf{tb}}(f))$. Result 2 is then a direct consequence of Observation 2.1: $\alpha^\mathsf{A}(q, W_{\mathsf{tb}}(f)) = \alpha^\mathsf{B}(q, W_{\mathsf{tb}}(f))$ as soon as all local interactions are valuable. Furthermore, result 1.a for Player B can be deduced from how strategies of values ε-close to α_q are built above.

Consider now result 1.b. Assume that $(\mathsf{K}, \mathsf{col})$ has finite representatives in Q and that the payoff function f is win/lose. Also assume for now that there is no stopping state in $\mathcal{G}$. We cannot proceed exactly like we did for result 1.a where, after exhibiting, for each state $q \in Q$, a Player-A strategy that is ε-optimal from q, we then glue these strategies together to form an ε-optimal strategy. The issue is that, to establish result 1.b, we need the obtained strategy to be $(\mathsf{K}, \mathsf{col})$-uniform. Hence, let $k \in \mathsf{K}$ and $Q_k := \mathsf{col}^{-1}[k] \subseteq Q$. We let $n_k := |Q_k| \in \mathbb{N}$. Let us build a game $\mathcal{G}_k$ that is identical to the game $\mathcal{G}$ except that we have added a trivial state q_k — whose color is new and has no impact on the winner of the game — from which there is probability $\frac{1}{n_k}$ to go to any state in Q_k. We denote by K' and col' the set of colors and coloring function in the game $\mathcal{G}_k$. Consider some $\varepsilon > 0$ and let $\varepsilon_k := \frac{\varepsilon}{n_k} > 0$. By Lemma 2.16 and what we have shown above applied to the game $\mathcal{G}_k$, Player A has a winning valuation w.r.t. $(q_k, \chi_{\mathcal{G}_k}[\mathsf{A}](q_k) - \frac{\varepsilon}{2})$ and the guard $\mathsf{gd} = \mathsf{A}$ that is $(\mathsf{K}', \mathsf{col}')$-uniform. It follows, from Lemma 2.12 that Player A has a $(\mathsf{K}', \mathsf{col}')$-uniform strategy $\mathsf{s}_{\mathsf{A}}^{\varepsilon_k}$ generated by $(S_q^{\mathsf{A}})_{q \in Q}$ whose value is at least $\chi_{\mathcal{G}_k}[\mathsf{A}](q_k) - \varepsilon_k$ from q_k. By definition of the game $\mathcal{G}_k$, the strategy $\mathsf{s}_{\mathsf{A}}^{\varepsilon_k, q_k}$ is therefore $(\mathsf{K}, \mathsf{col})$-uniform and ε-optimal from all states in Q_k. This can be done for all $k \in \mathsf{K}$. It follows that the Player-A strategy $\mathsf{s}_{\mathsf{A}}^{\varepsilon} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ such that, for all $q \in Q$ and $\rho \in Q^*$, we have $\mathsf{s}_{\mathsf{A}}^{\varepsilon} := \mathsf{s}_{\mathsf{A}}^{\varepsilon_{\mathsf{col}(q)}, \mathsf{col}(q)}(q \cdot \rho)$ is $(\mathsf{K}, \mathsf{col})$-uniform and ε-optimal. Finally, assume that the game $\mathcal{G}$ has stopping states. It suffices to add two fresh states $\top$ and $\bot$ colored with two fresh colors and to modify the win/lose payoff function f so that reaching $\top$ (resp. $\bot$) leads to value 1 (resp. 0). Furthermore, we replace each stopping state $q \in Q_{\mathsf{s}}$ with a trivial state that leads with probability $\mathsf{val}(q)$ to the state $\top$ and with probability $1 - \mathsf{val}(q)$ to the state $\bot$. This modification does not change the values of any states. We then can apply the result to this new game.

We proved result 1.b for Player A. However, the assumptions for this result — recall, that the game is win/lose and that $(\mathsf{K}, \mathsf{col})$ has finite representatives in Q — do not depend on the player considered. Hence, we can obtain the same result for Player B (up to reversing the roles of the players in all the proofs described in this section). □

We conclude by an application to turn-based games. Indeed, all turn-based interactions are supremized by deterministic GF-strategies. Hence, we obtain the corollary below.

Corollary 2.17. *All turn-based games are valuable and for all $\varepsilon > 0$, for both players, ε-optimal strategies can be found among deterministic strategies.*

This result was stated in [55, Theorem 1] and [56, Lemma 11]. In both cases, the authors suggest that it could be derived by “closely examining”

Martin's proof and realizing that turn-based game forms are supremized by deterministic GF-strategies (which is true). What we have done in this chapter formally proves this result.

2.5 Application: action-strategies

In this section, we consider standard games with richer strategies than the ones we have considered so far. Recall Definition 1.26, a strategy is a function mapping a finite non-empty sequence of states to a GF-strategy. Here, we consider strategies that not only depend on the sequence of states but also on the actions played by the players. These will be called action-strategies, whereas the strategies we have considered so far in this thesis will be called (in this section only) state-strategies. The goal of this section is to properly define action-strategies along with the corresponding notion of outcome. We then use Theorem 2.3 to show that, under specific conditions, concurrent games with action-strategies have a value and that this value is equal to the value with state-strategies. However, we exhibit a game where the values of state- and action-strategies are equal while there is an optimal strategy among action-strategies, but there is none among state-strategies.

We will use the definitions from Definition 1.8 of the projection function $\phi_{Q,Q_{\mathsf{Act}}}$ and of the payoff function $(f_\mathcal{C})_{Q,Q_{\mathsf{Act}}} : (Q \cup Q_{\mathsf{Act}})^\omega \to [0,1]$ obtained from a payoff function $f_\mathcal{C} : Q^\omega \to [0,1]$.

2.5.1 . Definitions

We first define below the set of admissible sequences on which action-strategies will be defined. Informally, these admissible sequences are the sequences of the following shape: $q \cdot (q,a,b) \cdot q' \cdot (q',a',b') \cdots$ with $q, q' \in Q$, $(a,b) \in \mathsf{Act}^q_\mathsf{A} \times \mathsf{Act}^q_\mathsf{B}$ and $(a',b') \in \mathsf{Act}^{q'}_\mathsf{A} \times \mathsf{Act}^{q'}_\mathsf{B}$.

Definition 2.12 (State and action sequences, Admissible sequences). *Consider a standard concurrent arena $\mathcal{C}$. We let:*

$$Q_{\mathsf{Act}} := \bigcup_{q \in Q} (\{q\} \times \mathsf{Act}^q_\mathsf{A} \times \mathsf{Act}^q_\mathsf{B})$$

We let $\mathsf{SeqAdm}^Q_\mathcal{C} \subseteq (Q \cdot Q_{\mathsf{Act}})^* \cdot Q$ *be such that:*

$$\begin{aligned}\mathsf{SeqAdm}^Q_\mathcal{C} :=& \{\rho = q_0 \cdot (q_0,a_0,b_0) \cdot q_1 \cdots (q_{n-1},a_{n-1},b_{n-1}) \cdot q_n \in (Q \cdot Q_{\mathsf{Act}})^* \cdot Q \\ & \mid \forall 0 \leq i \leq n-1,\ q_i \in Q, (a_i,b_i) \in \mathsf{Act}^{q_i}_\mathsf{A} \times \mathsf{Act}^{q_i}_\mathsf{B},\ q_n \in Q\}\end{aligned}$$

and $\mathsf{SeqAdm}^{Q_{\mathsf{Act}}}_\mathcal{C} \subseteq (Q \cdot Q_{\mathsf{Act}})^+$ *be such that:*

$$\begin{aligned}\mathsf{SeqAdm}^{Q_{\mathsf{Act}}}_\mathcal{C} :=& \{\rho = q_0 \cdot (q_0,a_0,b_0) \cdot q_1 \cdots (q_{n-1},a_{n-1},b_{n-1}) \in (Q \cdot Q_{\mathsf{Act}})^+ \\ & \mid \forall 0 \leq i \leq n-1,\ q_i \in Q, (a_i,b_i) \in \mathsf{Act}^{q_i}_\mathsf{A} \times \mathsf{Act}^{q_i}_\mathsf{B}\}\end{aligned}$$

We also let $\mathsf{SeqAdm}_{\mathcal{C}} := \mathsf{SeqAdm}_{\mathcal{C}}^{Q} \uplus \mathsf{SeqAdm}_{\mathcal{C}}^{Q_{\mathsf{Act}}}$. *Furthermore, we let* $\mathsf{SeqAdm}_{\mathcal{C}}^{\omega} \subseteq Q \cdot Q_{\mathsf{Act}}^{\omega}$ *be equal to:*

$$\mathsf{SeqAdm}_{\mathcal{C}}^{\omega} := \{\rho \in (Q \cdot Q_{\mathsf{Act}})^{\omega} \mid \forall i \in \mathbb{N}, \rho_{\leq i} \in \mathsf{SeqAdm}_{\mathcal{C}}\}$$

We can now define formally the notion of action-strategies.

Definition 2.13 (Action strategies)**.** *Consider a standard concurrent arena* $\mathcal{C}$. *Player-*A action-strategies *are maps*

$$\mathsf{s}_{\mathsf{A}} : \mathsf{SeqAdm}_{\mathcal{C}}^{Q} \to \cup_{q \in Q} \Sigma_{\mathsf{A}}(\mathsf{F}(q))$$

such that, for all $\rho \in \mathsf{SeqAdm}_{\mathcal{C}}^{Q}$, *we have* $\mathsf{s}_{\mathsf{A}}(\rho) \in \Sigma_{\mathsf{A}}(\mathsf{F}(\rho_{\mathsf{lt}}))$. *We denote by* $\mathsf{S}_{\mathsf{A}}^{\mathcal{C},\mathsf{Act}}$ *the corresponding set of strategies in the arena* $\mathcal{C}$. *From a Player-*A *strategy* $\mathsf{s}_{\mathsf{A}} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$, *we build the action-strategy* $\mathsf{s}_{\mathsf{A}}^{\mathsf{Sta}}$ *such that* $\mathsf{s}_{\mathsf{A}}^{\mathsf{Sta}} := \mathsf{s}_{\mathsf{A}} \circ \phi_{Q,Q_{\mathsf{Act}}} : \mathsf{SeqAdm}_{\mathcal{C}}^{Q} \to \cup_{q \in Q} \Sigma_{\mathsf{A}}(\mathsf{F}(q))$. *Such a strategy is called a* state-strategy. *We denote the set of state-strategies by* $\mathsf{S}_{\mathsf{A}}^{\mathcal{C},\mathsf{Sta}} := \{\mathsf{s}_{\mathsf{A}}^{\mathsf{Sta}} \mid \mathsf{s}_{\mathsf{A}} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}\} \subseteq \mathsf{S}_{\mathsf{A}}^{\mathcal{C},\mathsf{Act}}$. *This is analogous for Player* B.

The stochastic tree induced by action-strategies will be (Q, Q_{Act})-alternating (recall Definition 1.8). Let us define the probability to go from Q to Q_{Act} and vice versa.

Definition 2.14 (Probability transition given two GF-strategies)**.** *Consider a standard concurrent arena* $\mathcal{C}$, *a state* $q \in Q$, *another state* $(q', a, b) \in Q_{\mathsf{Act}}$ *and two* GF*-strategies* $(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \in \Sigma_{\mathsf{A}}(\mathsf{F}(q)) \times \Sigma_{\mathsf{B}}(\mathsf{F}(q))$. *The probability to go from* $q \in Q$ *to* $(q', a, b) \in Q_{\mathsf{Act}}$ *if the players play, in* q, σ_{A} *and* σ_{B}, *denoted* $\mathbb{P}^{\sigma_{\mathsf{A}},\sigma_{\mathsf{B}}}(q, (q', a, b))$, *is equal to:*

$$\mathbb{P}_{\mathcal{C},\mathsf{Act}}^{\sigma_{\mathsf{A}},\sigma_{\mathsf{B}}}(q, (q', a, b)) := \begin{cases} 0 & \text{if } q \neq q' \\ \sigma_{\mathsf{A}}(a) \cdot \sigma_{\mathsf{B}}(b) & \text{otherwise} \end{cases}$$

Furthermore, consider a state $(q, a, b) \in Q_{\mathsf{Act}}$ *and another* $q' \in Q$. *The probability to go from* (q, a, b) *to* q' *is equal to:*

$$\mathbb{P}_{\mathcal{C},\mathsf{Act}}^{\mathcal{D}}((q, a, b), q') := \varrho_q(a, b)(q')$$

We define below the stochastic trees action-induced by a pair of action-strategies.

Definition 2.15 (Probability distribution given two strategies)**.** *Consider a standard concurrent arena* $\mathcal{C}$ *and two arbitrary action-strategies* $(\mathsf{s}_{\mathsf{A}}, \mathsf{s}_{\mathsf{B}}) \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C},\mathsf{Act}} \times \mathsf{S}_{\mathsf{B}}^{\mathcal{C},\mathsf{Act}}$. *We denote by* $\mathbb{P}_{\mathcal{C},\mathsf{Act}}^{\mathsf{s}_{\mathsf{A}},\mathsf{s}_{\mathsf{B}}} : (Q \cup Q_{\mathsf{Act}})^{+} \to \mathcal{D}(Q \cup Q_{\mathsf{Act}})$ *the function giving the probability distribution over the next state of the arena given the sequence of states already seen. That is, for all finite admissible paths* $\pi \in$

$\mathsf{SeqAdm}_\mathcal{C}$ *and* $q \in (Q \cup Q_{\mathsf{Act}})$*, we have:*

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act}}(\pi)(q) := \begin{cases} 0 & \text{if } \pi_{\mathsf{lt}} \in Q \text{ and } q \in Q \\ \mathbb{P}^{\mathsf{s_A}(\pi),\mathsf{s_B}(\pi)}_{\mathcal{C},\mathsf{Act}}(\pi_{\mathsf{lt}}, q) & \text{if } \pi_{\mathsf{lt}} \in Q \text{ and } q \in Q_{\mathsf{Act}} \\ 0 & \text{if } \pi_{\mathsf{lt}} \in Q_{\mathsf{Act}} \text{ and } q \in Q_{\mathsf{Act}} \\ \mathbb{P}^{\mathcal{D}}_{\mathcal{C},\mathsf{Act}}(\pi_{\mathsf{lt}}, q) & \text{if } \pi_{\mathsf{lt}} \in Q_{\mathsf{Act}} \text{ and } q \in Q \end{cases}$$

For $\pi \in (Q \cup Q_{\mathsf{Act}})^+ \setminus \mathsf{SeqAdm}_\mathcal{C}$, $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act}}(\pi) : Q \cup Q_{\mathsf{Act}} \to [0,1]$ *is defined arbitrarily such that* $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act}}(\pi) \in \mathcal{D}(Q \cup Q_{\mathsf{Act}})$.

The stochastic tree $\mathcal{T}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act}}$ action-induced *by the pair of strategies* $(\mathsf{s_A}, \mathsf{s_B})$ *is then equal to* $\mathcal{T}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act}} := \langle Q \cup Q_{\mathsf{Act}}, \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act}} \rangle$.

Observation 2.3. *The stochastic tree* $\mathcal{T}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act}}$ *of the above definition ensures that, for all* $q \in Q$*, and* $\pi \in (Q \cup Q_{\mathsf{Act}})^* \setminus q^{-1} \cdot \mathsf{SeqAdm}_\mathcal{C}$*, we have:*

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q}(\pi) = 0$$

where $q^{-1} \cdot \mathsf{SeqAdm}_\mathcal{C} := \{\rho \in (Q \cup Q_{\mathsf{Act}})^* \mid q \cdot \rho \in \mathsf{SeqAdm}_\mathcal{C}\}$. *In particular, the stochastic tree* $\mathcal{T}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act}}$ *is* (Q, Q_{Act})*-alternating.*

One may wonder, given two state-strategies, how do the stochastic trees induced and action-induced by that pair of strategies relate in terms of the expected value of any measurable functions. In fact, these expected values are equal in both stochastic trees, since state-strategies do not depend on the actions seen, as stated in the lemma below.

Lemma 2.18 (Proof 2.7.4)**.** *Consider a standard concurrent game* $\mathcal{C}$*. Consider a measurable functions* $f : Q^\omega \to [0,1]$*, and any two strategies* $\mathsf{s_A} \in \mathsf{S}^\mathcal{C}_\mathsf{A}$ *and* $\mathsf{s_B} \in \mathsf{S}^\mathcal{C}_\mathsf{B}$*. Then, for all starting states* $q \in Q$*:*

$$\mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[f^q] = \mathbb{E}^{\mathsf{s_A^{Sta}},\mathsf{s_B^{Sta}}}_{\mathcal{C},\mathsf{Act},q}[(f_{Q,Q_{\mathsf{Act}}})^q]$$

Let us now define the value of the game where (action-)strategies are used.

Definition 2.16 ($X_\mathsf{A}, X_\mathsf{B}$-values of the game)**.** *Consider a concurrent game* $\mathcal{G} = \langle \mathcal{C}, f \rangle$ *and let* $g := (f_\mathcal{C})_{Q,Q_{\mathsf{Act}}} : (Q \cup Q_{\mathsf{Act}})^\omega \to [0,1]$*. For* $X_\mathsf{B} \in \{\mathsf{Sta}, \mathsf{Act}\}$ *and a Player-*A *strategy* $\mathsf{s_A} \in \mathsf{S}^{\mathcal{C},\mathsf{Act}}_\mathsf{A}$ *the vector* $\chi_{\mathcal{G},X_\mathsf{B}}[\mathsf{s_A}] : Q \to [0,1]$ *giving the* X_B*-value of the strategy* $\mathsf{s_A}$ *is such that, for all* $q \in Q$*, we have:*

$$\chi_{\mathcal{G},X_\mathsf{B}}[\mathsf{s_A}](q) := \inf_{\mathsf{s_B} \in \mathsf{S}^{\mathcal{C},X_\mathsf{B}}_\mathsf{B}} \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q}[g^q]$$

For $X_\mathsf{A} \in \{\mathsf{Sta}, \mathsf{Act}\}$*, the vector* $\chi_{\mathcal{G},X_\mathsf{A},X_\mathsf{B}}[\mathsf{A}] : Q \to [0,1]$ *giving the* $X_\mathsf{A}, X_\mathsf{B}$*-value for Player* A *is such that, for all* $q \in Q$*, we have:*

$$\chi_{\mathcal{G},X_\mathsf{A},X_\mathsf{B}}[\mathsf{A}](q) := \sup_{\mathsf{s_A} \in \mathsf{S}^{\mathcal{C},X_\mathsf{A}}_\mathsf{A}} \chi_{\mathcal{G},X_\mathsf{B}}[\mathsf{s_A}](q)$$

That is, Player A *uses* X_{A}*-strategies and Player* B *uses* X_{B}*-strategies. The value can be defined symmetrically for Player* B*. When the* $X_{\mathsf{A}}, X_{\mathsf{B}}$*-values of both players are the same, this defines the* $X_{\mathsf{A}}, X_{\mathsf{B}}$*-value of the game:* $\chi_{\mathcal{G},X_{\mathsf{A}},X_{\mathsf{B}}} := \chi_{\mathcal{G},X_{\mathsf{A}},X_{\mathsf{B}}}[\mathsf{A}] = \chi_{\mathcal{G},X_{\mathsf{A}},X_{\mathsf{B}}}[\mathsf{B}]$*. (If* $X_{\mathsf{A}} = X_{\mathsf{B}}$*, one of them is omitted.)*

2.5.2 . Expressive power

The existence of Act-value is ensured as soon as, at each state, one of the players set of actions is finite. In addition, when, at each state, both of the players set of actions are finite and if the game is win/lose, then the Act-value and the Sta-value are equal. This is the main result of this section stated below and it is a corollary of Theorem 2.3.

Corollary 2.19. *Consider a standard concurrent game* $\mathcal{G}$*. Assume that, for all* $q \in Q$*, we have either* $\mathsf{Act}_{\mathsf{A}}^q$ *or* $\mathsf{Act}_{\mathsf{B}}^q$ *finite (in particular, the game* $\mathcal{G}$ *is valuable by Proposition 1.12). In that case, the game* $\mathcal{G}$ *has an* Act*-value.*

If we additionally assume that $\mathcal{G}$ *is win/lose, and that, for all* $q \in Q$*, both* $\mathsf{Act}_{\mathsf{A}}^q$ *and* $\mathsf{Act}_{\mathsf{B}}^q$ *are finite, then the* Act*-value and the* Sta*-value of* $\mathcal{G}$ *exist and are equal.*

Note that the proof of this corollary is quite long. However, there is no real difficulty to deduce it from Theorem 2.3. However, the change of formalism that we consider in this section makes the proof technical.

To prove this corollary, we define a new standard concurrent game where the actions chosen by the players are encoded in its states. We define below such a game along with a way to translate strategies from the original game to this new game.

Definition 2.17 (Action-encoded-states game)**.** *Consider a standard concurrent game* $\mathcal{G}$*. We build the game* $\mathcal{G}_{\mathsf{Act}} = \langle \mathcal{C}_{\mathsf{Act}}, f_{\mathsf{Act}} \rangle$ *in the following way, for* $\mathcal{C}_{\mathsf{Act}} := \langle Q \cup Q_{\mathsf{Act}}, , \mathsf{F}_{\mathsf{Act}}, \mathsf{K}_{\mathsf{Act}}, \mathsf{col}_{\mathsf{Act}} \rangle$*:*

- *for all stopping states* $q \in Q_{\mathsf{s}}$*,* q *is still a stopping state in* $\mathcal{G}_{\mathsf{Act}}$ *with the same value;*
- *for all non-stopping states* $q \in Q_{\mathsf{ns}}$*, we let* $\mathsf{F}_{\mathsf{Act}}(q) := \langle \mathsf{Act}_{\mathsf{A}}^q, \mathsf{Act}_{\mathsf{B}}^q, Q_{\mathsf{Act}}, \varrho_q^{\mathsf{Act}} \rangle$ *such that, for all* $(a, b) \in \mathsf{Act}_{\mathsf{A}}^q \times \mathsf{Act}_{\mathsf{B}}^q$*,* $\varrho_q^{\mathsf{Act}}(a, b)(q, a, b) := 1$*;*
- *all states* $(q, a, b) \in Q_{\mathsf{Act}}$ *are trivial, specifically, we have* $\mathsf{F}_{\mathsf{Act}}((q, a, b)) := \langle *, *, Q, \varrho_{(q,a,b)}^{\mathsf{Act}} \rangle$ *with* $\varrho_{(q,a,b)}^{\mathsf{Act}}(*, *) := \varrho_q(a, b) \in \mathcal{D}(Q)$*;*
- $\mathsf{K}_{\mathsf{Act}} := Q$*;*
- *for all* $q \in Q_{\mathsf{ns}}$*, we let* $\mathsf{col}_{\mathsf{Act}}(q) := q$ *and for all* $(q, a, b) \in Q_{\mathsf{Act}}$*, we let* $\mathsf{col}_{\mathsf{Act}}((q, a, b)) := q$*;*
- *We let* $Q^{\mathsf{db}} := \cup_{q \in Q} \{q\} \cdot \{q\}$ *and* $f^{\mathsf{sg}} : Q^{\mathsf{db}} \to Q$ *be the canonical function from* Q^{db} *to* Q*. Then, we let* $f_{\mathsf{Act}} : Q^\omega \to [0, 1]$ *be such that,*

for all $\rho \in Q^\omega$*:*

$$f_{\mathsf{Act}}(\rho) := \begin{cases} 0 & \text{if } \rho \notin (Q^{\mathsf{db}})^\omega \\ f_{\mathcal{C}} \circ (\mathsf{f}^{\mathsf{sg}})^\omega(\rho) & \text{otherwise} \end{cases}$$

We define below how to translate strategies from $\mathcal{C}$ to $\mathcal{C}_{\mathsf{Act}}$.

Definition 2.18 (Translating strategies). *Consider a standard concurrent arena* $\mathcal{C}$*. Consider a Player* $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$*. We define* $\mathsf{f}^{\mathsf{C}} : \mathsf{S}_{\mathsf{C}}^{\mathcal{C},\mathsf{Act}} \to \mathsf{S}_{\mathsf{C}}^{\mathcal{C}_{\mathsf{Act}}}$ *in the following way. For all* $\mathsf{s}_{\mathsf{C}} \in \mathsf{S}_{\mathsf{C}}^{\mathcal{C},\mathsf{Act}}$*, we have for all* $\pi \in (Q \cup Q_{\mathsf{Act}})^+$*:*

$$\mathsf{f}^{\mathsf{C}}(\mathsf{s}_{\mathsf{C}})(\pi) := \begin{cases} \text{is arbitrary} & \text{if } \pi \notin \mathsf{SeqAdm}_{\mathcal{C}} \\ \mathsf{s}_{\mathsf{C}}(\pi) & \text{otherwise} \end{cases}$$

Furthermore, we define $\mathsf{g}^{\mathsf{C}} : \mathsf{S}_{\mathsf{C}}^{\mathcal{C}_{\mathsf{Act}}} \to \mathsf{S}_{\mathsf{C}}^{\mathcal{C},\mathsf{Act}}$ *in the following way, for all* $\pi \in \mathsf{SeqAdm}_{\mathcal{C}} \subseteq (Q \cup Q_{\mathsf{Act}})^+$*:*

$$\mathsf{g}^{\mathsf{C}}(\mathsf{s}_{\mathsf{C}})(\pi) := \mathsf{s}_{\mathsf{C}}(\pi)$$

Lemma 2.20 (Proof 2.7.5). *Consider a standard concurrent arena* $\mathcal{C}$*. Let* $\mathsf{s}_{\mathsf{A}}^{\mathsf{Sta}} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}_{\mathsf{Act}}}$ *be a strategy in* $\mathcal{C}_{\mathsf{Act}}$ *and let* $\mathsf{s}_{\mathsf{B}}^{\mathsf{Act}} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C},\mathsf{Act}}$ *be an action-strategy in* $\mathcal{C}$*. Let* $\mathsf{s}_{\mathsf{A}}^{\mathsf{Act}} := \mathsf{g}^{\mathsf{A}}(\mathsf{s}_{\mathsf{A}}) \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C},\mathsf{Act}}$ *and let* $\mathsf{s}_{\mathsf{B}} := \mathsf{f}^{\mathsf{B}}(\mathsf{s}_{\mathsf{B}}^{\mathsf{Act}}) \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}_{\mathsf{Act}}}$*.*

Then, for all $q \in Q$*:*

$$\mathbb{E}_{\mathcal{C},\mathsf{Act},q}^{\mathsf{s}_{\mathsf{A}}^{\mathsf{Act}},\mathsf{s}_{\mathsf{B}}^{\mathsf{Act}}}[((f_{\mathcal{C}})_{Q,Q_{\mathsf{Act}}})^q] = \mathbb{E}_{\mathcal{C}_{\mathsf{Act}},q}^{\mathsf{s}_{\mathsf{A}},\mathsf{s}_{\mathsf{B}}}[((f_{\mathsf{Act}})_{\mathcal{C}_{\mathsf{Act}}})^q]$$

Note that this also holds if we reverse the roles of Player A *and Player* B *strategies.*

We can now proceed to the proof of Corollary 2.19.

Proof. First assume that for all $q \in Q$, we have either $\mathsf{Act}_{\mathsf{A}}^q$ or $\mathsf{Act}_{\mathsf{B}}^q$ finite. In that case, the standard game $\mathcal{G}_{\mathsf{Act}}$ from Definition 2.17 is valuable. Indeed, for all $q \in Q_{\mathsf{ns}} \cup Q_{\mathsf{Act}}$, the game form $\mathsf{F}_{\mathsf{Act}}(q)$ has finitely many actions for at least one player, hence it is valuable by Proposition 1.12.

In the game $\mathcal{G}_{\mathsf{Act}}$, we only consider usual strategies, that is the ones that we have considered in this dissertation but in this section. Hence, we consider the usual value of the game and of the strategies. Consider some state $q \in Q$. Let us show that $\chi_{\mathcal{G},\mathsf{Act}}[\mathsf{A}](q) \geq \chi_{\mathcal{G}_{\mathsf{Act}}}[\mathsf{A}](q)$. Let $\varepsilon > 0$. Consider a Player-A strategy $\mathsf{s}_{\mathsf{A}} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}_{\mathsf{Act}}}$ such that $\chi_{\mathcal{G}_{\mathsf{Act}}}[\mathsf{s}_{\mathsf{A}}](q) \geq \chi_{\mathcal{G}_{\mathsf{Act}}}(q) - \varepsilon$. Let $\mathsf{s}_{\mathsf{A}}^{\mathsf{Act}} := \mathsf{g}^{\mathsf{A}}(\mathsf{s}_{\mathsf{A}}) \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C},\mathsf{Act}}$ be an action-strategy in the game $\mathcal{G}$. Let us show that the strategy $\mathsf{s}_{\mathsf{A}}^{\mathsf{Act}}$ has an Act-value at least $\chi_{\mathcal{G}_{\mathsf{Act}}}(q) - \varepsilon$ in the game $\mathcal{G}$. Consider any Player-B action-strategy $\mathsf{s}_{\mathsf{B}}^{\mathsf{Act}} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C},\mathsf{Act}}$ in the game $\mathcal{G}$. Then, letting $\mathsf{s}_{\mathsf{B}} := \mathsf{f}^{\mathsf{B}}(\mathsf{s}_{\mathsf{B}}^{\mathsf{Act}}) \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}_{\mathsf{Act}}}$ be a strategy in the game $\mathcal{G}_{\mathsf{Act}}$, we have by Lemma 2.20:

$$\mathbb{E}_{\mathcal{C},\mathsf{Act},q}^{\mathsf{s}_{\mathsf{A}}^{\mathsf{Act}},\mathsf{s}_{\mathsf{B}}^{\mathsf{Act}}}[((f_{\mathcal{C}})_{Q,Q_{\mathsf{Act}}})^q] = \mathbb{E}_{\mathcal{C}_{\mathsf{Act}},q}^{\mathsf{s}_{\mathsf{A}},\mathsf{s}_{\mathsf{B}}}[((f_{\mathsf{Act}})_{\mathcal{C}_{\mathsf{Act}}})^q] \geq \chi_{\mathcal{G}_{\mathsf{Act}}}[\mathsf{s}_{\mathsf{A}}](q) \geq \chi_{\mathcal{G}_{\mathsf{Act}}}(q) - \varepsilon$$

As this holds for all Player-B action-strategies $\mathsf{s}_\mathsf{B}^{\mathsf{Act}} \in \mathsf{S}_\mathsf{B}^{\mathcal{C},\mathsf{Act}}$, it follows that:

$$\chi_{\mathcal{G},\mathsf{Act}}[\mathsf{s}_\mathsf{A}^{\mathsf{Act}}](q) \geq \chi_{\mathcal{G}_{\mathsf{Act}}}(q) - \varepsilon$$

and therefore

$$\chi_{\mathcal{G},\mathsf{Act}}[\mathsf{A}](q) \geq \chi_{\mathcal{G}_{\mathsf{Act}}}(q) - \varepsilon$$

As this holds for all $\varepsilon > 0$, it follows that:

$$\chi_{\mathcal{G},\mathsf{Act}}[\mathsf{A}](q) \geq \chi_{\mathcal{G}_{\mathsf{Act}}}(q)$$

Symmetrically, we also obtain that $\chi_{\mathcal{G},\mathsf{Act}}[\mathsf{B}](q) \leq \chi_{\mathcal{G}_{\mathsf{Act}}}(q)$. Hence, the game $\mathcal{G}$ has an Act-value at state $q \in Q$ equal to $\chi_{\mathcal{G},\mathsf{Act}}(q) = \chi_{\mathcal{G}_{\mathsf{Act}}}(q)$ which holds for all $q \in Q$.

Assume now that, for all $q \in Q$, we have $\mathsf{Act}_\mathsf{A}^q$ and $\mathsf{Act}_\mathsf{B}^q$ finite and that the objective f is win/lose. It follows straightforwardly that f_{Act} is also win/lose. Consider now some $q \in \mathsf{K}_\mathsf{A} = Q$. Then, we have $(\mathsf{col}_{\mathsf{Act}})^{-1}[\{q\}] = \{q\} \cup \{(q,a,b) \mid (a,b) \in \mathsf{Act}_\mathsf{A}^q \times \mathsf{Act}_\mathsf{B}^q\}$ finite. Hence, the uniformizing pair $(\mathsf{K}_{\mathsf{Act}} = Q, \mathsf{col}_{\mathsf{Act}})$ has finite representatives in $Q \cup Q_{\mathsf{Act}}$. Now, consider some $\varepsilon > 0$. By Theorem 2.3, there exists a Player-A strategy $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}_{\mathsf{Act}}}$ that is $(\mathsf{K}_{\mathsf{Act}}, \mathsf{col}_{\mathsf{Act}})$-uniform such that $\chi_{\mathcal{G}_{\mathsf{Act}}}[\mathsf{s}_\mathsf{A}](q) \geq \chi_{\mathcal{G}_{\mathsf{Act}}}(q) - \varepsilon = \chi_{\mathcal{G},\mathsf{Act}}(q) - \varepsilon$. Let $\mathsf{s}_\mathsf{A}^{\mathsf{Act}} := \mathsf{g}^\mathsf{A}(\mathsf{s}_\mathsf{A}) \in \mathsf{S}_\mathsf{A}^{\mathcal{C},\mathsf{Act}}$ be an action-strategy in the game $\mathcal{G}$. We have shown above that this action-strategy $\mathsf{s}_\mathsf{A}^{\mathsf{Act}}$ is ε-optimal in the game $\mathcal{G}$. Let us show that this strategy s_A actually is a state-strategy. Since the strategy s_A is $(\mathsf{K}_{\mathsf{Act}}, \mathsf{col}_{\mathsf{Act}})$-uniform, it can be seen as a map $\mathsf{s}'_\mathsf{A} : \mathsf{K}_{\mathsf{Act}}^* \cdot (Q \cup Q_{\mathsf{Act}}) \to \cup_{q \in Q \cup Q_{\mathsf{Act}}} \Sigma_\mathsf{A}(\mathsf{F}_{\mathsf{Act}}(q))$. Furthermore, $\mathsf{K}_{\mathsf{Act}} = Q$. Now, we let $\mathsf{t}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}}$ be a Player-A strategy in the game $\mathcal{G}$ such that, for all $\pi \in Q^+$, we have $\mathsf{t}_\mathsf{A}(\pi) := \mathsf{s}_\mathsf{A}(\pi)$. Let us show that $\mathsf{s}_\mathsf{A}^{\mathsf{Act}} = \mathsf{t}_\mathsf{A}^{\mathsf{Sta}}$. Consider some $\pi \in \mathsf{SeqAdm}_{\mathcal{C}}^Q$. We have:

$$\mathsf{t}_\mathsf{A}^{\mathsf{Sta}}(\pi) = \mathsf{t}_\mathsf{A} \circ \phi_{Q,Q_{\mathsf{Act}}}(\pi) = \mathsf{s}'_\mathsf{A} \circ \phi_{Q,Q_{\mathsf{Act}}}(\pi) = \mathsf{s}_\mathsf{A}(\pi) = \mathsf{g}^\mathsf{A}(\mathsf{s}_\mathsf{A})(\pi) = \mathsf{s}_\mathsf{A}^{\mathsf{Act}}(\pi)$$

As this holds for all $\pi \in \mathsf{SeqAdm}_{\mathcal{C}}^Q$, it follows that $\mathsf{s}_\mathsf{A}^{\mathsf{Act}} = \mathsf{t}_\mathsf{A}^{\mathsf{Sta}}$. We have exhibited a Player-A strategy $\mathsf{s}_\mathsf{A}^{\mathsf{Act}}$ in the game $\mathcal{G}$ whose value ε-close to the Act-value of the game and that is a state-strategy. This holds for all $\varepsilon > 0$ and also for Player B. Hence, by Lemma 2.18, the Act-value and the Sta-value of $\mathcal{G}$ exist and are equal. □

Note that Lemma 2.18 ensures that if a Sta-value exists in $\mathcal{G}$ and if a value (as we have considered until this section) exist in $\mathcal{G}$, then they are equal.

We conclude this section by providing an example of a standard game with finitely many actions for both players at each state but where finding optimal strategies requires to consider action-strategies. This shows that although at the limit, knowing the actions does not improve what the strategies can do, it may be that achieving a specific value is only possible if the strategy knows the action played (by the other player in our example).

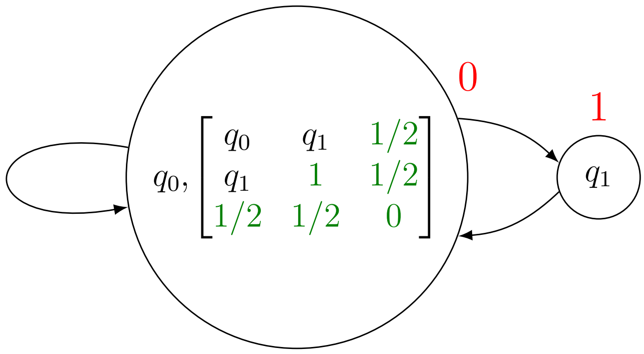

Figure 2.5: A co-Büchi win/lose game where Player-A optimal strategies can only be found among action-strategies. The colors are depicted in red near the states.

Definition 2.19 (Game described in Figure 2.5)**.** *The game $\mathcal{G} = \langle \mathcal{C}, W \rangle$ of Figure 2.5 is standard, deterministic and win/lose. They are only two non-stopping states: q_0 whose local interaction is not trivial and q_1 whose only successor state is q_0. The stopping states (recall Definition 1.18) are not drawn but are referred to by their values in the local interaction at state q_0. The set of colors considered is $\mathsf{K} := \{0, 1\}$ and the colors of the states q_0 and q_1 are given in red near them: $\mathsf{col}(q_0) := 0$ and $\mathsf{col}(q_1) := 1$. This game is win/lose, and the objective W is a co-Büchi objective (recall Definition 1.25): if no stopping state is reached, Player A wins if and only if the state q_1 is seen only finitely often. The Player-A set of actions at state q_0 is $\mathsf{Act}_{\mathsf{A}}^{q_0} := \{a_1, a_2, a_3\}$ where a_1 refers to the top row and a_3 refers to the bottom row and similarly we have $\mathsf{Act}_{\mathsf{B}}^{q_0} := \{b_1, b_2, b_3\}$ where b_1 refers to the leftmost column and b_3 refers to the rightmost column.*

We have presented a slight modification of the game described in Definition 2.19 above in [40] to illustrate another property ensured by concurrent games. We will discuss further this example in Chapters 3 and 5.

Proposition 2.21. *The co-Büchi standard finite deterministic game $\mathcal{G}$ of Figure 2.5 is such that:*

- *the game has value $\frac{1}{2}$ and Player B has an optimal positional strategy;*
- *Player A has an optimal action-strategy but has no optimal state-strategy.*

We decompose this proposition in three lemmas.

Lemma 2.22. *The value of the game described in Definition 2.19 is $1/2$. Furthermore, for all positive $\varepsilon > 0$, the Player-A positional strategy $\mathsf{s}_{\mathsf{A}}^{\varepsilon} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ such that $\mathsf{s}_{\mathsf{A}}^{\varepsilon}(q_0)(a_1) := 1 - \varepsilon$ and $\mathsf{s}_{\mathsf{A}}^{\varepsilon}(q_0)(a_3) := \varepsilon$ has value $\frac{1}{2} - \varepsilon$.*

Proof. Consider a Player-B strategy $\mathsf{s_B}$ playing positionally b_3 with probability 1. Then, regardless of Player-A's strategy, the game will stop after one step and a stopping state value of at most $\frac{1}{2}$ will be reached. Hence, $\chi_{\mathcal{G}}(\mathsf{s_B})(q_0) \leq \frac{1}{2}$. Now, consider some $\varepsilon > 0$ and the Player-A strategy $\mathsf{s}_\mathsf{A}^\varepsilon$ described in the statement. Then, regardless of Player-B's strategy, each time the game is at state q_0, a stopping state is reached in the next step with probability at least $\varepsilon > 0$. Hence, almost-surely a stopping state is reached. Furthermore, the expected value of the stopping states reached is at least $\frac{1}{2} - \varepsilon$. Hence, this Player-$\mathsf{A}$ strategy $\mathsf{s}_\mathsf{A}^\varepsilon$ has value at least $\frac{1}{2} - \varepsilon$. In fact, its value is exactly equal to $\frac{1}{2} - \varepsilon$ since Player B can play action b_3 with probability 1 and ensure that a stopping state is reached and that their expected value is equal to $\frac{1}{2} - \varepsilon$. That is, $\chi_{\mathcal{G}}(\mathsf{s}_\mathsf{A}^\varepsilon)(q_0) = \frac{1}{2} - \varepsilon$. As this holds for all $\varepsilon > 0$, it follows that $\chi_{\mathcal{G}}(q_0) = \frac{1}{2}$. □

Let us now describe informally an optimal Player-A action-strategy $\mathsf{s_A}$. First, note that it needs not be defined after b_3 is seen since in that case a stopping state is necessarily reached. While b_2 has not occurred, $\mathsf{s_A}$ plays with positive probability a_1 and a_2 and a_1 with very high and increasing probability such that, if action b_2 or b_3 never occurs, then almost-surely, the state q_1 is seen only finitely often. As soon as the action b_2 occurs, there is a positive probability to reach the stopping state of value 1. (This probability may be arbitrarily small if Player B waits long enough, but it is positive.) In that case, the strategy $\mathsf{s_A}$ switches to a strategy $\mathsf{s}_\mathsf{A}^\varepsilon$ for some small enough ε.[8] This $\varepsilon > 0$ is chosen so that the mean of the probability to reach the stopping state of value 1 and $\frac{1}{2} - \varepsilon$ is at least $\frac{1}{2}$. The formal arguments we give below on how to construct an optimal strategy have already been given in [40].

Lemma 2.23. *Consider the game described in Definition 2.19. We let* $\varphi_{\mathsf{Act}} : \mathsf{SeqAdm}_{\mathcal{C}} \to (\mathsf{Act}_\mathsf{B}^{q_0})^*$ *be such that, for all* $\rho = q \cdot (q, a, b) \cdot q' \cdot (q', a', b') \cdots \in \mathsf{SeqAdm}_{\mathcal{C}}$, *we have* $\varphi_{\mathsf{Act}}(\rho) := b \cdot b' \cdots \in (\mathsf{Act}_\mathsf{B}^{q_0})^*$. *Consider the Player-$\mathsf{A}$ action-strategy* $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C},\mathsf{Act}}$ *such that, for all* $\rho \in \mathsf{SeqAdm}_{\mathcal{C}}$ *such that* $\rho_{\mathsf{lt}} = q_0$, *denoting* $\pi := \varphi_{\mathsf{Act}}(\rho)$, *we have:*

$$\mathsf{s_A}(\rho) := \begin{cases} \{a_1 \mapsto 1 - \varepsilon_{|\pi|}, a_2 \mapsto \varepsilon_{|\pi|}\} & \text{if } \pi \in (b_1)^* \\ \mathsf{s}_\mathsf{A}^{\varepsilon'_n}(\rho) & \text{otherwise, for } n := \max\{k \in \mathbb{N} \mid (b_1)^k \sqsubset \pi\} \end{cases}$$

where, for all $n \in \mathbb{N}$, *we have* $\varepsilon_n := \frac{1}{2^{n+1}}$ *and* ε'_n *is chosen such that:*

$$(1 - \varepsilon_n) \cdot (\frac{1}{2} - \varepsilon'_n) + \varepsilon_n \geq \frac{1}{2}$$

For instance, $\varepsilon'_n := \frac{1}{2} - \frac{\frac{1}{2} - \varepsilon_n}{1 - \varepsilon_n} = \frac{\varepsilon_n}{2(1-\varepsilon_n)} > 0$. *This Player-$\mathsf{A}$ action-strategy* $\mathsf{s_A}$ *is optimal from* q_0, *i.e. it has value* $\frac{1}{2}$.

[8]Note that switching strategy is necessary: if we remove the action a_3 for Player A, then the value of the game is 0 from q_0.

Proof. Consider any Player-B strategy $\mathsf{s_B}$. We let $Q_a := Q \cup Q_{\mathsf{Act}}$. For all $k \in \mathbb{N}$, we denote by $\mathsf{R}_k := (Q_a)^k \cdot (Q_a)^* \cdot q_1 \cdot (Q_a)^\omega$ the event describing the infinite paths for which the state q_1 is seen after at least k steps. With some abuse of notations, all sequences of actions in $(\mathsf{Act}_{\mathsf{B}}^{q_0})^\uparrow$ are seen as events in $\mathsf{Borel}(Q_a)$ where we consider only the Player-B actions seen, similarly to what we did in the lemma with the function φ_{Act}.

First, consider what happens assuming that only the action b_1 is played. For all $k \in \mathbb{N}$, we have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[\mathsf{R}_k \cap (b_1)^\omega] \leq \sum_{n \geq k} \varepsilon_n = \frac{1}{2^k}$$

Hence:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[(\bigcap_{k \in \mathbb{N}} \mathsf{R}_k) \cap (b_1)^\omega] \leq \lim_{k \to \infty} \sum_{n \geq k} \varepsilon_n = 0$$

Furthermore, we have $\mathsf{coBuchi}_{Q,Q_{\mathsf{Act}}} = (Q_a)^\omega \setminus (\bigcap_{k \in \mathbb{N}} \mathsf{R}_k)$. It follows that:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[\mathsf{coBuchi}_{Q,Q_{\mathsf{Act}}} \cap (b_1)^\omega] = \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[(b_1)^\omega]$$

On the other hand, let us consider what happens if at some point the action b_2 occurs. Consider some $k \in \mathbb{N}$ such that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[(b_1)^k \cdot b_2] > 0$. Assuming the event $(b_1)^k \cdot b_2$, we have that with probability $1 - \varepsilon_k$ the game proceeds to q_1 and the Player-A strategy switches to a strategy of value $\frac{1}{2} - \varepsilon'_k$ (by Lemma 2.22) and with probability ε_k, a stopping state of value 1 is reached. Hence, we have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[\mathsf{coBuchi}_{Q,Q_{\mathsf{Act}}} \mid (b_1)^k \cdot b_2] = (1 - \varepsilon_k) \cdot (\frac{1}{2} - \varepsilon'_k) + \varepsilon_k = \frac{1}{2}$$

As this holds for all $k \in \mathbb{N}$, it follows that:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[\mathsf{coBuchi}_{Q,Q_{\mathsf{Act}}} \cap (b_1)^* \cdot b_2] \geq \frac{1}{2} \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[(b_1)^* \cdot b_2]$$

Finally, considering the case where action b_3 occurs after a sequence of actions b_1, we have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[\mathsf{coBuchi}_{Q,Q_{\mathsf{Act}}} \cap (b_1)^* \cdot b_3] = \frac{1}{2} \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[(b_1)^* \cdot b_3]$$

Overall, we obtain:

$$\begin{aligned}
\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[\mathsf{coBuchi}_{Q,Q_{\mathsf{Act}}}] &= \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[\mathsf{coBuchi}_{Q,Q_{\mathsf{Act}}} \cap (b_1)^\omega] \\
&+ \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[\mathsf{coBuchi}_{Q,Q_{\mathsf{Act}}} \cap (b_1)^* \cdot b_2] \\
&+ \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[\mathsf{coBuchi}_{Q,Q_{\mathsf{Act}}} \cap (b_1)^* \cdot b_3] \\
&\geq \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[(b_1)^\omega] + \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[(b_1)^* \cdot b_2] \cdot \frac{1}{2} \\
&+ \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\mathsf{Act},q_0}[(b_1)^* \cdot b_3] \cdot \frac{1}{2} \\
&\geq \frac{1}{2}
\end{aligned}$$

Hence, the Player-A strategy $\mathsf{s_A}$ is optimal. □

Let us now proceed to the proof of Proposition 2.21.

Proof. Given what we have proved in Lemma 2.22 and 2.23, it remains to show that no Player-A strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ (that is, a type of strategy that we have considered so far in this dissertation) can be optimal in this game. We let $Q := \{q_0, q_1\}$ and $W := \mathsf{coBuchi}_{Q,Q_{\mathsf{Act}}}$.

The strategy $\mathsf{s_A}$ can be seen as a function $\mathsf{s}'_\mathsf{A} : Q^* \cdot q_0 \to \Sigma_\mathsf{A}(\mathsf{F}(q_0))$. Let us build a Player-B strategy $\mathsf{s_B}$ such that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[\mathsf{coBuchi}] < \frac{1}{2}$. Let $\mathsf{Next}_{q_1} := \{\rho \in Q^* \cdot q_0 \mid \mathsf{s_A}(\rho)(a_1) = 1 \vee \mathsf{s_A}(\rho)(a_2) = 1\}$. We also let $\mathsf{Adm} := \{\rho \in Q^* \mid \forall i < |\rho| - 1,\ \rho_{\leq i} \in \mathsf{Next}_{q_1} \Rightarrow \rho_{\leq i+1} = \rho_{\leq i} \cdot q_1\}$. Finally, we also let $\mathsf{Err} := \{\rho \in Q^* \cdot q_0 \mid \mathsf{s_A}(q_0)(a_3) > 0\}$. We can now define the strategy $\mathsf{s_B}$. First, for all $\rho \in \mathsf{Next}_{q_1}$, we let $\mathsf{Resp_B}(\rho) \in \mathsf{Act}_\mathsf{B}^{q_0}$ to be equal to $\mathsf{Resp_B} := b_2$ if $\mathsf{s_A}(q_0)(a_1) = 1$ and $\mathsf{Resp_B} := b_1$ otherwise. Then, there are two cases:

- Assume that $\mathsf{Adm} \cap \mathsf{Err} \neq \emptyset$. In that case, consider some $\pi \in \mathsf{Adm} \cap \mathsf{Err}$ with no prefix in $\mathsf{Adm} \cap \mathsf{Err}$. We let $n := \pi$ and $\mathsf{NoPr}(\pi) := \{\rho \in Q^* \cdot q_0 \mid \rho \not\sqsubseteq \pi\}$. We define $\mathsf{s_B}$ in the following way, for all $\rho \neq \pi \in Q^* \cdot q_0$:

$$\mathsf{s_B}(\rho) := \begin{cases} \{b_3 \mapsto 1\} & \text{if } \rho \in \mathsf{NoPr}(\pi) \\ \{\mathsf{Resp_B}(\rho) \mapsto 1\} & \text{otherwise, if } \rho \sqsubset \pi, \text{ and } \rho \in \mathsf{Next}_{q_1} \\ \{q_0 \mapsto 1\} & \text{otherwise, if } \rho \sqsubset \pi, \text{ and } \rho \notin \mathsf{Next}_{q_1} \end{cases}$$

  By definition of Adm and Err and of π, for all $\rho \in Q^* \cdot q_0$ such that $\rho \sqsubset \pi$, we have:

  - if $\rho \in \mathsf{Next}_{q_1}$, then $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}}(\rho)(q_1) = 1$ and $\rho \cdot q_1 \sqsubseteq \pi$;
  - if $\rho \notin \mathsf{Next}_{q_1}$, then $\mathsf{s_B}(\rho)(b_0) = 1$ and $\mathsf{s_A}(\pi)(a_3) = 0$ and $\mathsf{s_A}(a_1), \mathsf{s_A}(a_2) > 0$. Hence, $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}}(\rho)(\pi_{|\rho|}) > 0$.

  Therefore, $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}(\pi) > 0$. We let $p_\pi := \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}(\pi)$. Since, $\pi \in \mathsf{Err}$, we have $\mathsf{s_A}(\pi)(a_3) > 0$. Since $\mathsf{s_B}(\pi)(a_3) = 1$, it follows that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}(\rho)[W \mid \pi] < \frac{1}{2}$. We obtain:

$$\begin{aligned} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[W] &= \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[W \cap \pi] + \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[W \cap \mathsf{NoPr}(\pi)] \\ &= \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[W \mid \pi] \cdot p_\pi + \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[W \mid \mathsf{NoPr}(\pi)] \cdot (1 - p_\pi) \\ &< \frac{1}{2} \cdot p_\pi + \frac{1}{2} \cdot (1 - p_\pi) = \frac{1}{2} \end{aligned}$$

- Assume now that $\mathsf{Adm} \cap \mathsf{Err} = \emptyset$. Consider a sequence $(\varepsilon_k)_{k \in \mathbb{N}}$ such that $\sum_k \varepsilon_k < \frac{1}{2}$. Then, for all $\rho \in Q^* \cdot q_0$, we let $|\rho|_1$ denote the number of q_1 in ρ. Then, for all $\rho \in Q^* \cdot q_0$, we let:

$$\mathsf{s_B}(\rho) := \begin{cases} \{\mathsf{Resp_B}(\rho) \mapsto 1\} & \text{if } \rho \in \mathsf{Next}_{q_1} \\ \{b_1 \mapsto 1 - \varepsilon_k; b_2 \mapsto \varepsilon_k\} & \text{otherwise, for } k := |\rho|_1 \end{cases}$$

Clearly, for all $\rho \notin \mathsf{Adm}$, we have $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}(\rho) = 0$. Hence, $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[Q^* \cdot (0 \cup \frac{1}{2})] = 0$ (where 0 and $\frac{1}{2}$ refer to the stopping states of the same value), since $\mathsf{Adm} \cap \mathsf{Err} = \emptyset$. Furthermore (1 referring to the stopping state of value 1):

$$\begin{aligned}\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[Q^* \cdot 1] &\leq \sum_{k \in \mathbb{N}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[\{\rho \in Q^* \mid |\rho|_1 = k\} \cdot 1] \\ &\leq \sum_{k \in \mathbb{N}} \varepsilon_k < \frac{1}{2}\end{aligned}$$

In addition:

$$\begin{aligned}\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[Q^* \cdot (q_0)^\omega] &\leq \sum_{k \in \mathbb{N}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[Q^k \cdot (q_0)^\omega] \\ &\leq \sum_{k \in \mathbb{N}} \lim_{n \to \infty} (1 - \varepsilon_k)^n = 0\end{aligned}$$

Hence:

$$\begin{aligned}\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[W] = \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[Q^* \cdot 1] + \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[W \cap (Q^* \cdot q_1)^\omega] \\ \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[Q^* \cdot 1] < \frac{1}{2}\end{aligned}$$

In any case, the Player-A strategy $\mathsf{s_A}$ is not optimal from q_0. □

2.6 Discussion and open question

This chapter is mainly devoted to the proof of Theorem 2.3, which is a generalization of Martin's result [12]. As we discussed in this chapter, the way we prove this generalization uses Martin's central idea: building, from a concurrent game $\mathcal{G}$, a turn-based game $\mathcal{G}_{\mathsf{tb}}$. Then, from winning strategies in $\mathcal{G}_{\mathsf{tb}}$, one can design almost-optimal strategies in $\mathcal{G}$. By closely examining how these almost-optimal strategies are obtained, we were able to establish extensions (1.a) and (1.b) of Theorem 2.3. Quite frustratingly, though the conclusion of result (1.b) seems unsurprising, we need two additional assumptions to establish it. This leaves as open question if, by further exploiting the properties ensured by the game $\mathcal{G}_{\mathsf{tb}}$, we could prove that this result (1.b) still holds even if one or two of these additional assumptions are weakened, or even dropped.

Open Question 2.1. *Does result* (1.b) *of Theorem 2.3 still holds if we do not assume anymore that* $(\mathsf{K}, \mathsf{col})$ *has a finite representative in* Q *and/or that* $\mathcal{G}$ *is win/lose?*

We believe that another benefit of this chapter, besides the results established in Theorem 2.3, is how the proof of this theorem, which generally speaking is not new, is presented. One of our goal was to explain Martin's underlying ideas with intermediate lemmas and examples. However, because we wanted to establish Theorem 2.3 in all its generality, the intermediate lemmas we introduced and the turn-based games indexed by guards we defined are heavy on notations and assumptions. This is particularly noticeable when considering Lemmas 2.11 and 2.12: because we prove result (1.a), we consider a collection of GF-strategies supremizing the game $\mathcal{G}$; because we prove result (1.b), we consider a uniformizing pair; and because we prove that both results (1.a) and (1.b) hold even if the local interactions in $\mathcal{G}$ are not valuable, we consider guards. This last constraint is the most apparent when considering the turn-based games indexed by guards that we consider, instead of only considering one turn-based game, as Martin did. Hence, it could be relevant to give a detailed proof only of result (2) from Theorem 2.3. That way, we could detail the underlying ideas behind the proof, without the burden of, what would then be, unnecessary complications.

2.7 Appendix

2.7.1 . Proof of Proposition 2.4

Proof. First, for all $j \in \mathbb{N}$, $\bowtie \in \{\leq, <, \geq, >, =, \neq\}$ and $u \in \mathbb{R}$, we denote by $V(j, \bowtie, u)$ the measurable set:

$$V(j, \bowtie, u) := \cup_{\rho \in Q^j, v(\rho) \bowtie u} \mathsf{Cyl}(\pi)$$

Consider some $\alpha \in [0, 1]$. We have:

$$(\mathsf{limsup}_v)^{-1}([0, \alpha]) = \bigcap_{n \in \mathbb{N}} \bigcup_{k \in \mathbb{N}} \bigcap_{j \geq k} V(j, \leq, \alpha + \frac{1}{n})$$

Hence, $\mathsf{limsup}_v^{-1}([0, \alpha])$ is Borel. Furthermore:

$$(\mathsf{liminf}_v)^{-1}([0, \alpha]) = \bigcap_{n \in \mathbb{N}} \bigcap_{k \in \mathbb{N}} \bigcup_{j \geq k} V(j, \leq, \alpha + \frac{1}{n})$$

Hence, $\mathsf{liminf}_v^{-1}([0, \alpha])$ is Borel. It follows that limsup_v and liminf_v are both measurable functions. □

2.7.2 . Proof of Proposition 2.5

Proof. For the strict comparison, we have:

$$\{f < g\} = \bigcup_{q < r \in \mathbb{Q}} (\{f \leq q\} \cap \{r \leq g\})$$

Furthermore, $\{f \leq g\} = Q^\omega \setminus \{f > g\}$. Then, $\{f \neq g\} = \{f < g\} \cup \{f > g\}$ and $\{f = g\} = Q^\omega \setminus \{f \neq g\}$. □

2.7.3 . Proof of Lemma 2.11

Proof. We want to define v_ε and $\mathsf{s_A}$ as (U, m)-uniform functions. Hence, it suffices to define them on $\mathsf{U}^* \times Q$. Since the valuation v is (U, m)-uniform, it is well defined as a function from $\mathsf{U}^* \times Q$ to $[0, 1]$. That is, we have for all $(\gamma, q) \in \mathsf{U}^* \times Q$:

$$v_\varepsilon(\gamma, q) := \max(v(\gamma, q) - \frac{\varepsilon}{2^{|\gamma|}}, 0)$$

For all $w : \mathsf{U}^* \times Q \to [0, 1]$ and $(\gamma, q) \in \mathsf{U}^* \times Q_{\mathsf{ns}}$, we denote by $\mathcal{F}^w_{\gamma,q}$ the game in normal form $\mathcal{F}^w_{\gamma,q} := \langle \mathsf{F}(q), w(\gamma \cdot \mathsf{col}(q), \cdot)\rangle$. We also let $\mathsf{s_A}(\gamma, q) \in S^q_{\mathsf{A}}$ be a $0 < \frac{\varepsilon}{2^{|\gamma|+1}}$-optimal strategy in the game in normal form $\mathcal{F}^v_{\gamma,q}$. Note that we can indeed choose $\mathsf{s_A}(\gamma, q)$ in S^q_{A} since this set supremizes the game form $\mathsf{F}(q)$. Straightforwardly, the strategy $\mathsf{s_A}$ is generated by the collection $(S^q_{\mathsf{A}})_{q\in Q}$. Let us show that it dominates the valuation v_ε.

Let $(\gamma, q) \in \mathsf{U}^* \times Q_{\mathsf{ns}}$ and $n := |\gamma|$. If $v_\varepsilon(\gamma, q) = 0$, then $v_\varepsilon(\gamma, q) \leq \mathsf{val}[\mathcal{F}^{v_\varepsilon}_{\gamma,q}]$ straightforwardly holds. Assume that it is not the case, i.e. $v_\varepsilon(\gamma, q) = v(\gamma, q) - \frac{\varepsilon}{2^n}$. For all Player B GF-strategies $\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}(\mathsf{F}(q))$, by Lemma 1.10 for the first inequality:

$$\begin{aligned}
\mathsf{out}[\mathcal{F}^{v_\varepsilon}_{\gamma,q}](\mathsf{s_A}(\gamma, q), \sigma_{\mathsf{B}}) &= \mathsf{out}[\langle \mathsf{F}(q), v_\varepsilon(\gamma \cdot \mathsf{col}(q), \cdot)\rangle](\mathsf{s_A}(\gamma, q), \sigma_{\mathsf{B}}) \\
&\geq \mathsf{out}[\langle \mathsf{F}(q), v(\gamma \cdot \mathsf{col}(q), \cdot)\rangle](\mathsf{s_A}(\gamma, q), \sigma_{\mathsf{B}}) - \frac{\varepsilon}{2^{n+1}} \\
&= \mathsf{out}[\mathcal{F}^v_{\gamma,q}](\mathsf{s_A}(\gamma, q), \sigma_{\mathsf{B}}) - \frac{\varepsilon}{2^{n+1}} \\
&\geq \mathsf{val}[\mathcal{F}^v_{\gamma,q}](\mathsf{s_A}(\gamma, q)) - \frac{\varepsilon}{2^{n+1}} \\
&\geq \mathsf{val}[\mathcal{F}^v_{\gamma,q}] - \frac{\varepsilon}{2^{n+1}} - \frac{\varepsilon}{2^{n+1}} = \mathsf{val}[\mathcal{F}^v_{\gamma,q}] - \frac{\varepsilon}{2^n}
\end{aligned}$$

This last inequality comes from the fact that $\mathsf{s_A}(\gamma, q)$ is a $\frac{\varepsilon}{2^{n+1}}$-optimal strategy in the game in normal form $\mathcal{F}^v_{\gamma,q}$. Furthermore, since the valuation v is non-decreasing, we have $\mathsf{val}[\mathcal{F}^v_{\gamma,q}] \geq v(\gamma, q)$. Hence, we obtain:

$$\mathsf{out}[\mathcal{F}^{v_\varepsilon}_{\gamma,q}](\mathsf{s_A}(\gamma, q), \sigma_{\mathsf{B}}) \geq v(\gamma, q) - \frac{\varepsilon}{2^n} = v_\varepsilon(\gamma, q)$$

As this holds for all $(\gamma, q) \in \mathsf{U}^* \times Q$, the strategy $\mathsf{s_A}$ dominates the valuation v_ε. □

2.7.4 . Proof of Lemma 2.18

Proof. We want to apply Lemma 1.7. We let $\mathcal{T} := \mathcal{T}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}}$ and $\mathcal{T}^{\mathsf{Act}} := \mathcal{T}^{\mathsf{s^{Sta}_A},\mathsf{s^{Sta}_B}}_{\mathcal{C},\mathsf{Act}}$. We use similar notations for the corresponding probability functions and measures. For instance, the probability measure from any state $q \in Q$ is denoted $\mathbb{P}_q$ for the stochastic tree $\mathcal{T}$ whereas is denoted $\mathbb{P}^{\mathsf{Act}}_q$ for the stochastic tree $\mathcal{T}'$. As mentioned in Observation 2.3, the stochastic tree $\mathcal{T}^{\mathsf{Act}}$ is (Q, Q_{Act})-alternating. For all $q \in Q$ and $\pi \in Q^*$, we denote by $\mathsf{PreIm}_q(\pi) \subseteq (Q_{\mathsf{Act}} \cdot Q)^*$ the set:

$$\mathsf{PreIm}_q(\pi) := (\{q\} \times \mathsf{Act}^q_{\mathsf{A}} \times \mathsf{Act}^q_{\mathsf{B}}) \cdot (\phi_{Q,Q_{\mathsf{Act}}})^{-1}[\{\pi\}] \cap q^{-1} \cdot \mathsf{SeqAdm}^Q_{\mathcal{C}}$$

Furthermore, let $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$ be a player. Recall that by definition of the strategy $\mathsf{s}_{\mathsf{C}}^{\mathsf{Sta}}$, we have, for all $\rho \in \mathsf{SeqAdm}_{\mathcal{C}}^{Q}$:

$$\mathsf{s}_{\mathsf{C}}^{\mathsf{Sta}}(\rho) = \mathsf{s}_{\mathsf{C}} \circ \phi_{Q,Q_{\mathsf{Act}}}(\rho) \tag{2.10}$$

Now, consider some state $q \in Q$. Let us show by induction on $n \in \mathbb{N}$ the property $\mathcal{P}(n)$: for all $\pi \in Q^{\leq n}$, we have:

$$\mathbb{P}_q[\mathsf{Cyl}(\pi)] = \mathbb{P}_q^{\mathsf{Act}}[\cup_{\rho \in \mathsf{T}(\pi)} \mathsf{Cyl}(\rho)]$$

where $\mathsf{T}(\pi) = Q_{\mathsf{Act}} \cdot (\phi_{Q,Q_{\mathsf{Act}}})^{-1}[\{\pi\}]$ comes from Lemma 1.7. This holds for $n = 0$ since the stochastic tree $\mathcal{T}^{\mathsf{Act}}$ is (Q, Q_{Act})-alternating.Consider now some $q' \in Q$. We have, recalling Definition 1.28:

$$\begin{aligned}
\mathbb{P}_q[\mathsf{Cyl}(q')] &= \mathbb{P}_q(q') = \mathsf{out}[\langle \mathsf{F}(q), q' \rangle](\mathsf{s}_{\mathsf{A}}(q), \mathsf{s}_{\mathsf{B}}(q)) \\
&= \sum_{(a,b) \in \mathsf{Act}_{\mathsf{A}} \times \mathsf{Act}_{\mathsf{B}}} \mathsf{s}_{\mathsf{A}}(q)(a) \cdot \mathsf{s}_{\mathsf{B}}(q)(b) \cdot \varrho_q(a,b)(q') \\
&= \sum_{(a,b) \in \mathsf{Act}_{\mathsf{A}} \times \mathsf{Act}_{\mathsf{B}}} \mathbb{P}_q^{\mathsf{Act}}(q, (q,a,b)) \cdot \mathbb{P}_q^{\mathsf{Act}}((q,a,b), q') \\
&= \mathbb{P}_q^{\mathsf{Act}}[Q_{\mathsf{Act}} \cdot \mathsf{Cyl}(q')] = \mathbb{P}_q^{\mathsf{Act}}[\cup_{\rho \in \mathsf{T}(\pi)} \mathsf{Cyl}(q)]
\end{aligned}$$

Hence, $\mathcal{P}(n)$ also holds. Assume now that $\mathcal{P}(n)$ holds for some $n \geq 1$. Consider some $\pi \in Q^{n+1}$. By Observation 2.3, for all $\rho \notin q^{-1} \cdot \mathsf{SeqAdm}_{\mathcal{C}}$, we have $\mathbb{P}_q^{\mathsf{Act}}(\rho) = \mathbb{P}_q^{\mathsf{Act}}[\mathsf{Cyl}(\rho)] = 0$. Hence:

$$\mathbb{P}_q^{\mathsf{Act}}[\cup_{\rho \in \mathsf{T}(\pi)} \mathsf{Cyl}(\rho)] = \mathbb{P}_q^{\mathsf{Act}}[\cup_{\rho \in \mathsf{PreIm}_q(\pi)} \mathsf{Cyl}(\rho)] \tag{2.11}$$

Furthermore, by definition, we have:

$$\mathsf{PreIm}_q(\pi) = \{\rho \cdot (\rho_{\mathsf{lt}}, a, b) \cdot \pi_{\mathsf{lt}} \mid \rho \in \mathsf{PreIm}_q(\mathsf{tl}(\pi)),\ (a,b) \in \mathsf{Act}_{\mathsf{A}}^{\rho_{\mathsf{lt}}} \times \mathsf{Act}_{\mathsf{B}}^{\rho_{\mathsf{lt}}}\}$$

Since holds because, since $n \geq 1$, we have $\mathsf{tl}(\pi) \neq \epsilon$ and therefore, for all $\rho \in \mathsf{PreIm}_q(\mathsf{tl}(\pi))$, we have $\rho_{\mathsf{lt}} \in Q$. It follows that:

$$\begin{aligned}
\mathbb{P}_q^{\mathsf{Act}}[\bigcup_{\rho \in \mathsf{PreIm}_q(\pi)} \mathsf{Cyl}(\rho)] &= \sum_{\rho \in \mathsf{PreIm}_q(\mathsf{tl}(\pi))} \sum_{(a,b) \in \mathsf{Act}_{\mathsf{A}}^{\rho_{\mathsf{lt}}} \times \mathsf{Act}_{\mathsf{B}}^{\rho_{\mathsf{lt}}}} \mathbb{P}_q^{\mathsf{Act}}[\mathsf{Cyl}(\rho \cdot (\rho_{\mathsf{lt}}, a, b) \cdot \pi_{\mathsf{lt}})] \\
&= \sum_{\rho \in \mathsf{PreIm}_q(\mathsf{tl}(\pi))} \sum_{(a,b) \in \mathsf{Act}_{\mathsf{A}}^{\rho_{\mathsf{lt}}} \times \mathsf{Act}_{\mathsf{B}}^{\rho_{\mathsf{lt}}}} \mathbb{P}_q^{\mathsf{Act}}(\rho \cdot (\rho_{\mathsf{lt}}, a, b) \cdot \pi_{\mathsf{lt}})
\end{aligned}$$

Consider some $\rho \in \mathsf{PreIm}_q(\mathsf{tl}(\pi))$ and $(a,b) \in \mathsf{Act}_{\mathsf{A}}^{\rho_{\mathsf{lt}}} \times \mathsf{Act}_{\mathsf{B}}^{\rho_{\mathsf{lt}}}$. In particular, we have $(\mathsf{tl}(\pi))_{\mathsf{lt}} = \rho_{\mathsf{lt}}$. We have, recalling Definition 2.15 and Equation (2.10):

$$\begin{aligned}
\mathbb{P}_q^{\mathsf{Act}}(\rho \cdot (\rho_{\mathsf{lt}}, a, b) \cdot q') &= \mathbb{P}_q^{\mathsf{Act}}(\rho) \cdot \mathbb{P}_{q \cdot \rho}^{\mathsf{Act}}((\rho_{\mathsf{lt}}, a, b)) \cdot \mathbb{P}_{q \cdot \rho \cdot (\rho_{\mathsf{lt}}, a, b)}^{\mathsf{Act}}(q') \\
&= \mathbb{P}_q^{\mathsf{Act}}(\rho) \cdot \mathsf{s}_{\mathsf{A}}^{\mathsf{Sta}}(q \cdot \rho)(a) \cdot \mathsf{s}_{\mathsf{B}}^{\mathsf{Sta}}(q \cdot \rho)(b) \cdot \mathbb{P}_{q \cdot \rho \cdot (\rho_{\mathsf{lt}}, a, b)}^{\mathsf{Act}}(q') \\
&= \mathbb{P}_q^{\mathsf{Act}}(\rho) \cdot \mathsf{s}_{\mathsf{A}}^{\mathsf{Sta}}(q \cdot \rho)(a) \cdot \mathsf{s}_{\mathsf{B}}^{\mathsf{Sta}}(q \cdot \rho)(b) \cdot \varrho_{\rho_{\mathsf{lt}}}(a,b)(q') \\
&= \mathbb{P}_q^{\mathsf{Act}}(\rho) \cdot \mathsf{s}_{\mathsf{A}} \circ \phi_{Q,Q_{\mathsf{Act}}}(q \cdot \rho)(a) \cdot \mathsf{s}_{\mathsf{B}} \circ \phi_{Q,Q_{\mathsf{Act}}}(q \cdot \rho)(b) \cdot \varrho_{(\mathsf{tl}(\pi))_{\mathsf{lt}}}(a,b)(q') \\
&= \mathbb{P}_q^{\mathsf{Act}}(\rho) \cdot \mathsf{s}_{\mathsf{A}}(q \cdot \mathsf{tl}(\pi))(a) \cdot \mathsf{s}_{\mathsf{B}}(q \cdot \mathsf{tl}(\pi))(b) \cdot \varrho_{(\mathsf{tl}(\pi))_{\mathsf{lt}}}(a,b)(q')
\end{aligned}$$

Furthermore, denoting $\pi' := \mathsf{tl}(\pi)$ and $q' := \pi_{\mathsf{lt}}$, we have:

$$\sum_{(a,b) \in \mathsf{Act}_{\mathsf{A}}^{\rho_{\mathsf{lt}}} \times \mathsf{Act}_{\mathsf{B}}^{\rho_{\mathsf{lt}}}} \mathsf{s}_{\mathsf{A}}(q \cdot \pi')(a) \cdot \mathsf{s}_{\mathsf{B}}(q \cdot \pi')(b) \cdot \varrho_{\pi'_{\mathsf{lt}}}(a,b)(q') = \mathsf{out}[\langle \mathsf{F}(\pi'_{\mathsf{lt}}), q' \rangle](\mathsf{s}_{\mathsf{A}}(q \cdot \pi'), \mathsf{s}_{\mathsf{B}}(q \cdot \pi'))$$

By Definitions 1.28 and 1.29:

$$\mathsf{out}[\langle \mathsf{F}(\pi'_{\mathsf{lt}}), q' \rangle](\mathsf{s}_{\mathsf{A}}(q \cdot \pi'), \mathsf{s}_{\mathsf{B}}(q \cdot \pi')) = \mathbb{P}_{\mathcal{C}}^{\mathsf{s}_{\mathsf{A}}(q \cdot \pi'), \mathsf{s}_{\mathsf{B}}(q \cdot \pi')}(\pi'_{\mathsf{lt}}, q') = \mathbb{P}_{q \cdot \pi'}(q')$$

Overall, we obtain, by $\mathcal{P}(n)$ applied to $\pi' = \mathsf{tl}(\pi) \in Q^n$, recalling that $q' = \pi_{\mathsf{lt}}$:

$$\begin{aligned}
\mathbb{P}_q^{\mathsf{Act}}[\bigcup_{\rho \in \mathsf{PreIm}_q(\pi)} \mathsf{Cyl}(\rho)] &= \sum_{\rho \in \mathsf{PreIm}_q(\mathsf{tl}(\pi))} \sum_{(a,b) \in \mathsf{Act}_{\mathsf{A}}^{\rho_{\mathsf{lt}}} \times \mathsf{Act}_{\mathsf{B}}^{\rho_{\mathsf{lt}}}} \mathbb{P}_q^{\mathsf{Act}}(\rho \cdot (\rho_{\mathsf{lt}}, a, b) \cdot \pi_{\mathsf{lt}}) \\
&= \sum_{\rho \in \mathsf{PreIm}_q(\mathsf{tl}(\pi))} \mathbb{P}_q^{\mathsf{Act}}(\rho) \cdot \mathbb{P}_{q \cdot \mathsf{tl}(\pi)}(\pi_{\mathsf{lt}}) \\
&= \mathbb{P}_q[\mathsf{Cyl}(\mathsf{tl}(\pi))] \cdot \mathbb{P}_{q \cdot \mathsf{tl}(\pi)}(\pi_{\mathsf{lt}}) = \mathbb{P}_q(\mathsf{tl}(\pi)) \cdot \mathbb{P}_{q \cdot \mathsf{tl}(\pi)}(\pi_{\mathsf{lt}}) \\
&= \mathbb{P}_q(\mathsf{tl}(\pi) \cdot \pi_{\mathsf{lt}}) = \mathbb{P}_q[\mathsf{Cyl}(\pi)]
\end{aligned}$$

By Equation (2.11), we obtain:

$$\mathbb{P}_q[\mathsf{Cyl}(\pi)] = \mathbb{P}_q^{\mathsf{Act}}[\cup_{\rho \in \mathsf{T}(\pi)} \mathsf{Cyl}(\rho)]$$

Hence, $\mathcal{P}(n+1)$ holds. We conclude by applying Lemma 1.7. □

2.7.5 . Proof of Lemma 2.20

Proof. First, note that the state space in both stochastic trees $\mathcal{T}_{\mathcal{C},\mathsf{Act},q}^{\mathsf{s}_{\mathsf{A}}^{\mathsf{Act}},\mathsf{s}_{\mathsf{B}}^{\mathsf{Act}}}$ and $\mathcal{T}_{\mathcal{C}_{\mathsf{Act}},q}^{\mathsf{s}_{\mathsf{A}},\mathsf{s}_{\mathsf{B}}}$ is the same: it is equal to $(Q \cup Q_{\mathsf{Act}})$. Let us show by induction on $n \in \mathbb{N}$ the property $\mathcal{P}(n)$: for all $\pi \in (Q \cup Q_{\mathsf{Act}})^{\leq n}$, we have $\mathbb{P}_{\mathcal{C},\mathsf{Act},q}^{\mathsf{s}_{\mathsf{A}}^{\mathsf{Act}},\mathsf{s}_{\mathsf{B}}^{\mathsf{Act}}}(\pi) = \mathbb{P}_{\mathcal{C}_{\mathsf{Act}},q}^{\mathsf{s}_{\mathsf{A}},\mathsf{s}_{\mathsf{B}}}(\pi)$.

This straightforwardly holds for $n = 0$. Assume now that $\mathcal{P}(n)$ holds for some $n \in \mathbb{N}$. Let $\pi \in (Q \cup Q_{\mathsf{Act}})^{n+1}$. Clearly, if $\pi \notin q^{-1} \cdot \mathsf{SeqAdm}_{\mathcal{C}}$, we have $\mathbb{P}_{\mathcal{C}_{\mathsf{Act}},q}^{\mathsf{s}_{\mathsf{A}},\mathsf{s}_{\mathsf{B}}}(\pi) = 0$ (recall Definition 2.17). Similarly, $\mathbb{P}_{\mathcal{C},\mathsf{Act},q}^{\mathsf{s}_{\mathsf{A}}^{\mathsf{Act}},\mathsf{s}_{\mathsf{B}}^{\mathsf{Act}}}(\pi) = 0$ (recall Observation 2.3). Assume now that $\pi \in q^{-1} \cdot \mathsf{SeqAdm}_{\mathcal{C}}$. Let $\rho := q \cdot \mathsf{tl}(\pi)$. Assume that $\pi_{\mathsf{lt}} \in Q$ and $\rho_{\mathsf{lt}} = (q', a, b) \in Q_{\mathsf{Act}}$. In that case:

$$\begin{aligned}
\mathbb{P}_{\mathcal{C},\mathsf{Act},q}^{\mathsf{s}_{\mathsf{A}}^{\mathsf{Act}},\mathsf{s}_{\mathsf{B}}^{\mathsf{Act}}}(\pi) &= \mathbb{P}_{\mathcal{C},\mathsf{Act},q}^{\mathsf{s}_{\mathsf{A}}^{\mathsf{Act}},\mathsf{s}_{\mathsf{B}}^{\mathsf{Act}}}(\mathsf{tl}(\pi)) \cdot \mathbb{P}_{\mathsf{Act}}^{\mathcal{D}}(\rho_{\mathsf{lt}}, \pi_{\mathsf{lt}}) && \text{by Definition 2.15} \\
&= \mathbb{P}_{\mathcal{C},\mathsf{Act},q}^{\mathsf{s}_{\mathsf{A}}^{\mathsf{Act}},\mathsf{s}_{\mathsf{B}}^{\mathsf{Act}}}(\mathsf{tl}(\rho)) \cdot \varrho_{q'}^{\mathsf{Act}}(a,b)(\pi_{\mathsf{lt}}) && \text{by Definition 2.14} \\
&= \mathbb{P}_{\mathcal{C}_{\mathsf{Act}},q}^{\mathsf{s}_{\mathsf{A}},\mathsf{s}_{\mathsf{B}}}(\mathsf{tl}(\pi)) \cdot \varrho_{q'}(a,b)(\pi_{\mathsf{lt}}) && \text{by } \mathcal{P}(n) \\
&= \mathbb{P}_{\mathcal{C}_{\mathsf{Act}},q}^{\mathsf{s}_{\mathsf{A}},\mathsf{s}_{\mathsf{B}}}(\mathsf{tl}(\pi)) \cdot \mathsf{out}[\langle \mathsf{F}_{\mathsf{Act}}(\rho_{\mathsf{lt}}), \pi_{\mathsf{lt}} \rangle](\mathsf{s}_{\mathsf{A}}(\rho), \mathsf{s}_{\mathsf{B}}(\rho)) && \text{by Definition 2.17} \\
&= \mathbb{P}_{\mathcal{C}_{\mathsf{Act}},q}^{\mathsf{s}_{\mathsf{A}},\mathsf{s}_{\mathsf{B}}}(\mathsf{tl}(\pi)) \cdot \mathbb{P}_{\mathcal{C}_{\mathsf{Act}}}^{\mathsf{s}_{\mathsf{A}}(\rho),\mathsf{s}_{\mathsf{B}}(\rho)}(\rho_{\mathsf{lt}}, \pi_{\mathsf{lt}}) && \text{by Definition 1.28} \\
&= \mathbb{P}_{\mathcal{C}_{\mathsf{Act}},q}^{\mathsf{s}_{\mathsf{A}},\mathsf{s}_{\mathsf{B}}}(\rho) && \text{by Definition 1.29}
\end{aligned}$$

Assume now that $\rho_{\mathsf{lt}} \in Q$ and $\pi_{\mathsf{lt}} = (\rho_{\mathsf{lt}}, a, b) \in Q_{\mathsf{Act}}$. In that case:

$$\begin{aligned}
\mathbb{P}^{\mathsf{s}_\mathsf{A}^{\mathsf{Act}},\mathsf{s}_\mathsf{B}^{\mathsf{Act}}}_{\mathcal{C},\mathsf{Act},q}(\pi) &= \mathbb{P}^{\mathsf{s}_\mathsf{A}^{\mathsf{Act}},\mathsf{s}_\mathsf{B}^{\mathsf{Act}}}_{\mathcal{C},\mathsf{Act},q}(\mathsf{tl}(\pi)) \cdot \mathbb{P}^{\mathsf{s}_\mathsf{A}^{\mathsf{Act}}(\rho),\mathsf{s}_\mathsf{B}^{\mathsf{Act}}(\rho)}_{\mathsf{Act}}(\rho_{\mathsf{lt}}, \pi_{\mathsf{lt}}) && \text{by Definition 2.15} \\
&= \mathbb{P}^{\mathsf{s}_\mathsf{A}^{\mathsf{Act}},\mathsf{s}_\mathsf{B}^{\mathsf{Act}}}_{\mathcal{C},\mathsf{Act},q}(\mathsf{tl}(\pi)) \cdot \mathsf{s}_\mathsf{A}^{\mathsf{Act}}(\rho)(a) \cdot \mathsf{s}_\mathsf{B}^{\mathsf{Act}}(\rho)(b) && \text{by Definition 2.14} \\
&= \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C}_{\mathsf{Act}},q}(\mathsf{tl}(\pi)) \cdot \mathsf{s}_\mathsf{A}^{\mathsf{Act}}(\rho)(a) \cdot \mathsf{s}_\mathsf{B}^{\mathsf{Act}}(\rho)(b) && \text{by } \mathcal{P}(n) \\
&= \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C}_{\mathsf{Act}},q}(\mathsf{tl}(\pi)) \cdot \mathsf{s}_\mathsf{A}(\rho)(a) \cdot \mathsf{s}_\mathsf{B}(\rho)(b) && \text{by Definition 2.18} \\
&= \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C}_{\mathsf{Act}},q}(\mathsf{tl}(\pi)) \cdot \mathsf{out}[\langle \mathsf{F}_{\mathsf{Act}}(\rho_{\mathsf{lt}}), \pi_{\mathsf{lt}} \rangle](\mathsf{s}_\mathsf{A}(\rho), \mathsf{s}_\mathsf{B}(\rho)) && \text{by Definition 2.17} \\
&= \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C}_{\mathsf{Act}},q}(\mathsf{tl}(\pi)) \cdot \mathbb{P}^{\mathsf{s}_\mathsf{A}(\rho),\mathsf{s}_\mathsf{B}(\rho)}_{\mathcal{C}_{\mathsf{Act}}}(\rho_{\mathsf{lt}}, \pi_{\mathsf{lt}}) && \text{by Definition 1.28} \\
&= \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C}_{\mathsf{Act}},q}(\pi) && \text{by Definition 1.29}
\end{aligned}$$

Overall, $\mathcal{P}(n+1)$ and therefore $\mathcal{P}(n)$ holds for all $n \in \mathbb{N}$. It follows, by Theorem 1.2, that $\mathbb{P}^{\mathsf{s}_\mathsf{A}^{\mathsf{Act}},\mathsf{s}_\mathsf{B}^{\mathsf{Act}}}_{\mathcal{C},\mathsf{Act},q} = \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C}_{\mathsf{Act}},q} : \mathsf{Borel}((Q \cup Q_{\mathsf{Act}})) \to [0,1]$. Let us denote by $\mathbb{P}_q$ this probability measure.

As mentioned above, for all $\rho \in (Q \cup Q_{\mathsf{Act}})^*$ such that $\rho \notin Q^{-1} \cdot \mathsf{SeqAdm}_{\mathcal{C}}$, we have $\mathbb{P}_q(\rho) = 0$. Hence, $\mathbb{P}[(Q \cup Q_{\mathsf{Act}})^\omega \setminus q^{-1} \cdot \mathsf{SeqAdm}^\omega_{\mathcal{C}}] = 0$. Consider some $\rho \in q^{-1} \cdot \mathsf{SeqAdm}^\omega_{\mathcal{C}}$. Let us write ρ as $\rho = (q_0, a_0, b_0) \cdot q_1 \cdot (q_1, a_1, b_1) \cdot q_2 \cdot (q_2, a_2, b_2) \cdots$ (with $q_0 := q$). If there is some $i \in \mathbb{N}$ such that $q_i \in Q_\mathsf{s}$, then considering the least index $i_\mathsf{s} \in \mathbb{N}$ such that $q_{i_\mathsf{s}} \in Q_\mathsf{s}$, then we have:

$$\begin{aligned}
(f_{\mathcal{C}})_{Q,Q_{\mathsf{Act}}}(q \cdot \rho) &= f_{\mathcal{C}} \circ \phi_{Q,Q_{\mathsf{Act}}}(q \cdot \rho) && \text{by Definition 1.8} \\
&= f_{\mathcal{C}}(q_0 \cdot q_1 \cdot q_2 \cdots) && \text{by Definition 1.8} \\
&= \mathsf{val}(q_{i_\mathsf{s}}) && \text{by Definition 1.30} \\
&= (f_{\mathsf{Act}})_{\mathcal{C}_{\mathsf{Act}}}(q \cdot \rho) && \text{by Definition 1.30}
\end{aligned}$$

If that is not the case, i.e. for all $i \in \mathbb{N}$, we have $q_i \in Q_{\mathsf{ns}}$, then we have:

$$\begin{aligned}
(f_{\mathcal{C}})_{Q,Q_{\mathsf{Act}}}(q \cdot \rho) &= f_{\mathcal{C}} \circ \phi_{Q,Q_{\mathsf{Act}}}(q \cdot \rho) && \text{by Definition 1.8} \\
&= f_{\mathcal{C}}(q_0 \cdot q_1 \cdot q_2 \cdots) && \text{by Definition 1.8} \\
&= f_{\mathcal{C}} \circ (\mathsf{f}^{\mathsf{sg}})^\omega(q_0 \cdot q_0 \cdot q_1 \cdot q_1 \cdots) && \text{by Definition 2.17} \\
&= f_{\mathsf{Act}}(q_0 \cdot q_0 \cdot q_1 \cdot q_1 \cdots) && \text{by Definition 2.17} \\
&= f_{\mathsf{Act}} \circ \mathsf{col}^\omega_{\mathsf{Act}}(q_0 \cdot (q_0, a_0, b_0) \cdot q_1 \cdot (q_1, a_1, b_1) \cdots) && \text{by Definition 2.17} \\
&= (f_{\mathsf{Act}})_{\mathcal{C}_{\mathsf{Act}}}(q \cdot \rho) && \text{by Definition 1.30}
\end{aligned}$$

That is, $((f_{\mathcal{C}})_{Q,Q_{\mathsf{Act}}})^q$ and $(f_{\mathsf{Act}} \circ \mathsf{col}^\omega_{\mathsf{Act}})^q$ coincide on $\mathsf{SeqAdm}^\omega_{\mathcal{C}}$. Hence: $\mathbb{E}^{\mathsf{s}_\mathsf{A}^{\mathsf{Act}},\mathsf{s}_\mathsf{B}^{\mathsf{Act}}}_{\mathcal{C},\mathsf{Act},q}[((f^Q)_{Q,Q_{\mathsf{Act}}})^q] = \mathbb{E}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C}_{\mathsf{Act}},q}[((f_{\mathsf{Act}})_{\mathcal{C}_{\mathsf{Act}}})^q]$, thus proving the lemma. □

3 - On subgame optimal strategies

In this chapter, we study subgame (ε-)optimal strategies. These are strategies that are not only (ε-)optimal from any state, but that also are (ε-)optimal after any sequence of states is seen. In the specific case of games with stopping states, we only consider sequences of states that stop when reaching a stopping state. Hence, in this chapter, we will frequently use the notations below, recalling Definition 1.18: $Q_{\mathsf{s}} \subseteq Q$ (resp. $Q_{\mathsf{ns}} \subseteq Q$) refers to the set of stopping states (resp. non-stopping states) in Q.

Definition 3.1. *Given a concurrent arena $\mathcal{C}$, for all $n \in \mathbb{N}$, we denote by $\Omega_{\mathcal{C}}^*$, $\Omega_{\mathcal{C}}^+$, $\Omega_{\mathcal{C}}^n$, $\Omega_{\mathcal{C}}^{\leq n}$ and $\Omega_{\mathcal{C}}^\omega$ the following sets: $\Omega_{\mathcal{C}}^* := (Q_{\mathsf{ns}})^* \cup (Q_{\mathsf{ns}})^* \cdot Q_{\mathsf{s}}$, $\Omega_{\mathcal{C}}^+ := (Q_{\mathsf{ns}})^+ \cup (Q_{\mathsf{ns}})^* \cdot Q_{\mathsf{s}}$, $\Omega_{\mathcal{C}}^n := \Omega_{\mathcal{C}}^* \cap Q^n$, $\Omega_{\mathcal{C}}^{\leq n} := \Omega_{\mathcal{C}}^* \cap Q^{\leq n}$ and $\Omega_{\mathcal{C}}^\omega := (Q_{\mathsf{ns}})^\omega \cup (Q_{\mathsf{ns}})^* \cdot Q_{\mathsf{s}}$.*

We now define the notion of strategy which guarantees a valuation, which allows us to define the notion of subgame (ε-)optimal strategies.

Definition 3.2 (Strategy which guarantees a valuation)**.** *Consider a concurrent game $\mathcal{G} = \langle \mathcal{C}, f \rangle$. For all $\pi \in (Q_{\mathsf{ns}})^*$, we denote by $\mathcal{G}^\pi$ the game $\mathcal{G}^\pi := \langle \mathcal{C}, f^{\mathsf{col}^+(\pi)} \rangle$. Recall that $f^{\mathsf{col}^+(\pi)} : \mathsf{K}^\omega \to [0,1]$ is the residual function such that, for all $\rho \in \mathsf{K}^\omega$, we have $f^{\mathsf{col}^+(\pi)} = f(\mathsf{col}^+(\pi) \cdot \rho)$. (In particular, $\mathcal{G}^\epsilon = \mathcal{G}$.) Consider some $v : \Omega_{\mathcal{C}}^+ \to [0,1]$. A Player-$\mathsf{A}$ strategy $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$* guarantees *the valuation v if for all $\pi \in \Omega_{\mathcal{C}}^+$, the value of the residual strategy $\mathsf{s}_{\mathsf{A}}^{\mathsf{tl}(\pi)} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ from π_{lt} is at least $v(\pi)$: $\chi_{\mathcal{G}^{\mathsf{tl}(\pi)}}[\mathsf{s}_{\mathsf{A}}^{\mathsf{tl}(\pi)}](\pi_{\mathsf{lt}}) \geq v(\pi)$. This is symmetrical for Player B.*

Definition 3.3 (Subgame (ε-)optimal strategies)**.** *Consider a concurrent game $\mathcal{G} = \langle \mathcal{C}, f \rangle$. We extend the Player-$\mathsf{A}$ valuation of the states into a valuation of finite sequences of states not continuing after a stopping state: $\chi_{\mathcal{G}}[\mathsf{A}] : \Omega_{\mathcal{C}}^+ \to [0,1]$ such that, for all $\pi \in \Omega_{\mathcal{C}}^+$, we have $\chi_{\mathcal{G}}[\mathsf{A}](\pi) := \chi_{\mathcal{G}^{\mathsf{tl}(\pi)}}[\mathsf{A}](\pi_{\mathsf{lt}})$. We define similarly $\chi_{\mathcal{G}}[\mathsf{s_A}] : \Omega_{\mathcal{C}}^+ \to [0,1]$ for a Player-A strategy $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$.*

Then, for all $\varepsilon \geq 0$, a Player-A strategy $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ is subgame ε-optimal *if it guarantees the valuation $\max(\chi_{\mathcal{G}}[\mathsf{A}] - \varepsilon, 0) : (\Omega_{\mathcal{C}})^+ \to [0,1]$. When $\varepsilon = 0$, the strategy $\mathsf{s_A}$ is simply said to be* subgame optimal*. This is symmetrical for Player B.*

Given an arbitrary strategy, being subgame optimal is stronger than being optimal since being subgame optimal requires to be optimal after every finite sequence of states. The difference in strength between these two notions is particularly visible after finite histories where the other player has made a mistake, i.e. has not played optimally against the strategy considered. In such a situation, an optimal strategy could (and, in fact, sometimes should, as it will be seen in item 2.b below) also make a mistake as long as it is small enough. On the other hand, a subgame optimal strategy cannot do so and needs to still be optimal. Therefore, subgame optimality can be seen as more robust than optimality. In addition, since subgame optimal strategies satisfy more properties than optimal strategies, it is easier to characterize properties related to them. This is why, in this chapter, we study them, along with subgame ε-optimal strategies.

Let us first give the big picture of what is done in this chapter, we will give the details afterwards. There are four sections in this chapter. The first section is the only one where we deal with subgame ε-optimal strategies. Contrary to subgame optimal strategies, subgame ε-optimal always exist, for all $\varepsilon > 0$. This is formally stated and proved in that section. In the three remaining sections, we deal with subgame optimal strategies. In the second section, we give two characterizations related to subgame optimal strategies that will be used afterwards: one stating at which conditions a strategy is subgame optimal; another one stating at which conditions there exist subgame optimal strategies. These characterizations are then used in the two following sections. The last two sections can be seen as applications of the results proved in this second section. In the third section, we look at how to use these results in the context of standard finite concurrent (possibly turn-based) games. In the fourth and last section, we study some conditions under which we can transfer results existing in finite turn-based games to the context of standard finite concurrent games.

Let us be more specific. In the following, arbitrary payoff functions will always refer to measurable functions taking their values in $[0, 1]$. As mentioned above, in Section 3.1, we focus on subgame ε-optimal strategies for $\varepsilon > 0$. We show the following:

1.a. It is already known (see [57, Proposition 11, Lemma 12]) that, if at each state both players have finitely many actions, then for all positive $\varepsilon > 0$, both players have subgame ε-optimal strategies. We generalize this result to arbitrary games, see Theorem 3.1, while keeping essentially the same proof, of which we explain the main ideas, namely reset strategies. We use Theorem 2.3 (item $(1.a)$) to prove this result.

1.b. We then use Theorem 3.1 to deduce a result on prefix-independent (PI) win/lose games. Namely, in all PI win/lose (possibly infinite) games where the infimum of the states values is positive, Player A has subgame almost-surely winning strategies. We have already proved this result in [41, Theorem 3] in the context of finite-state games. In fact, we show a slightly more general result by only assuming arbitrary PI upward well-founded (notion to be defined) payoff functions, see Corollary 3.6.

In Section 3.2, we focus on subgame optimal strategies in PI games. This section is almost entirely based on [41] except that we do not only consider win/lose objectives, but more general payoff functions: sometimes arbitrary PI payoff functions, sometimes only upward well-founded PI ones (notion to be defined). We show several results:

2.a. We provide a characterization of subgame optimal strategies in PI games: a Player-A strategy is subgame optimal if and only if 1) it is locally optimal and 2) for every Player-B deterministic strategy, after every history, almost-surely the (superior) limit of the Player-A value of the states visited is less than or equal to the payoff function, see Theorem 3.12. We then consider what Theorem 3.12 amounts to in the special cases where the game is finite-state (Corollary 3.14) and if we additionally assume that the Player-A strategy considered is positional (Corollary 3.16). These are the results we mentioned at the beginning of this part as being key results in concurrent games.

2.b. In [58], the authors have shown that in finite PI win/lose turn-based games, there always exist optimal strategies and the memory sufficient to play optimally is equal to the memory sufficient to play almost-surely, in games where this is possible. We generalize this result to arbitrary finite-state concurrent games with PI upward well-founded payoff functions f. In such a context, subgame optimal strategies do not always exist (and neither do optimal strategies), however we exhibit necessary and sufficient conditions for the existence of subgame optimal strategies. We give the intuition behind these necessary and sufficient conditions with the help of an example where there is an optimal strategy, but there is no subgame optimal one, see Page 134[1]. As a bi-product of the proof that these conditions are indeed necessary and sufficient, we deduce that if every game with a win/lose objective obtained from f via a threshold that has a subgame almost-surely winning strategy also has a positional one, then every game that has a subgame optimal strategy also has a positional one. Note that this transfer result also holds with finite memory, see Theorem 3.17.

Third, Section 3.4, which is also based on [41], we focus on subgame optimal strategies in standard concurrent games. We consider two different issues:

3.a. We first focus on turn-based games and apply Theorem 3.17 discussed above to finite turn-based games. That is, we recover the results proved in [58] — dealing with the existence of (subgame) optimal strategies. That is, we show that in finite turn-based games with PI upward well-founded payoff function, Player A has a subgame optimal strategy along

[1]This is a game where, as hinted above, playing optimally requires "making a mistake". Formally, this means that no locally optimal strategy (notion to be defined) is optimal.

with the previously-mentioned memory transfer. We then build a payoff function from the parity objective to illustrate our result, see Corollary 3.27.

3.b. We then come back to standard finite concurrent games with PI upward well-founded payoff functions. In general, it does not hold, as as mentioned in item 2.b., that the existence of optimal strategies implies the existence of subgame optimal strategies. Here, we exhibit a structural condition — that we call being positively bounded — on strategies so that the existence of a positively bounded optimal strategies is equivalent to the existence of a positively bounded subgame optimal strategies, see Theorem 3.28. This structural condition only depends on what the strategy does in the arena, regardless of the payoff function. In this case, the structural condition we consider is for a strategy to not play arbitrarily small yet positive probabilities.

Finally, Section 3.4 is based on [41] but also borrows ideas from [38]. In that section, we focus on transferring already existing result in turn-based games to the case of standard concurrent games. We proceed in two steps

4.a. First, in standard concurrent games, we define the notion of sequentialization, that is we build a turn-based game from standard concurrent games where Player A plays first, and then Player B responds. We translate strategies back and forth between the two games, which allows us to consider how and when the values change between the two games. Note that the results of this subsection will be used in Chapter 6.

4.b. Then, we introduce another type of strategies, namely finite-choice strategies. Informally, a strategy has finite choice if it uses only finitely many GF-strategies at each state. This is stronger than being positively bounded (in finite-state arenas). We use the sequentialization from item 4.a to show that when such strategies exist, we can transfer already existing results in turn-based games to concurrent games, for some payoff functions. Note that the condition on the payoff functions is unrelated with being PI or upward well founded. As a corollary, we obtain that in finite concurrent games with a parity objective, if there is a subgame optimal strategy that has finite choice, then there is one that is positional, see Corollary 3.38.

3.1 Subgame almost-optimal strategies

In this section, we focus on subgame almost-optimal strategies, that is subgame ε-optimal strategies for all $\varepsilon > 0$. It is shown in [57, Proposition

11, Lemma 12] that subgame almost-optimal strategies always exist in standard concurrent games[2] where, at each state, both players have finitely many actions. We adapt the proof of [57] to show the existence of subgame almost-optimal strategies in all concurrent games. Furthermore, one can realize that in this proof subgame almost-optimal strategies are built from almost-optimal strategies. Since we have exhibited in Theorem 2.3 a restriction on the class of strategies we need to consider to find almost-optimal strategies, we prove the theorem below:

Theorem 3.1. *Consider an arbitrary concurrent game $\mathcal{G}$. Let $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$ be a player and assumed that the game $\mathcal{G}$ is supremized w.r.t. Player C by a collection $(S^q_{\mathsf{C}})_{q \in Q}$ of sets of GF-strategies. Then, for all positive $\varepsilon > 0$, Player C has a subgame ε-optimal strategy generated by $(S^q_{\mathsf{C}})_{q \in Q}$.*

In the next subsection below, we discuss the proof of this theorem, whereas Subsection 3.1.2 is dedicated to an application of this theorem.

3.1.1 . Proof of Theorem 3.1

In this subsection, we prove Theorem 3.1. The result is proved for Player A, the case of Player B being analogous. For the remainder of this subsection we fix an arbitrary concurrent game $\mathcal{G} = \langle \mathcal{C}, g \rangle$. We also let $f := g_{\mathcal{C}} : Q^\omega \to [0, 1]$.

The idea behind the construction of subgame ε-optimal strategies is to make use of reset strategies. Informally, these are strategies that are initialized at the beginning of the game and updated whenever the property they are supposed to ensure does not hold after some history (i.e. finite sequence of states). The goal is then to show that almost-surely there are only finitely many updates. Hence, almost-surely, if we consider long enough history, the strategy is not changed anymore which means that it ensures the property for all finite histories thereafter. For all positive $\varepsilon > 0$ and finite history $\rho \in \Omega^+_{\mathcal{C}}$, we denote by $\mathsf{s}_{\varepsilon,\rho} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{A}}$ a Player-A strategy that is ε-optimal in the game $\mathcal{G}^{\mathsf{tl}(\rho)}$ from the state ρ_{lt}. We define formally below the reset strategies we consider.

Definition 3.4 (Reset strategy). *Consider some positive $\varepsilon > 0$. We define inductively a function $\mathsf{U}_\varepsilon : \Omega^+_{\mathcal{C}} \to \Omega^+_{\mathcal{C}}$ that ensures that, for all $\rho \in \Omega^+_{\mathcal{C}}$, we have $\mathsf{U}_\varepsilon(\rho) \sqsubseteq \rho$. We also denote by $\mathsf{Suf}(\rho) \in \Omega^*_{\mathcal{C}}$ the finite path such that $\rho = \mathsf{U}_\varepsilon(\rho) \cdot \mathsf{Suf}(\rho)$ and by $\mathsf{Pl}(\rho) := \mathsf{U}_\varepsilon(\rho)_{\mathsf{lt}} \cdot \mathsf{Suf}(\rho) \in \Omega^+_{\mathcal{C}}$. For all $q \in Q$, we let $\mathsf{U}_\varepsilon(q) := q$. Furthermore, for all $\rho \cdot q \in \Omega^+_{\mathcal{C}}$, we let:*

$$\mathsf{U}_\varepsilon(\rho \cdot q) := \begin{cases} \mathsf{U}_\varepsilon(\rho) & \text{if } \chi_{\mathcal{G}^\rho}[\mathsf{s}^{\mathsf{Pl}(\rho)}_{\varepsilon,\mathsf{U}_\varepsilon(\rho)}](q) \geq \chi_{\mathcal{G}^\rho}[\mathsf{A}](q) - 2 \cdot \varepsilon \\ \rho \cdot q & \text{otherwise} \end{cases}$$

Then, the Player-A reset strategy $\mathsf{s}_{\varepsilon,\mathsf{Rst}} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{A}}$ is defined by, for all $\rho \in \Omega^+_{\mathcal{C}}$:

$$\mathsf{s}_{\varepsilon,\mathsf{Rst}}(\rho) := \mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\rho)}(\mathsf{Pl}(\rho))$$

[2] Recall that by definition of a concurrent game, a payoff function is bounded and measurable.

It is defined arbitrarily on finite paths not in $\Omega_{\mathcal{C}}^{+}$.

This definition ensures the lemma below:

Lemma 3.2. *For all positive $\varepsilon > 0$, the Player-A reset strategy $\mathsf{s}_{\varepsilon,\mathsf{Rst}}$ is subgame 2ε-optimal.*

Proof sketch. Consider some $\rho \in (Q_{\mathsf{ns}})^{+}$. Whenever there is an update after ρ (i.e. when $\mathsf{U}_{\varepsilon}(\rho \cdot \pi) = \rho \cdot \pi$ for $\pi \in \Omega_{\mathcal{C}}^{+}$), it means that the value of the strategy at that history $\rho \cdot \pi$ is less than the value of the history $\rho \cdot \pi$ w.r.t. Player A minus 2ε. Furthermore, when there is an update at $\rho \cdot \pi$, what the strategy $\mathsf{s}_{\varepsilon,\mathsf{Rst}}$ plays at $\rho \cdot \pi$ comes from a strategy that is ε-optimal at $\rho \cdot \pi$. Hence, with the update, the value of the residual strategy has increased by at least ε. From this observation, we can in fact deduce that almost-surely there are only finitely many updates (since $\varepsilon > 0$). Furthermore, when there is no more updates, the residual strategy is at least 2ε-optimal. We can then conclude by realizing that, from any finite path $\rho \cdot \pi \in \Omega_{\mathcal{C}}^{+}$, the expected value of the finite paths after which there is no more update is, roughly, the value of the residual strategy at $\rho \cdot \pi$. □

The formal proof of Lemma 3.2 is very heavy on notations due to the use of residual strategies and is quite technical. Hence we do not give this proof in the main sections of this chapter. A complete and detailed proof can be found in Appendix 3.6.1.

We can now proceed to the proof of Theorem 3.1.

Proof. Consider some Player $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$. Consider a collection $(S_{\mathsf{C}}^{q})_{q \in Q}$ of GF-strategies that supremizes the game $\mathcal{G}$ w.r.t. Player C. Then, by Theorem 2.3, for all $\rho \in \Omega_{\mathcal{C}}^{+}$, there is a Player-$\mathsf{C}$ strategy $\mathsf{s}_{\varepsilon,\rho} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ generated by $(S_{\mathsf{C}}^{q})_{q \in Q}$ that is ε-optimal from ρ. Considering such strategies, the Player-C reset strategy $\mathsf{s}_{\varepsilon,\mathsf{Rst}}$ defined from $\mathsf{s}_{\varepsilon,\rho}$ is also generated by the collection $(S_{\mathsf{C}}^{q})_{q \in Q}$. Lemma 3.2 ensures that it is subgame 2ε-optimal. □

3.1.2 . Application of Theorem 3.1

In this subsection, we present an application of Theorem 3.1 with PI (recall, prefix-independent) upward well-founded payoff functions (we will define this notion below). First, we consider the probability of PI Borel sets in stochastic trees. As stated in [59, Theorem 5], we have the adaptation below of Levy's 0-1 Law to the context of stochastic trees:

Theorem 3.3 (Levy's 0-1 Law for PI Borel sets in stochastic trees)**.** *Consider a stochastic tree $\mathcal{T} = \langle Q, \mathbb{P} \rangle$ and a Borel set $W \in \mathsf{Borel}(Q)$ that is prefix-independent. Then, from all finite paths $\pi \in Q^{+}$, the sets W and $W_{\lim 1}^{\pi} := \{\rho \in Q^{\omega} \mid \lim_{n \to \infty} \mathbb{P}_{\pi \cdot \rho_{\leq n}}(W) = 1\}$ are equal up to a null set. That is, for all $\pi \in Q^{+}$, we have $\mathbb{P}_{\pi}[W \cap W_{\lim 1}^{\pi}] = \mathbb{P}_{\pi}[W] = \mathbb{P}_{\pi}[W_{\lim 1}^{\pi}]$.*

In particular, this theorem above implies the lemma below:

Lemma 3.4. *Consider a stochastic tree $\mathcal{T} = \langle Q, \mathbb{P} \rangle$ and a Borel set $W \in \mathsf{Borel}(Q)$ that is prefix-independent. We have:*

$$\inf_{\rho \in Q^+} \mathbb{P}_\rho[W] > 0 \Leftrightarrow \inf_{\rho \in Q^+} \mathbb{P}_\rho[W] = 1$$

Proof. Assume that $\inf_{\rho \in Q^+} \mathbb{P}_\rho[W] > 0$. Consider the set $W_{\lim 0} := \{\rho \in Q^\omega \mid \lim_{n\to\infty} \mathbb{P}_{\rho_{\leq n}}(W) = 0\}$. Clearly, $W_{\lim 0} = \emptyset$. By Theorem 3.3, it follows that for all $\rho \in Q^+$, we have $\mathbb{P}_\rho[Q^\omega \setminus W] = \mathbb{P}_\rho[W_{\lim 0}] = 0$. That is, for all $\rho \in Q^+$, we have $\mathbb{P}_\rho[W] = 1$. □

Note that Lemma 3.4 also comes from Lemma 2 in [49]. In the same paper [49], the author studies win/lose PI objectives in standard concurrent games (with standard finite local interactions). They show [49, Theorem 1] that if a state[3] has value less than 1, then the infimum of the values of all states is 0. Equivalently, if the infimum of the states values is positive, then all states have value 1.

In [58, Theorem 3.2] in the context of (standard) turn-based finite (recall with finitely many states, and finitely many actions at each state) PI games, the authors have improved this result: if the infimum of the state values is positive, then Player A has an almost-surely winning strategy from every state. That is, not only all states have value 1, but also there is a strategy that achieves this value from every state. Interestingly, to prove this result, the authors have built an almost-surely winning strategy with reset strategies — similarly to what we presented in the previous subsection. We will discuss further other results proved in that paper [58] in Section 3.3.

In [41], we have adapted (almost verbatim) the reset-strategies-arguments from [58] to obtain an analogous result in finite standard concurrent games [41, Theorem 3]. In fact, we have even realized that the almost-surely winning strategy built in [58] was subgame almost-surely winning.

All these results can be generalized to arbitrary concurrent arenas with more general PI payoff functions than win/lose ones. Specifically, this holds for well-founded[4] payoff functions (upward or downward, depending on the player considered).

Definition 3.5 (Well-founded payoff functions)**.** *Consider a set of colors K and a payoff function $f : \mathsf{K}^\omega \to [0,1]$ and let $E := f[\mathsf{K}^\omega]$. The payoff function f is* upward well-founded *if there is no infinite ascending chain in E. That is, there is no sequence $(x_n)_{n\in\mathbb{N}} \in E^\mathbb{N}$ such that $x_n < x_{n+1}$ for all $n \in \mathbb{N}$. Symmetrically, the payoff function f is* downward well-founded *if there is no infinite descending chain in E.*

---

[3]We need not consider all finite sequences of states since the objective considered is prefix-independent.

[4]The term well-founded comes from its use on order relations.

In particular, any payoff function taking finitely many values is upward well-founded. In fact, well-foundedness can be alternatively defined as follows.

Lemma 3.5 (Proof 3.6.2). *Consider a set of colors* K *and a payoff function* $f : \mathsf{K}^\omega \to [0, 1]$ *and let* $E := f[\mathsf{K}^\omega]$. *It is upward (resp. downward) well-founded if and only if, for all* $x \in [0, 1]$, *there is some* $\varepsilon > 0$ *such that* $[x - \varepsilon, x) \cap E = \emptyset$ *(resp.* $(x, x + \varepsilon] \cap E = \emptyset$*).*

We state below the more general version of all the results discussed above in the context of arbitrary concurrent arenas and PI upward well-founded payoff function. This result could be proved with reset strategies, but there is no need as it is in fact a straightforward corollary of Theorem 3.1 (and Theorem 3.3).

Corollary 3.6. *Consider a concurrent game* $\mathcal{G}$ *where all stopping states have value 1 and with a PI upward well-founded payoff function* f. *Let* $E := f[\mathsf{K}^\omega] \subseteq [0, 1]$ *and* $c := \inf_{q \in Q} \chi_{\mathcal{G}}[\mathsf{A}](q)$. *Let*

$$d := \begin{cases} \inf E \cap [c, 1] & \text{if } E \cap [c, 1] \neq \emptyset \\ c & \text{otherwise} \end{cases}$$

Then, Player A *has a subgame almost-surely winning strategy w.r.t. the objective* $\{f \geq d\} = \{\rho \in \mathsf{K}^\omega \mid f(\rho) \geq d\} \in \mathsf{Borel}(\mathsf{K})$.

This is symmetrical for Player B *(upward is replaced by downward).*

We state a simpler version with the context of a win/lose objective.

Corollary 3.7. *Consider a concurrent game* $\mathcal{G}$ *where all stopping states have value 1 with a PI win/lose function* f. *If* $\inf_{q \in Q} \chi_{\mathcal{G}}[\mathsf{A}](q) > 0$, *then Player* A *has a subgame almost-surely winning strategy in* $\mathcal{G}$.

Proof. Since f is upward well-founded and by Lemma 3.5, there is some $\varepsilon > 0$ such that $[c - \varepsilon, c) \cap E = \emptyset$. Consider a Player-$\mathsf{A}$ subgame $\varepsilon/2$-optimal strategy $\mathsf{s_A}$, whose existence is ensured by Theorem 3.1. Let us show that this strategy is subgame almost-surely winning w.r.t. the winning objective $W_d := \{f \geq d\} \in \mathsf{Borel}(\mathsf{K})$ and $X_d := (\mathsf{col}^\omega)^{-1}[W_q] \subseteq (Q_{\mathsf{ns}})^\omega$. Consider any Player-$\mathsf{B}$ strategy $\mathsf{s_B}$ and $\rho \in (Q_{\mathsf{ns}})^+$. We let $W \subseteq Q^\omega$ be a winning objective for Player A such that, for all $\rho \in Q^\omega$

- If $\rho \in Q^* \cdot (Q_{\mathsf{ns}})^\omega$, then $\rho \in W$ if and only if a suffix of ρ is in X_d;
- Otherwise, $\rho \in W$ if and only if there is some $q \in Q_{\mathsf{s}}$ such that $\rho \in Q^* \cdot q^\omega$.

Note that the set W is prefix-independent. Since, in the stochastic tree $\mathcal{T}_{\mathcal{C}}^{\mathsf{s_A},\mathsf{s_B}}$ all stopping states are self-looping, we have $\mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}[W] = \mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}[X_d \cup (Q_{\mathsf{ns}})^* \cdot Q_{\mathsf{s}}]$.

Furthermore, by choice of the strategy $\mathsf{s_A}$, we have $\mathbb{E}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}[f_{\mathcal{C}}] \geq c - \frac{\varepsilon}{2}$. In addition, for all $\rho \in \mathsf{K}^\omega$, by definition of d, if $f(\rho) < d$, then $f(\rho) \leq c - \varepsilon$.

Hence:

$$\begin{aligned}
c - \varepsilon/2 \leq \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[f_{\mathcal{C}}] &= \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[f_{\mathcal{C}} \cdot \mathbb{1}_{X_d \cup Q^* \cdot Q_{\mathsf{s}}}] + \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[f_{\mathcal{C}} \cdot \mathbb{1}_{(Q_{\mathsf{ns}})^\omega \setminus X_d}] \\
&\leq \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[X_d \cup Q^* \cdot Q_{\mathsf{s}}] + (c - \varepsilon) \cdot (1 - \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[X_d \cup Q^* \cdot Q_{\mathsf{s}}]) \\
&= \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[W] + (c - \varepsilon) \cdot (1 - \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[W])
\end{aligned}$$

We obtain:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[W] \geq \frac{\varepsilon}{2 \cdot (1 - c + \varepsilon)} > 0$$

This holds for all $\rho \in Q^+$. Hence, by Lemma 3.4, we have that for all $\rho \in Q^+$, $1 = \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[W] = \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[X_d \cup Q^* \cdot Q_{\mathsf{s}}]$.

If f is win/lose, then for all $d > 0$, we have $\{f \geq d\} = \{f = 1\}$. □

We conclude this section by providing an example where Corollary 3.6 fails for a PI payoff function that is not upward well-founded.

Example 3.1. *Consider a game $\mathcal{G} = \langle \mathcal{C}, f \rangle$ where $\mathcal{C}$ is a turn-based deterministic arena on the set of colors $\mathsf{K} := \{0, 1\}$ with two states q_0 and q_1 that are colored with 0 and 1 respectively. Player A plays alone and decides at each step to which state she wants to go. The payoff function f maps each infinite sequences of 0 and 1 to the superior limit of the mean of the values seen, except if it is 1, in that case it maps it to 0. More formally, for all $\rho \in \mathsf{K}^\omega$, we have:*

$$f(\rho) := \begin{cases} \limsup_n(\frac{1}{n+1}\sum_{i=0}^n \rho_n) & \text{if } \limsup_n(\frac{1}{n+1}\sum_{i=0}^n \rho_n) < 1 \\ 0 & \text{otherwise} \end{cases}$$

Note that this payoff function is prefix-independent and not upward well-founded. Clearly, the value of both states q_0 and q_1 is 1, since Player A can ensure that the superior limit is as close as desired to 1 and yet less than 1. However, she has no almost-surely winning strategy since no infinite path has value 1 in this deterministic turn-based arena.

3.2 Subgame optimal strategies in arbitrary concurrent games

The remainder of this chapter, that is this section and the two following ones, is based on [41], where we focus on subgame optimal strategies. However, note that whereas in [41] we considered only PI win/lose objectives, in this section we generalize these results to PI payoff functions in arbitrary arenas as we did in Subsection 3.1.2. More specifically, in Subsection 3.2.1, we discuss a simple (and well-known) example where there is no optimal (subgame) strategy which we will use later to justify the conditions considered for subgame optimality. Second, in Subsection 3.2.2, we establish a sufficient condition for a strategy to guarantee a valuation of the states, which turns out to also be

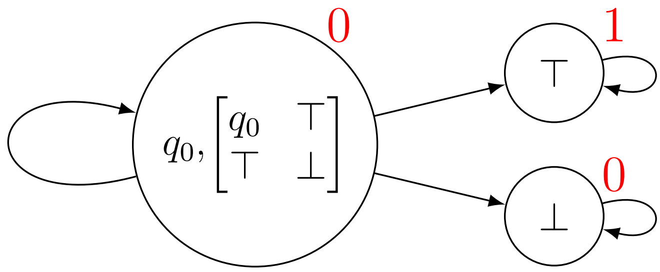

Figure 3.1: A deterministic standard concurrent reachability game $\mathcal{G} = \langle \mathcal{C}, \mathsf{Reach} \rangle$ where Player A wants to reach the target $\{\top\}$.

necessary if the valuation considered is the value of the game (for Player A). Finally, in Subsection 3.2.3, we use the previous result to exhibit a necessary and sufficient condition for the existence of subgame optimal strategies in games where the PI payoff function is upward well-founded.

3.2.1 . A simple game without optimal strategies

In this subsection, we focus on the reachability game of Definition 3.6. Note that we have already considered that game in Chapter 2 to exemplify the construction used to prove Theorem 2.3. This game is well-known, for instance called the snow-ball game in [60]. In the remainder of this dissertation, we will also refer to this game as the snow-ball game.

Definition 3.6 (Game described in Figure 3.1)**.** *Consider the game depicted in Figure 3.1. This game $\mathcal{G} = \langle \mathcal{C}, \mathsf{Reach} \rangle$ is standard. There is only one non-trivial state: q_0. (Alternatively, $\top$ could be a stopping state of value 1 and $\bot$ could be a stopping state of value 0.) The set of colors considered is $\mathsf{K} := \{0, 1\}$ and the colors of the states $q_0, \top, \bot$ are given in red next to them: $\mathsf{col}(q_0) := 0$, $\mathsf{col}(\bot) := 0$ and $\mathsf{col}(\top) := 1$. This game is win/lose, and the objective Reach is a reachability objective (recall Definition 1.25): Player A wins if and only if the state $\top$ is reached. The Player-A set of actions at state q_0 is $\mathsf{Act}_\mathsf{A}^{q_0} := \{a_1, a_2\}$ where a_1 refers to the top row and a_2 refers to the bottom row and similarly we have $\mathsf{Act}_\mathsf{B}^{q_0} := \{b_1, b_2\}$ where b_1 refers to the leftmost column and b_2 refers to the rightmost column.*

This game ensures the following properties.

Proposition 3.8. *The reachability game from Definition 3.6 is such that:*

- *the state q_0 has value 1: $\chi_\mathcal{G}(q_0) = 1$;*
- *for all positive $\varepsilon > 0$, the Player-A positional strategy $\mathsf{s}_\mathsf{A}^\varepsilon$ such that $\mathsf{s}_\mathsf{A}^\varepsilon(q_0)(a_1) := 1-\varepsilon$ and $\mathsf{s}_\mathsf{A}^\varepsilon(q_0)(a_2) := \varepsilon$ has value $1-\varepsilon$: $\chi_\mathcal{G}[\mathsf{s}_\mathsf{A}^\varepsilon](q_0) = 1-\varepsilon$;*
- *no Player-A strategy is optimal (i.e. has value 1).*

The proof of this proposition is not complicated, however it is much easier with the help of the results we will show in the next subsection. Hence, we prove it at the end of that subsection in Page 133.

Important remark: What we show in the next subsection below is quite straightforward to prove given the results we have shown in Subsection 2.3.1 in Chapter 2. However, note that it is central in this dissertation as we will often use the theorems of this subsection to show that a strategy we have defined is optimal, almost-optimal, subgame optimal (or that it is not).

3.2.2 . Sufficient condition for a strategy to guarantee a valuation

In this subsection, we present a pair of conditions sufficient for a Player-A strategy to guarantee a valuation, formally stated as Theorem 3.12. Furthermore, when the valuation is equal to the value of the game for Player A (i.e. for the Player-A strategy to be subgame optimal), this pair of conditions turns out to be also necessary.

Consider a game with a PI payoff function on an arbitrary concurrent arena (in particular, it may not be valuable) and a valuation of finite paths $v : \Omega_{\mathcal{C}}^+ \to [0,1]$. Recall Definition 3.1: $\Omega_{\mathcal{C}}^+$ refers to the non-empty sequences of states that stop once a stopping state is reached.

We explain informally the ideas behind Theorem 3.12 below for the case of the valuation $v := \chi_{\mathcal{G}}[\mathsf{A}]$. The first condition is local: it specifies how a Player-A strategy s_A should behave at each local interaction of the game. First, one can realize that after a history of non-stopping states $\rho \in (Q_{\mathsf{ns}})^+$, the Player-$\mathsf{A}$ value of the game in normal form $\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle$ is equal to the Player-A value after history ρ. Note that this holds even with a payoff function that is not PI. We state this formally below in Proposition 3.9.

Proposition 3.9 (Proof 3.6.4)**.** *Consider an arbitrary concurrent game $\mathcal{G}$. For all $\rho \in (Q_{\mathsf{ns}})^+$, we have $\chi_{\mathcal{G}}[\mathsf{A}](\rho) = \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_{\mathcal{G}}[\mathsf{A}]^\rho \rangle](\mathsf{A})$.*

This suggests that, for all finite sequences of non-stopping states $\rho \in (Q_{\mathsf{ns}})^+$, the GF-strategy $\mathsf{s}_\mathsf{A}(\rho)$ needs to be optimal in the game in normal form $\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle$ for the residual strategy $\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho)}$ to be optimal from ρ_{lt}. When considering this property with arbitrary valuations v, strategies ensuring that property are said to be dominating the valuation v. When $v := \chi_{\mathcal{G}}[\mathsf{A}]$, such strategies are called locally optimal. Note that a definition of strategies dominating valuations was given in Chapter 2 in Definition 2.6 suited for the proof of Theorem 2.3[5]. We give below a new definition of strategies dominating valuations that coincides with Definition 2.6 when $\mathsf{gd} = \mathsf{A}$.

[5]Specifically, Definition 2.6 made use of the notion of guards.

Definition 3.7 (Dominating a valuation, Locally optimal strategies). *Consider an arbitrary concurrent game $\mathcal{G}$ and a valuation $v : \Omega_{\mathcal{C}}^+ \to [0,1]$. A Player-$\mathsf{A}$ strategy $\mathsf{s_A}$* dominates *the valuation v if, for all $\rho \in (Q_{\mathsf{ns}})^+$, the GF-strategy $\mathsf{s_A}(\rho)$ is such that $\mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle](\mathsf{s_A}(\rho)) \geq v(\rho)$. When the valuation $v = \chi_{\mathcal{G}}[\mathsf{A}]$, the strategy $\mathsf{s_A}$ is said to be* locally optimal. *The definition is symmetrical for Player B.*

Although dominating a valuation in general is not necessary for guaranteeing it (recall Definition 3.2), it turns out that being locally optimal is a necessary condition for being subgame optimal, as stated below.

Lemma 3.10 (Proof 3.6.4). *In an arbitrary concurrent game $\mathcal{G}$, for all Player-A strategies $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ and $\rho \in Q^+$, we have $\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s_A}](\rho_{\mathsf{lt}}) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_{\mathcal{G}}[\mathsf{s_A}]^\rho \rangle](\mathsf{s_A}(\rho)) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_{\mathcal{G}}[\mathsf{A}]^\rho \rangle](\mathsf{s_A}(\rho))$. As a corollary, if $\mathsf{s_A}$ is subgame optimal, then for all $\rho \in (Q_{\mathsf{ns}})^+$, we have $\chi_{\mathcal{G}}[\mathsf{A}](\rho) = \chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s_A}](\rho_{\mathsf{lt}})$. Therefore, the strategy $\mathsf{s_A}$ is also locally optimal.*

Dominating a valuation v does not ensure guaranteeing v. However, it does ensure nice properties. Indeed, the simple yet crucial remark we can make is that given a Player-A strategy $\mathsf{s_A}$ dominating a valuation v, for all Player-B strategies $\mathsf{s_B}$, the valuation v is non-decreasing (recall Definition 2.3) in the stochastic tree induced by $\mathsf{s_A}$ and $\mathsf{s_B}$. However, in Definition 2.3 of non-decreasing valuation in stochastic trees, the valuations considered are of the type $Q^+ \to [0,1]$ instead of $\Omega_{\mathcal{C}}^+ \to [0,1]$. Hence, we define below a canonical way to transform valuations before stating Lemma 3.11.

Definition 3.8 (Canonical transformation of valuations). *Consider an arbitrary concurrent game $\mathcal{G}$ and a valuation $v : \Omega_{\mathcal{C}}^+ \to [0,1]$. We denote by $v_{\mathsf{s}} : Q^+ \to [0,1]$ the valuation such that, for all $\rho \in Q^+$, denoting $\pi_\rho \in \Omega_{\mathcal{C}}^+$ the longest prefix of ρ in $\Omega_{\mathcal{C}}^+$ (which is therefore equal to ρ if $\rho \in \Omega_{\mathcal{C}}^+$):*

$$v_{\mathsf{s}}(\rho) := v(\pi_\rho)$$

Lemma 3.11 (Proof 3.6.5). *Consider an arbitrary concurrent game $\mathcal{G}$, a valuation $v : \Omega_{\mathcal{C}}^+ \to [0,1]$ and a Player-A strategy $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ dominating the valuation v. For all Player-B strategies $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$, in the stochastic tree $\mathcal{T}_{\mathcal{C}}^{\mathsf{s_A},\mathsf{s_B}}$, for all $\pi \in \Omega_{\mathcal{C}}^+$, the valuation $(v_{\mathsf{s}})^\pi$ is non-decreasing from π. This is symmetrical for Player B.*

By Proposition 2.9, we have that given a Player-A strategy dominating a valuation v and any Player-B strategy, in the stochastic tree induced by both strategies, the value $v(\rho)$ of any finite paths $\rho \in \Omega_{\mathcal{C}}^+$ is less than or equal to the expected value of $\mathsf{limsup}_{v_{\mathsf{s}}} : Q^\omega \to [0,1]$ (recall Proposition 2.4) from ρ.

To obtain that $\mathsf{s_A}$ guarantees the valuation v, it would then suffice that, for any Player-B strategy and after any finite paths $\rho \in (Q_{\mathsf{ns}})^+$, almost-surely the superior limit $\mathsf{limsup}_{v_{\mathsf{s}}}$ is less than or equal to f. This constitutes the second condition that, along with dominating v, is sufficient for the strategy $\mathsf{s_A}$ to guarantee the valuation v. This is stated in Theorem 3.12 below.

Interestingly, when $v = \chi_{\mathcal{G}}[\mathsf{A}]$, this second condition is also a necessary condition for subgame optimality. Indeed, assume that $\mathsf{s_A}$ does not ensure this condition. That is, there is a Player-B strategy and a finite path $\rho \in (Q_{\mathsf{ns}})^+$ after which there is some $r \in \mathbb{Q} \cap [0,1]$ and $\delta > 0$ for which there is a positive probability that $f \leq r$ and $r + \delta \leq \mathsf{limsup}_{v_\mathsf{s}}$. We denote this event $E_{r,\delta} \subseteq Q^\omega$. Since f and $\mathsf{limsup}_{v_\mathsf{s}}$ are both prefix-independent, it follows that $E_{r,\delta}$ is prefix-independent. Hence, by Lemma 3.4 (from the previous section) and by definition of the superior limit function $\mathsf{limsup}_{v_\mathsf{s}}$, we can show that there are continuations of this path ρ ending at states of value at least $r + \delta/2$ for which the probability of $E_{r,\delta}$ is arbitrarily close to 1. In particular, there is probability arbitrarily close to 1 that $f \leq r$. It follows that we can exhibit a path $\rho' \in \rho \cdot (Q_{\mathsf{ns}})^+$ such that, from ρ' the value of the Player-A strategy $\mathsf{s_A}$ is less than the value of ρ' (w.r.t. $\chi_{\mathcal{G}}[\mathsf{A}]$).

We obtain the theorem below.

Theorem 3.12. *Consider an arbitrary PI concurrent game $\mathcal{G} = \langle \mathcal{C}, f \rangle$, a valuation $v : \Omega_{\mathcal{C}}^+ \to [0,1]$ and a Player-A strategy $\mathsf{s_A} \in \mathsf{S}_{\mathcal{C}}^{\mathsf{A}}$. Assume that the strategy $\mathsf{s_A}$ satisfies the pair of conditions below:*

- *$\mathsf{s_A}$ dominates the valuation v;*
- *for all $\rho \in (Q_{\mathsf{ns}})^+$ and Player-B strategies $\mathsf{s_B} \in \mathsf{S}_{\mathcal{C}}^{\mathsf{B}}$, we have: $\mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}[\mathsf{limsup}_{v_\mathsf{s}} \leq f_{\mathcal{C}}] = 1$.*

Then the Player-A strategy $\mathsf{s_A}$ guarantees the valuation v.

Conversely, if $\mathsf{s_A}$ guarantees $\chi_{\mathcal{G}}[\mathsf{A}]$ (i.e. $\mathsf{s_A}$ is subgame optimal), then the strategy $\mathsf{s_A}$ satisfies that pair of conditions for $v = \chi_{\mathcal{G}}[\mathsf{A}]$.

Proof. First, note that since f is PI, then for all $\rho \in (Q_{\mathsf{ns}})^+$, we have $(f_{\mathcal{C}})^\rho = f_{\mathcal{C}}$. Now, assume that the Player-A strategy $\mathsf{s_A}$ satisfies that pair of conditions. Let $\rho \in \Omega_{\mathcal{C}}^+$. If $\rho_{\mathsf{lt}} \in Q_{\mathsf{s}}$, then clearly the strategy $\mathsf{s_A}$ is optimal from ρ. Assume now that $\rho \in (Q_{\mathsf{ns}})^+$ and consider a Player-B strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$. By Lemma 3.11, the valuation $(v_\mathsf{s})^\rho : Q^* \to [0,1]$ is non-decreasing from ρ in the stochastic tree $\mathcal{T}_{\mathcal{C}}^{\mathsf{s_A},\mathsf{s_B}}$. Therefore, Proposition 2.9 ensures that $v(\rho) \leq \mathbb{E}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}[\mathsf{limsup}_{v_\mathsf{s}}]$. Furthermore, by assumption $\mathbb{E}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}[\mathsf{limsup}_{v_\mathsf{s}}] \leq \mathbb{E}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}[f_{\mathcal{C}}]$. Hence, we obtain $v(\rho) \leq \mathbb{E}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}[f_{\mathcal{C}}]$. As this holds for all Player-B strategies $\mathsf{s_B}$, it follows that $\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho)}](\rho_{\mathsf{lt}}) \geq v(\rho)$. In fact, the Player-$\mathsf{A}$ strategy $\mathsf{s_A}$ guarantees the valuation v.

Assume now that $v = \chi_{\mathcal{G}}[\mathsf{A}]$ and that the Player-A strategy $\mathsf{s_A}$ guarantees the valuation v, i.e. that the Player-A strategy $\mathsf{s_A}$ is subgame optimal. Lemma 3.10 ensures that this strategy must be locally optimal. Assume now that there is some $\rho \in (Q_{\mathsf{ns}})^+$ and Player-B strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$ such that $\mathbb{P}^{\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho)},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}[\mathsf{limsup}_{v_\mathsf{s}} \leq f_\mathcal{C}] < 1$. We denote the stochastic tree $\mathcal{T}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}$ simply by $\mathcal{T}$, and the corresponding probability function by $\mathbb{P}$. There is some $p, \delta \in \mathbb{Q} \cap [0,1]$ with $\delta > 0$ such that $\mathbb{P}[E_{p,\delta}] > 0$. where $E_{p,\delta} := \{f_\mathcal{C} \leq p \cap p + \delta \leq \mathsf{limsup}_{v_\mathsf{s}}\} \subseteq Q^\omega$. However, both $f_\mathcal{C}$ and $\mathsf{limsup}_{v_\mathsf{s}}$ are not PI since stopping states are taken into account. Hence, we let $E'_{p,\delta} := \{\pi \cdot \rho \in Q^\omega \mid \pi \in Q^*, \rho \in (Q_\mathsf{ns})^\omega,\ f \circ \mathsf{col}^\omega(\rho) \leq p \cap p+\delta \leq \mathsf{limsup}_v(\rho)\}$. This event is measurable and PI and has the same measure as the event $E_{p,\delta}$ since, once a stopping state is reached it is never left. Hence: $\mathbb{P}[E'_{p,\delta}] > 0$.

Consider now some $x, y \in [0,1]$ such that $\frac{1-p-\frac{\delta}{2}}{1-p} < x < y < 1$. By Lemma 3.4, we have that there is some $\pi \in (Q_\mathsf{ns})^+$ such that $\mathbb{P}_\pi[E'_{p,\delta}] \geq y$. We let $\mathsf{Always}_{\leq p+\delta/2} := \{\theta \in (Q_\mathsf{ns})^\omega \mid \forall i \in \mathbb{N}, v^\rho(\pi \cdot \theta_{\leq i}) \leq p + \delta/2\}$. We have $\mathbb{P}_\pi[E'_{p,\delta} \cap \mathsf{Always}_{\leq p+\delta/2}] = 0$ by definition of the function limsup_v. Therefore, we have:

$$\mathbb{P}_\pi[E'_{p,\delta}] \leq \sup_{\substack{\theta \in Q^* \\ v^\rho(\pi\cdot\theta) \geq p+\delta/2}} \mathbb{P}_{\pi\cdot\theta}[E'_{p,\delta}]$$

Hence, there is some $\theta \in Q^*$ such that $v^\rho(\pi \cdot \theta) \geq p + \delta/2$ and $\mathbb{P}_{\pi\cdot\theta}[E_{p,\delta}] \geq x$. In particular, $\mathbb{P}_{\pi\cdot\theta}[f_\mathcal{C} \leq p] \geq x$. Then, we have:

$$\begin{aligned}\mathbb{E}_{\pi\cdot\theta}[f_\mathcal{C}] &\leq \mathbb{P}_{\pi\cdot\theta}[f_\mathcal{C} \leq p] \cdot p + 1 - \mathbb{P}_{\pi\cdot\theta}[f_\mathcal{C} \leq p] \leq x \cdot p + 1 - x \\ &< p + \frac{\delta}{2} \leq v(\rho \cdot \pi \cdot \theta)\end{aligned}$$

Hence, the strategy s_A is not optimal from $\rho \cdot \pi \cdot \theta$, it is therefore not subgame optimal. □

Special cases. We have given in several articles weaker versions of Theorem 3.12, for instance if we assume that s_A is positional. We would like to recall some of these versions since they will be useful in the following.

First, in [41], we considered the case where the valuation $v : Q \to [0,1]$ values the states, not the finite sequences of states, with Q finite. This implies that the valuation v takes only finitely many values. Consider a Player-A strategy s_A dominating this valuation v — straightforwardly extended to finite sequences of states by considering the last element of the sequence — and any Player-B strategy s_B. By Lemma 3.11, this valuation v is non-decreasing in the stochastic tree induced by these strategies. Hence, Proposition 2.6 gives that almost-surely all infinite paths have a limit w.r.t. the valuation v. In this context this implies that almost-surely, all states seen infinitely often have the same value [41, Lemma 2]. This is stated formally below in Corollary 3.13, after the definition of valuations considering the last state of a sequence.

Definition 3.9. *Consider an arbitrary concurrent game $\mathcal{G}$ and a valuation of the states $v : Q \to [0,1]$. This valuation is extended into $v_\mathsf{lt} : Q^+ \to [0,1]$*

(it can also be seen as $v_{\mathsf{lt}} : \Omega^+_{\mathcal{C}} \to [0,1]$) such that, for all $\rho \in Q^+$, we have $v_{\mathsf{lt}}(\rho) := v(\rho_{\mathsf{lt}})$.

Corollary 3.13. *Consider an arbitrary concurrent game $\mathcal{G}$ and a valuation $v : Q \to [0,1]$ taking finitely many values. For all $u \in v[Q]$ we let $Q^v_u := v^{-1}[\{u\}]$. Consider a Player-A strategy dominating the valuation $v_{\mathsf{lt}} : \Omega^+_{\mathcal{C}} \to [0,1]$ and any Player-B strategy $\mathsf{s_B}$. For all $\rho \in (Q_{\mathsf{ns}})^+$, we have:*

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[\bigcup_{u \in v[Q]} (Q^* \cdot (Q^v_u)^\omega)] = 1$$

Proof. Consider any $q \in Q$ and Player-B strategy $\mathsf{s_B} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$. By Lemma 3.11, the valuation $v_{\mathsf{lt}} = ((v_{\mathsf{lt}})_{\mathsf{s}})^\rho : Q^+ \to [0,1]$ is non-decreasing from ρ in the stochastic tree $\mathcal{T}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}}$. Hence, by Lemma 2.6, almost-surely from ρ the superior and inferior limit of v_{lt} are equal. Since this valuation v_{lt} takes only finitely many values, it follows that almost-surely, the game settles in a unique value slice Q^v_u for some $u \in v[Q]$. □

In this context, Theorem 3.12 amounts to the corollary below, and corresponds to [41, Theorem 1]:

Corollary 3.14. *Consider an arbitrary PI concurrent game $\mathcal{G} = \langle \mathcal{C}, f \rangle$ and a valuation $v : Q \to [0,1]$ taking finitely many values. For all $u \in v[Q]$ we let $Q^v_u := v^{-1}[\{u\}]$. Let $\mathsf{s_A} \in \mathsf{S}^{\mathsf{A}}_{\mathcal{C}}$ be a Player-A strategy. Assume that the strategy $\mathsf{s_A}$ satisfies the pair of conditions below:*

- *$\mathsf{s_A}$ dominates the valuation v_{lt};*
- *for all $\rho \in (Q_{\mathsf{ns}})^+$ and Player-B strategies $\mathsf{s_B} \in \mathsf{S}^{\mathsf{B}}_{\mathcal{C}}$, for all $u \in v[Q]$, we have*
$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[Q^* \cdot (Q^v_u)^\omega \cap \{f \geq u\}] = \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[Q^* \cdot (Q^v_u)^\omega]$$

Then the Player-A strategy $\mathsf{s_A}$ guarantees the valuation v.

Conversely, if $\mathsf{s_A}$ guarantees $\chi_{\mathcal{G}}[\mathsf{A}]$ (i.e. $\mathsf{s_A}$ is subgame optimal), then it also satisfies that pair of conditions for $v = \chi_{\mathcal{G}}[\mathsf{A}]$.

Proof. Proving this corollary only amounts to proving that the second condition is equivalent to the second condition of Theorem 3.12. Consider some $\rho \in (Q_{\mathsf{ns}})^+$ and a Player-B strategy $\mathsf{s_B} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$. By Corollary 3.13, assuming the condition of this corollary, we have:

$$\begin{aligned}
\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[\mathsf{limsup}_{(v_{\mathsf{lt}})_{\mathsf{s}}} \leq f] &= \sum_{u \in v[Q]} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[\{\mathsf{limsup}_{(v_{\mathsf{lt}})_{\mathsf{s}}} \leq f\} \cap Q^* \cdot (Q^v_u)^\omega] \\
&= \sum_{u \in v[Q]} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[\{u \leq f\} \cap Q^* \cdot (Q^v_u)^\omega] \\
&= \sum_{u \in v[Q]} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[Q^* \cdot (Q^v_u)^\omega] = 1
\end{aligned}$$

Reciprocally, assuming that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[\mathsf{limsup}_{(v_{\mathsf{lt}})_\mathsf{s}} \leq f] = 1$, for all $u \in v[Q]$ such that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[Q^* \cdot (Q^v_u)^\omega] > 0$, it must be that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[\{f \geq u\} \mid Q^* \cdot (Q^v_u)^\omega] = 1$ since, as already used above, $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[\{\mathsf{limsup}_{(v_{\mathsf{lt}})_\mathsf{s}} = u\} \mid Q^* \cdot (Q^v_u)^\omega] = 1$. □

Let us consider a further special case in the case where a positional Player-A strategy $\mathsf{s_A}$ dominates a valuation of the states taking finitely many values. In that case, as above almost-surely all states seen infinitely often have the same values w.r.t. v. In addition, states $q \in Q$ where the local value of the strategy is more than the value of the state, that is such that $v(q) < \mathsf{val}[\langle \mathsf{F}(q), v\rangle][\mathsf{s_A}(q)]$, are seen only finitely often. Furthermore, in the special case of a finite standard arena, in all end components, all states have the same values. This is stated below in Corollary 3.15. It is a slight generalization of [39, Proposition 18].

Corollary 3.15. *Consider an arbitrary arena $\mathcal{C}$, a valuation of the states $v : Q \to [0,1]$ taking finitely many values and a positional Player-A strategy $\mathsf{s_A} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{A}}$ dominating the valuation $v_{\mathsf{lt}} : \Omega^+_{\mathcal{C}} \to [0,1]$. For all Player-$\mathsf{B}$ strategies and from all finite paths $\pi \in \Omega^+_{\mathcal{C}}$, almost-surely, the set of states seen infinitely often is included in $\{q \in Q \mid v(q) = \mathsf{val}[\langle \mathsf{F}(q), v\rangle][\mathsf{s_A}(q)]\}$.*

If we assume additionally that the arena $\mathcal{C}$ is standard and finite, then in the MDP $\Gamma^{\mathsf{s_A}}_{\mathcal{C}}$, for all end components $H \in \mathcal{E}_{\Gamma^{\mathsf{s_A}}_{\mathcal{C}}}$, there is a value $u(v,H) \in [0,1]$ such that $v[Q_H] = \{u(v,H)\}$.

Proof. For all $q \in Q$, we let $d_q := \mathsf{val}[\langle \mathsf{F}(q), v\rangle][\mathsf{s_A}(q)] - v(q)$. Since the strategy $\mathsf{s_A}$ dominates the valuation v_{lt}, for all $q \in Q$, we have $d_q \geq 0$. Consider some state $q \in Q$ such that $d_q > 0$ and let $u := v(q) \in [0,1)$. We let $Q^v_{>u} := \{q \in Q \mid v(q) > u\}$, $Q^v_{\leq u} := Q \setminus Q^v_{>u}$ and $Q^v_u := \{q \in Q \mid v(q) = u\}$. For any Player-$\mathsf{B}$ GF-strategy $\sigma_{\mathsf{B}} \in \Sigma^q_{\mathsf{B}}$, recalling Definition 1.28, we have:

$$\begin{aligned} u + d_q \leq \mathsf{out}[\langle \mathsf{F}(q), v\rangle](\mathsf{s_A}(q), \sigma_{\mathsf{B}}) &= \sum_{q' \in Q} \varrho_q(\mathsf{s_A}(q), \sigma_{\mathsf{B}})(q') \cdot v(q') \\ &\leq \sum_{q' \in Q^v_{>u}} \varrho_q(\mathsf{s_A}(q), \sigma_{\mathsf{B}})(q') + \sum_{q' \in Q^v_{\leq u}} \varrho_q(\mathsf{s_A}(q), \sigma_{\mathsf{B}})(q') \cdot u \\ &= \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[Q^v_{>u}] \cdot (1-u) + u \end{aligned}$$

Hence, $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[Q^v_{>u}] \geq \frac{d_q}{1-u} \geq d_q > 0$. This holds for all Player-B GF-strategies $\sigma_{\mathsf{B}} \in \Sigma^q_{\mathsf{B}}$. Hence, from any finite path $\pi \in \Omega^+_{\mathcal{C}}$, for any Player-$\mathsf{B}$ strategy with the Player-A strategy $\mathsf{s_A}$, if the state q is seen infinitely often, almost surely, the set $Q^v_{>u}$ is seen infinitely often almost-surely. Furthermore, by Corollary 3.13, almost-surely all states visited infinitely often have the same values w.r.t. v. That is, for any Player-B strategy, almost-surely the state q is seen only finitely often.

Assume now that the arena $\mathcal{C}$ is standard and consider an EC $H \in \mathcal{E}_{\Gamma^{\mathsf{s_A}}_{\mathcal{C}}}$. Since it is an end component, all states in it may be seen infinitely often with probability 1 for a Player-B strategy playing at each state $q \in Q_H$ uniformly over all actions in $\beta_H(q)$. Hence, since Corollary 3.13 ensures that for any Player-B strategy almost-surely all states seen infinitely have the same value, it follows that all states in H have the same value. □

Finally, we have below what Theorem 3.12 amounts to in the context of finite standard concurrent game for a win/lose objective. We already proved this result in the context of reachability games [39, Proposition 17] and generalized it to more general objectives (but still not all PI objectives) in [40, Lemma 16].

Corollary 3.16. *Consider a finite standard concurrent PI win/lose game $\mathcal{G}$ and a valuation $v : Q \to [0,1]$. Let $\mathsf{s_A} \in \mathsf{S}^{\mathsf{A}}_{\mathcal{C}}$ be a positional Player-A strategy. Assume that the strategy $\mathsf{s_A}$ satisfies the pair of conditions below:*

- *it dominates the valuation v;*
- *for all end components H in the MDP induced by the strategy $\mathsf{s_A}$, if $u(v,H) > 0$, then for all $q \in Q_H$, we have $\chi_{\mathcal{C}_H^{\mathsf{s_A}}}(q) = 1$.*

Then the Player-A strategy $\mathsf{s_A}$ guarantees the valuation v.

Conversely, if $\mathsf{s_A}$ guarantees $\chi_{\mathcal{G}}[\mathsf{A}]$ (i.e. $\mathsf{s_A}$ is subgame optimal), then it also satisfies that pair of conditions for $v = \chi_{\mathcal{G}}[\mathsf{A}]$.

Proof. As for the proof of Corollary 3.14 which used Theorem 3.12, we prove this corollary by showing that the second condition is equivalent to the second condition of Corollary 3.14. Assume that the above conditions hold. By Theorem 2.3, since all local interactions in MDPs are supremized by deterministic GF-strategies, almost-optimal strategies for Player B against the strategy $\mathsf{s_A}$ can be found among deterministic strategies. Furthermore, Lemma 1.17 ensures that for all deterministic Player-B strategies, the game almost-surely settles in an EC. By assumption, it follows that if the game settles in a value slice of positive value, then almost-surely Player A wins, which implies the second condition of Corollary 3.14.

Assume now that the second condition of Corollary 3.14. For any EC H in the MDP induced by the strategy $\mathsf{s_A}$ such that $u(v,H) > 0$, it must be against all Player-B (deterministic) strategies compatible with that EC, the game has value at least $u(v,H) > 0$ from any state in Q_H. Since the game is win/lose, this implies that for all $q \in Q_H$, we have $\chi_{\mathcal{C}_H^{\mathsf{s_A}}}(q) = 1$. □

Proof of Proposition 3.8. With the help of the results proved in this subsection, let us show Proposition 3.8.

Proof. Consider some positive $\varepsilon > 0$. Let $v_\varepsilon : Q \to [0,1]$ be such that $v_\varepsilon(q_0) := 1 - \varepsilon$, $v_\varepsilon(\top) := 1$ and $v_\varepsilon(\bot) := 0$. The strategy $\mathsf{s}^\varepsilon_\mathsf{A}$ dominates this valuation. Furthermore, the only end components compatible with this strategy are $\{\top\}$ — the target of value 1 — and $\{\bot\}$ — of value 0. Hence, the strategy $\mathsf{s_A}$

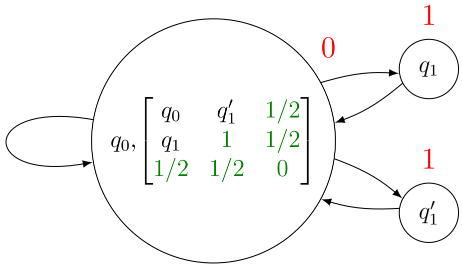

Figure 3.2: A co-Büchi game.

satisfies the conditions of Corollary 3.16 and therefore guarantees the valuation v_ε. Since this holds for all positive $\varepsilon > 0$, it follows that the value of the state q_0 is 1.

Consider now a Player-A strategy $\mathsf{s_A} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{A}}$. Let us show that its value is less than 1. There are two cases. First, if, for all $n \geq 1$, we have $\mathsf{s_A}(q_0^n)(a_1) = 1$, then for a Player-B strategy $\mathsf{s_B} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$ that plays positionally action b_1 with probability 1, then with $\mathsf{s_A}$ and $\mathsf{s_B}$ surely the game will loop on q_0. Otherwise, consider the least $n_0 \geq 1$ such that $\mathsf{s_A}(q_0^{n_0})(a_2) > 0$. Consider then a Player-B strategy $\mathsf{s_B} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$ such that, for all $n \geq 1$:

$$\mathsf{s_B}(q_0^n) := \begin{cases} \{b_1 \mapsto 1,\ b_2 \mapsto 0\} & \text{if } n < n_0 \\ \{b_1 \mapsto 0,\ b_2 \mapsto 1\} & \text{otherwise} \end{cases}$$

With both strategies $\mathsf{s_A}, \mathsf{s_B}$, we have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[q_0^{n_0}] = 1$$

and

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[q_0^{n_0} \cdot \bot] = \mathsf{s_A}(q_0^{n_0})(a_2) > 0$$

Hence, the value of the strategy $\mathsf{s_A}$ is less than 1. In fact, Player A has no optimal strategy in this game. □

3.2.3 . Necessary and sufficient condition for the existence of subgame optimal strategies

In the previous subsection, we have (in particular) studied a necessary and sufficient pair of conditions for a Player-A strategy to be subgame optimal. In this subsection, we focus on the existence of subgame optimal strategies in arbitrary finite-state games. This section is an adaptation of [41, Section 6] to the case of arbitrary local interactions.

In [58, Theorem 4.5], the authors have proved a transfer result in PI win/lose turn-based games: the amount of memory sufficient to play optimally at every state of value 1 of every game is also sufficient to play optimally in every game. This result does not hold in concurrent games as is. First, although there are always optimal strategies in PI turn-based games (as proved in the same paper [58, Theorem 4.3]), there are PI concurrent games without optimal strategies, as discussed above in Proposition 3.8. Second, although almost-surely winning strategies can be found among positional strategies in standard concurrent co-Büchi games (we will discuss it further in Chapter 5), infinite memory may be required to play optimally in co-Büchi standard concurrent games. This is witnessed by the game of Figure 3.2. Note that this game is very close to the co-Büchi game of Figure 2.5[6]. The difference with the game of Figure 2.5 is that there are two states q_1 and q_1' instead of only one state q_1. In that way, Player A can now know when Player B has played action b. Hence, the infinite-memory Player-A strategy described in the proof of Lemma 2.23 can be translated in this setting to obtain an optimal strategy (i.e. a strategy of value $\frac{1}{2}$). Let us recall quickly how this strategy plays. To play optimally, Player A may play the top row with probability $1 - \varepsilon_k$ and the middle row with probability ε_k for $\varepsilon_k > 0$ that goes (fast) to 0 when k goes to ∞ (where k denotes the number of steps). The ε_k is chosen so that, if Player B always plays the left column with probability 1, then the state q_1 is seen finitely often with probability 1. Furthermore, as soon as the state q_1' is visited, Player A switches to a positional strategy playing the bottom row with probability ε_k' small enough (where k denotes the number of steps before the state q_1' was seen) and the two top rows with probability $(1 - \varepsilon_k')/2$.

Therefore, the transfer of memory from almost-surely winning to optimal does not hold in concurrent games even if it is assumed that optimal strategies exist. However, one can note that although the strategy described above is optimal, it is not subgame optimal. Indeed, when the strategy switches, the value of the residual strategy is $1/2 - \varepsilon_k' < 1/2$. In fact, there is no subgame optimal strategy in that game. Actually, if we assume that, not only optimal but subgame optimal strategies exist, then the transfer of memory will hold.

The aim of this subsection is twofold: first, we identify a necessary and sufficient condition for the existence of subgame optimal strategies[7]. Second, we establish the above-mentioned memory transfer that relates the amount of memory to play subgame optimally and to be subgame almost-surely winning. Furthermore, this is done with any PI upward well-founded payoff function. Note that although we generalize some of the results from [58] — that we have discussed above — the method we use here is different from what the authors of [58] did to prove the transfer of memory in turn-based games. Namely, they showed that there is a live and self-consistent permutation of the distribution

[6]Recall, in the game of Figure 2.5, Player A has an optimal strategy among action-strategies but none among classical strategies (the ones we consider in this dissertation). A description of this game in provided in Definition 2.19.

[7]Note that this is different from what we did in the previous section: there, we established a necessary and sufficient condition for a specific strategy to be subgame optimal. Here, given a game, we consider necessary and sufficient conditions on the game for the existence of a subgame optimal strategy.

over states that both players can agree on. They can then play according to this permutation (which can be used to specify their preferences over these distributions over states).

Before formally stating the main theorem of this section, we first need to define the winning objectives that can be obtained from payoff functions. This is done below.

Definition 3.10 (Winning objective obtained from a payoff function)**.** *For all sets of colors* K*, payoff functions* $f : \mathsf{K}^\omega \to [0,1]$ *and* $u \in [0,1]$*, the set* $\{f \geq u\} := \{\rho \in \mathsf{K}^\omega \mid f(\rho) \geq u\}$ *is a winning objective* obtained *from* f*.*

To establish this transfer of memory, we will actually modify the game forms occurring in the game. Specifically, we will keep all Player-B GF-strategies while disregarding any Player-A GF-strategy that is not optimal in a specific game in normal form. We define this change of game forms below, we will illustrate it later on a standard game form.

Definition 3.11 (Only optimal GF-strategies in Game forms)**.** *Consider a set of outcomes* O*, a game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ *on that set of outcomes and a valuation* $v : \mathsf{O} \to [0,1]$ *such that* $\mathsf{Opt}_\mathsf{A}(\langle \mathcal{F}, v \rangle) \neq \emptyset$*. We let* $\mathsf{Opt}(\mathcal{F}, v) \in \mathsf{Form}(\mathsf{O})$ *be the game form defined by* $\mathsf{Opt}(\mathcal{F}, v) := \langle \mathsf{Opt}_\mathsf{A}(\langle \mathcal{F}, v \rangle), \Sigma_\mathsf{B}, \mathsf{O}, \varrho \rangle$*.*

Given a set of outcomes O *and any set of game forms* $E \subseteq \mathsf{Form}(\mathsf{O})$*, we let* $\mathsf{Opt}(E) \subseteq \mathsf{Form}(\mathsf{O})$ *denote the set of game forms* $\mathsf{Opt}(E) := \{\mathsf{Opt}(\mathcal{F}, v) \mid \mathcal{F} \in E,\ v : \mathsf{O} \to [0,1],\ \mathsf{Opt}_\mathsf{A}(\langle \mathcal{F}, v \rangle) \neq \emptyset\}$*. Note that* $\mathsf{Opt}(E)$ *is not empty as soon as* $E \neq \emptyset$ *since, for all* $\mathcal{F} \in E$*, we have* $\mathsf{Opt}_\mathsf{A}(\langle \mathcal{F}, v \rangle) \neq \emptyset$ *for all constant valuations* $v : \mathsf{O} \to [0,1]$*.*

Given a set of game forms E and a memory skeleton M, we now introduce below the definition of (E, M)-subgame almost-surely winnable payoff functions, i.e. payoff functions for which, for all win/lose objectives that can be obtained from them, in all games built on E, subgame almost-surely winning strategies can be found among strategies that can be implemented with M.

Definition 3.12 ((E, M)-subgame almost-surely winnable objective)**.** *Consider a non-empty finite set of colors* K*, a PI payoff function* $f : \mathsf{K}^\omega \to [0,1]$ *and a memory skeleton* $\mathsf{M} = \langle M, m_{\mathsf{init}}, \mu \rangle$ *on* K*. The payoff function* f *is said to be* (E, M)*-subgame almost-surely winnable (*(E, M)*-SAW for short) if the following holds: for all* $u \in [0,1]$*, in all finite-state concurrent games* $\mathcal{G} = \langle \mathcal{C}, \{f \geq u\} \rangle$ *built on* E *where there is a subgame almost-surely winning strategy, there is one that is* M*-implementable. If* $|M| = 1$*, then the payoff function* f *is said to be* E*-positionally subgame almost-surely winnable (*E*-PSAW for short).*

We can now state the main theorem of this section. Recall the notation $V_\mathsf{A}^\mathcal{G}$ and Q_u^A from Definition 1.32.

Theorem 3.17. *Consider an arbitrary finite-state concurrent game* $\mathcal{G} = \langle \mathcal{C}, f \rangle$ *with a PI upward well-founded payoff function* $f : \mathsf{K}^\omega \to [0, 1]$*. The five following assertions are equivalent:*

a. there exists a Player-A *subgame optimal strategy;*

b. there exists a locally optimal Player-A *strategy that is optimal;*

c. for all positive $\varepsilon > 0$*, there exists a locally optimal Player-*A *strategy that is* ε*-optimal;*

d. there exists a locally optimal Player-A *strategy* $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ *such that, for all* $u \in V_\mathsf{A}^\mathcal{G}$ *and* $q \in Q_u^\mathsf{A}$*, we have* $\chi_{\langle \mathcal{C}, \{f \geq u\} \rangle}[\mathsf{s}_\mathsf{A}](q) > 0$*;*

e. there exists a locally optimal Player-A *strategy* $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ *such that, for all* $u \in V_\mathsf{A}^\mathcal{G}$ *and* $q \in Q_u^\mathsf{A}$*, we have* $\chi_{\langle \mathcal{C}_{\mathsf{Exit}(Q_u^\mathsf{A})}, \{f \geq u\} \rangle}[\mathsf{s}_\mathsf{A}](q) > 0$*.*

where $\mathcal{C}_{\mathsf{Exit}(Q_u^\mathsf{A})}$ *corresponds to the arena* $\mathcal{C}$ *where all states outside of* Q_u^A *are stopping states of value 1.*

If this holds and if, for some finite memory skeleton M*, the payoff function* f *is* $(\mathsf{Opt}(\{\mathsf{F}(q) \mid q \in Q\}), \mathsf{M})$*-SAW, then there exists a subgame optimal* M*-implementable strategy.*

First, note that the equivalence is stated in terms of existence of strategies, not on the strategies themselves. In particular, any subgame optimal strategy is both optimal and locally optimal, however, an optimal strategy that is locally optimal is not necessarily a subgame optimal strategy. An example is provided in Appendix 3.6.3. We would also like to point out that in the arena $\mathcal{C}$ with a win/lose objective, e.g. $\{f \geq u\}$ for some $u \in V_\mathsf{A}^\mathcal{G}$, the stopping states are still taken into account. That is, if a state $q \in Q_\mathsf{s}$ is reached, the game stops and the value $\mathsf{val}(q)$ occurs.

Second, we would like to highlight what we believe is an important take-away from this theorem. Beside the memory transfer, this theorem tells at which condition there is a subgame optimal strategy. Although items b., c., d. and e. are different, they have the same generic form: there is an assumption that locally optimal strategies satisfy a specific property w.r.t. the objective. This specific property obviously matters for the equivalence to hold, however we would like to focus on the local optimality assumption. What this theorem suggests is that, the reason why, in concurrent games, there does not always exist (subgame) optimal strategies is that if one only considers locally optimal strategies, then the value of the game may drop. For instance, in the snow-ball game of Definition 3.6, the value of the state q_0 is 1, but if one only considers locally optimal strategies (i.e. strategies that always play the top row with probability 1), then the value of that state becomes 0. This can also be witnessed in the co-Büchi game of Figure 3.2 that we discussed above. There is an optimal strategy in this game, but there is no subgame optimal ones. If one only considers locally optimal strategies (i.e. that plays the bottom row with probability 0), then the value of the game is 0. In other words, for a strategy to be optimal it must, after some history, switch to a sub-optimal strategy. Interestingly, as we will discuss in the next section, this cannot occur in standard finite turn-based games, which explains why subgame optimal strategies always exist[8] in that setting.

Third, it is straightforward that item a. implies item b. (from Lemma 3.10) and that item b. implies item c.. It is also straightforward that item d. implies item e.. However, the implication item c. implies item d. is less direct and uses the well-foundedness assumption. Note that it is only for this implication that the well-foundedness assumption is used. Let us formally prove this implication below.

Lemma 3.18. *Consider an arbitrary finite-state concurrent game $\mathcal{G} = \langle \mathcal{C}, f \rangle$ with a PI upward well-founded payoff function $f : \mathsf{K}^\omega \to [0,1]$. If for all positive $\varepsilon > 0$, there is a Player-A ε-optimal strategy that is locally optimal, then, for all $u \in V_{\mathsf{A}}^{\mathcal{G}}$ and $q \in Q_u^{\mathsf{A}}$, there is a Player-A locally optimal strategy $\mathsf{s_A}$ such that $\chi_{\langle \mathcal{C}, \{f \geq u\}\rangle}[\mathsf{s_A}](q) > 0$.*

Proof. Consider some $u \in V_{\mathsf{A}}^{\mathcal{G}}$ and $q \in Q_u^{\mathsf{A}}$. Because the payoff function f is upward well-founded, there is some $0 < \delta \leq u$ such that $[u - \delta, u) \cap f[\mathsf{K}^\omega] = \emptyset$. Therefore, we have for all $\rho \in \mathsf{K}^\omega$, $f(\rho) \geq u$ if and only if $f(\rho) \geq u - \delta$. Consider any Player-A strategy $\mathsf{s_A}$ that is ε-optimal for some $0 < \varepsilon < \delta$. Assume towards a contradiction that this Player-A strategy is such that $\chi_{\langle \mathcal{C}, \{f \geq u\}\rangle}[\mathsf{s_A}](q) = 0$. By definition of δ, we also have $\chi_{\langle \mathcal{C}, \{f \geq u - \delta\}\rangle}[\mathsf{s_A}](q) = 0$. Consider a Player-$\mathsf{B}$ strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$ such that, for $x := \frac{\delta - \varepsilon}{2} > 0$, we have:

$$\sum_{q \in Q_{\mathsf{s}}} \mathbb{P}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[(Q_{\mathsf{ns}})^* \cdot \{q\}] \cdot \mathsf{val}(q) + \mathbb{P}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[\{f \geq u - \delta\} \cap (Q_{\mathsf{ns}})^\omega] \leq x$$

Hence, we have:

$$\begin{aligned}
\mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[f_{\mathcal{C}}] &= \sum_{q \in Q_{\mathsf{s}}} \mathbb{P}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[(Q_{\mathsf{ns}})^* \cdot \{q\}] \cdot \mathsf{val}(q) + \mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[f_{\mathcal{C}} \cdot \mathbb{1}_{(Q_{\mathsf{ns}})^\omega}] \\
&\leq x - \mathbb{P}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[\{f \geq u - \delta\} \cap (Q_{\mathsf{ns}})^\omega] + \mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[f_{\mathcal{C}} \cdot \mathbb{1}_{(Q_{\mathsf{ns}})^\omega}] \\
&\leq x - \mathbb{P}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[\{f \geq u - \delta\} \cap (Q_{\mathsf{ns}})^\omega] + \mathbb{P}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[\{f \geq u - \delta\} \cap (Q_{\mathsf{ns}})^\omega] \\
&+ \mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[f \cdot \mathbb{1}_{\{f < u - \delta\} \cap (Q_{\mathsf{ns}})^\omega}] \\
&\leq x + u - \delta < u - \varepsilon
\end{aligned}$$

This is in contradiction with the fact that the Player-A strategy $\mathsf{s_A}$ is ε-optimal from the state q of Player-A value $\chi_{\mathcal{G}}[\mathsf{A}](q) = u$. □

[8]Assuming the payoff function is PI upward well-founded.

In the remainder of this subsection, we explain the constructions leading to the proof of Theorem 3.17, i.e. to the proof that item e. implies item a. The transfer of memory is a direct consequence of the way this theorem is proven. We fix an arbitrary finite-state concurrent game $\mathcal{G} = \langle \mathcal{C}, f \rangle$ with a PI payoff function f for the remainder of the subsection — the upward well-foundedness assumption is not used for the implication item e. implies item a.

The idea is as follows. Recall that, given any locally optimal Player-A strategy, almost-surely the game settles in a value slice Q_u^{A} for some $u \in V_{\mathsf{A}}^{\mathcal{G}}$ as stated in Corollary 3.13. Furthermore, as stated in Corollary 3.14 in finite-state arenas, subgame optimal strategies are exactly the strategies that are locally optimal and such that, for all Player-B strategies almost-surely, the value w.r.t. f of infinite paths is at least the value of the value slice in which the game settles. Our idea is therefore to consider, for all $u \in V_{\mathsf{A}}^{\mathcal{G}} \setminus \{0\}$, subgame almost-surely winning strategies in a derived game $\mathcal{G}_u := \langle \mathcal{C}_u, \{f \geq u\} \rangle$ with $\mathcal{C}_u$ a "restriction" of the arena $\mathcal{C}$ to Q_u. We can then glue together these subgame almost-surely winning strategies – defined for all $u \in V_{\mathsf{A}}^{\mathcal{G}} \setminus \{0\}$ – into a subgame optimal strategy. However, there are some issues:

1) there must exist a subgame almost-surely winning strategy in $\mathcal{G}_u$;

2) this subgame almost-surely winning strategy in $\mathcal{G}_u$ should be locally optimal when considered in the whole game $\mathcal{G}$.

Let us first deal with issue 2). Let $u \in V_{\mathsf{A}}^{\mathcal{G}}$. One can ensure that the almost-surely winning strategies in the game $\mathcal{G}_u$ are all locally optimal in $\mathcal{G}$ by properly defining the arena $\mathcal{C}_u$. More specifically, this is done by enforcing that the only Player-A possible strategies in $\mathcal{C}_u$ are locally optimal in the game $\mathcal{G}$. To do so, we construct the arena $\mathcal{C}_u$ such that its set of states with non-trivial interaction is Q_u and the local interaction at state $q \in Q_u$ is equal to $\mathsf{Opt}(\mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}])$ (recall Definition 3.11).

We illustrate this construction on a standard finite game form: a part of a concurrent game is depicted in Figure 3.3 and the change of the interaction of the players at state q_0 is depicted in Figures 3.4, 3.5, 3.6 and 3.7.

Furthermore, since we want from all the states the existence of subgame almost-surely winning strategies in $\mathcal{G}_u$ — recall issue 1) — we will build the game $\mathcal{G}_u$ such that any edge leading to a state not in Q_u in $\mathcal{G}$ now leads to a stopping state of value 1.

Definition 3.13 (Game $\mathcal{G}_u$)**.** *Consider a positive value $u \in V_{\mathsf{A}}^{\mathcal{G}} \setminus \{0\}$. We define the game $\mathcal{G}_u = \langle \mathcal{C}_u, \{f \geq u\} \rangle$ with $\mathcal{C}_u = \langle Q, \mathsf{F}^{\mathsf{Opt}}, \mathsf{K}, \mathsf{col} \rangle$ with:*

- *all states $q \in Q \setminus Q_u^{\mathsf{A}}$ are stopping states of value 1: $\mathsf{val}(q) \leftarrow 1$;*

- *The values of all stopping states in $Q_{\mathsf{s}} \cap Q_u^{\mathsf{A}}$ — whose values in $\mathcal{G}$ are all u since they are in Q_u — are changed to 1;*

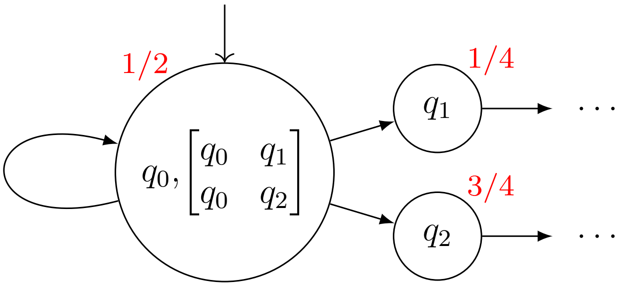

Figure 3.3: A part of a deterministic standard concurrent game $\mathcal{G}$ with $\mathsf{Act}_{\mathsf{A}}^{q_0} = \{a_1, a_2\}$. The values are depicted in red near the states.

$$\begin{matrix} a_1 \\ a_2 \end{matrix} \begin{bmatrix} q_0 & q_1 \\ q_0 & q_2 \end{bmatrix}$$

Figure 3.4: The local interaction $\mathsf{F}(q_0)$ at state q_0 in the game of Figure 3.3.

$$\begin{matrix} a_1 \\ a_2 \end{matrix} \begin{bmatrix} \frac{1}{2} & \frac{1}{4} \\ \frac{1}{2} & \frac{3}{4} \end{bmatrix}$$

Figure 3.5: The game in normal form $\langle \mathsf{F}(q_0), \chi_{\mathcal{G}} \rangle$ from the game $\mathcal{G}$ of Figure 3.3.

$$\begin{matrix} \frac{a_1+a_2}{2} \\ a_2 \end{matrix} \begin{bmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{3}{4} \end{bmatrix}$$

Figure 3.6: The game in normal form from Figure 3.5 with only optimal strategies available for Player A.

$$\begin{matrix} \frac{a_1+a_2}{2} \\ a_2 \end{matrix} \begin{bmatrix} q_0 & \frac{q_1+q_2}{2} \\ q_0 & q_2 \end{bmatrix}$$

Figure 3.7: The game form obtained from the game form of Figure 3.4 with only the optimal strategies from Figure 3.6.

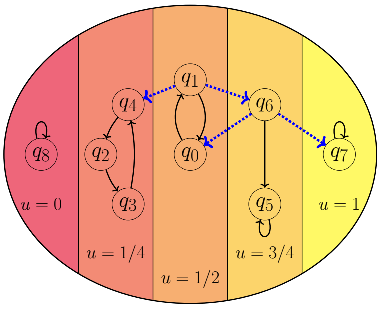

Figure 3.8: The depiction of a PI concurrent game with its value slices.

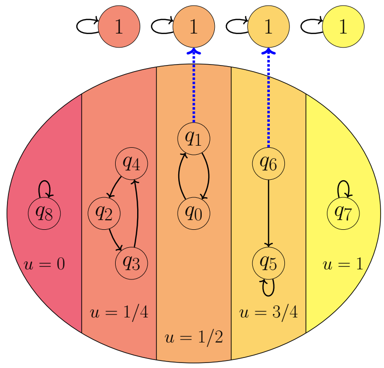

Figure 3.9: The PI concurrent game after the modifications of Definition 3.13.

- *For all* $q \in Q_{\mathsf{ns}} \cap Q_u^{\mathsf{A}}$, *we set* $\mathsf{F}^{\mathsf{Opt}}(q) := \mathsf{Opt}(\mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}])$.

An illustration of this construction can be found in Figures 3.8 and 3.9. The blue dotted arrows are the ones that need to be redirected when the game is changed. With such a definition, we have made some progress w.r.t. the issue 1) cited previously (regarding the existence of subgame almost-surely winning strategies): the values of all states of the game $\mathcal{G}_u$ are positive (for positive u).

Lemma 3.19. *Consider the game* $\mathcal{G}_u$ *for some positive* $u \in V_{\mathsf{A}}^{\mathcal{G}} \setminus \{0\}$ *and assume that, in* $\mathcal{G}$*, there exists a strategy that is locally optimal such that, for all* $q \in Q_u^{\mathsf{A}}$*, we have* $\chi_{\langle \mathcal{C}, \{f \geq u\} \cup \mathsf{Exit}(Q_u^{\mathsf{A}}) \rangle}[\mathsf{s_A}](q) > 0$*. Then, for all states* q *in* $\mathcal{G}_u$ *we have* $\chi_{\mathcal{G}_u}[\mathsf{A}](q) > 0$.

Proof Sketch. Consider a state $q \in Q_u^{\mathsf{A}}$ and a Player-A locally optimal strategy $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ in $\mathcal{G}$ such that $\chi_{\langle \mathcal{C}, \{f \geq u\} \cup \mathsf{Exit}(Q_u^{\mathsf{A}}) \rangle}[\mathsf{s_A}](q) > 0$. Then, the strategy $\mathsf{s_A}$ (restricted to $(Q_u^{\mathsf{A}})^+$) can be seen as a strategy in $\mathcal{G}_u$. Note that this is only possible because the strategy $\mathsf{s_A}$ is locally optimal (due to the definition of $\mathcal{G}_u$).

Consider a Player-B strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}_u}$. This strategy can be seen as a strategy in $\mathcal{C}$, assuming it is defined arbitrarily once the game has exited Q_u^{A}. Since if the play never reaches a stopping states (of value 1, since all stopping states in $\mathcal{G}_u$ have value 1), what happens in $\mathcal{G}_u$ and $\mathcal{G}$ is identical, it follows that $\mathbb{P}_{\mathcal{C}_u,q}^{\mathsf{s_A},\mathsf{s_B}}[\{f \geq u\} \cup (Q_{\mathsf{ns}})^* \cdot Q_{\mathsf{s}}] = \mathbb{P}_{\mathcal{C}_{\mathsf{Exit}(Q_u^{\mathsf{A}})},q}^{\mathsf{s_A},\mathsf{s_B}}[\{f \geq u\}] > 0$. Thus, the value of the state q is positive in $\mathcal{G}_u$. □

In fact, Lemma 3.19 suffices to deal with issue 1). Indeed, as stated in Corollary 3.7 in the previous subsection, it is a general result that in a finite-state PI win/lose concurrent game where all stopping states have value 1, if all states have positive values, then there is a subgame almost-surely winning strategy.

However, there is an additional difficulty when considering the transfer of memory. Consider some finite memory skeleton $\mathsf{M} = \langle M, m_{\text{init}}, \mu \rangle$ and assume that the payoff function f is $(\mathsf{Opt}(\{\mathsf{F}(q) \mid q \in Q\}), \mathsf{M}) - \mathsf{SAW}$. Subgame almost-surely winning strategies in $\mathcal{G}_u$ can be found among M-implementable strategies. However, when we will glue these pieces of strategies together below, we will need for the strategies to be subgame almost-surely winning strategies in $\mathcal{G}_u$ regardless of their starting memory state. This is due to the fact that, since the colors will be seen in the whole game, not only in a specific value slice, then the memory state, when entering the value slice Q_u^{A} for the last time, may be different to the initial memory state m_{init}. We first introduce below in Definition 3.14 the set of memory states reachable from a starting memory state given the set of colors that could occur.

Definition 3.14. *Consider a finite memory skeleton* $\mathsf{M} = \langle M, m_{\text{init}}, \mu \rangle$ *on the set of colors* K. *For any finite subset of colors* $\mathsf{K}' \subseteq \mathsf{K}$, *we let* $\mathsf{Reach}(\mathsf{M}, \mathsf{K}') := \{m \in M \mid \exists \rho \in (\mathsf{K}')^*,\ \mu^*(m_{\text{init}}, \rho) = m\} \subseteq M$ *be the set of memory states of* M *reachable from* m_{init} *with finite sequences colors in* K'.

For all $m \in M$, *we denote by* M^m *the memory skeleton* $\mathsf{M}^m := \langle M, m, \mu \rangle$.

We now state in Lemma 3.20 below: when subgame almost-surely winning strategies exist in the game $\mathcal{G}_u$, then there is an action map that implements a subgame almost-surely winning strategies regardless of the starting memory state.

Lemma 3.20. *Consider a finite memory skeleton* $\mathsf{M} = \langle M, m_{\text{init}}, \mu \rangle$ *and assume that the payoff function* f *is* $(\mathsf{Opt}(\{\mathsf{F}(q) \mid q \in Q\}), \mathsf{M}) - \mathsf{SAW}$. *Let* $u \in V_{\mathsf{A}}^{\mathcal{G}} \setminus \{0\}$ *and assume that there is subgame almost-surely winning strategy in the game* $\mathcal{G}_u$. *Then, for all finite set of colors* $\mathsf{K}' \subseteq \mathsf{K}$, *there is an action map* $\lambda : M \times Q_u^{\mathsf{A}} \to \cup_{q \in Q_u^{\mathsf{A}}} \Sigma_{\mathsf{A}}^q$ *such that, for all* $m \in \mathsf{Reach}(\mathsf{M}, \mathsf{K}')$, *the strategy implemented by* M^m *and* λ *is subgame almost-surely winning in* $\mathcal{G}_u$.

Proof. For all $m \in \mathsf{Reach}(\mathsf{M}, \mathsf{K}')$, we let $\rho_m \in \mathsf{K}'$ be a finite sequence of colors from m_{init} to m: $\mu^*(m_{\mathsf{init}}, \rho_m) = m$. We modify the game $\mathcal{G}_u$ into a game $\mathcal{G}'_u$ as follows. We add, before actually entering the arena $\mathcal{C}_u$, and for all $m \in \mathsf{Reach}(\mathsf{M}, \mathsf{K}')$, a sequence of states — with trivial local interaction[9] — whose corresponding sequence of colors is equal to ρ_m. Since the payoff function f is PI, adding finitely many colors before entering the game does not change its value. Hence, there is still a subgame almost-surely winning strategy in the game $\mathcal{G}'_u$. Since the payoff function f is $(\mathsf{Opt}(\{\mathsf{F}(q) \mid q \in Q\}), \mathsf{M}) - \mathsf{SAW}$, there is an action map λ that, along with M, implements such a subgame almost-surely winning strategy in $\mathcal{G}'_u$. By definition of the modification of the game $\mathcal{G}_u$ into $\mathcal{G}'_u$, the action map λ ensures that, for all $m \in \mathsf{Reach}(\mathsf{M}, \mathsf{K}')$, the strategy

[9]Note that a trivial interaction can be obtained from any game form by mapping every outcome to the same state.

implemented by M^m and λ is subgame almost-surely winning in $\mathcal{G}_u$. $\square$

We can now glue together pieces of strategies s_A^u defined in all games $\mathcal{G}_u$ into a single strategy $\mathsf{s}_\mathsf{A}[(\mathsf{s}_\mathsf{A}^u)_{u \in V_\mathsf{A}^\mathcal{G} \setminus \{0\}}]$. Informally, the glued strategy mimics the strategy on $(Q_u^\mathsf{A})^+$ and switches strategy when a value slice is left and another one is reached.

Definition 3.15 (Gluing strategies)**.** *For all values $u \in V_\mathsf{A}^\mathcal{G} \setminus \{0\}$, consider a Player-*A *strategy s_A^u in the game $\mathcal{G}_u$. Then, we glue these strategies into the strategy $\mathsf{s}_\mathsf{A}[(\mathsf{s}_\mathsf{A}^u)_{u \in V_\mathsf{A}^\mathcal{G} \setminus \{0\}}] : Q^+ \to \cup_{q \in Q} \Sigma_\mathsf{A}^q$ simply written s_A such that, for all $\rho \in Q^+$:*

$$\mathsf{s}_\mathsf{A}(\rho) := \begin{cases} \mathsf{s}_\mathsf{A}^u(\pi) & \textit{if } u = \chi_\mathcal{G}[\mathsf{A}](\rho_{\mathsf{lt}}) > 0 \textit{ for } \pi \textit{ the longest suffix of } \rho \textit{ in } (Q_u^\mathsf{A})^+ \\ \textit{is arbitrary} & \textit{if } \chi_\mathcal{G}[\mathsf{A}](q) = 0 \end{cases}$$

As stated in Lemma 3.21 below, the construction described in Definition 3.15 transfers almost-surely winning strategies in $\mathcal{G}_u$ into a subgame optimal strategy in $\mathcal{G}$.

Lemma 3.21. *Consider a Player-*A *strategy s_A locally optimal such that, for all $u \in V_\mathsf{A}^\mathcal{G} \setminus \{0\}$, for all $\rho \in Q^+$, we have $\mathsf{s}_\mathsf{A}^\rho \in \mathsf{S}_\mathsf{A}^{\mathcal{C}_u}$ subgame almost-surely winning in $\mathcal{G}_u$. Then, the strategy s_A is subgame optimal in $\mathcal{G}$.*

It is in particular the case for the glued strategy $\mathsf{s}_\mathsf{A}[(\mathsf{s}_\mathsf{A}^u)_{u \in V_\mathcal{G} \setminus \{0\}}]$ as soon as, for all $u \in V_\mathsf{A}^\mathcal{G} \setminus \{0\}$, $\mathsf{s}_\mathsf{A}^u \in \mathsf{S}_\mathsf{A}^{\mathcal{C}_u}$ is a subgame almost-surely winning strategy in $\mathcal{G}_u$.

Proof. We apply Corollary 3.14. The strategy s_A is locally optimal. In addition, if the game eventually settles in a value slice Q_u^A for some $u > 0$, from then on the strategy s_A is almost-surely winning in $\mathcal{G}_u$, whose win/lose objective is $\{f \geq u\}$. This holds for all $u \in V_\mathsf{A}^\mathcal{G} \setminus \{0\}$, so the second condition of Corollary 3.14 holds.

Now, consider for all $u \in V_\mathsf{A}^\mathcal{G} \setminus \{0\}$, a Player-A strategy $\mathsf{s}_\mathsf{A}^u \in \mathsf{S}_\mathsf{A}^{\mathcal{C}_u}$ subgame almost-surely winning strategy in $\mathcal{G}_u$. Let s_A be the glued strategy $\mathsf{s}_\mathsf{A}[(\mathsf{s}_\mathsf{A}^u)_{u \in V_\mathcal{G} \setminus \{0\}}]$. Then, the strategy s_A is locally optimal. Indeed, by Lemma 3.9, for all $q \in Q$, we have $\chi_\mathcal{G}[\mathsf{A}](q) = \mathsf{val}[\langle \mathsf{F}(q), \chi_\mathcal{G}[\mathsf{A}] \rangle]$. Hence, for all $u \in V_\mathsf{A}^\mathcal{G} \setminus \{0\}$, for all states in Q_u^A, by the strategy restriction done to define the game $\mathcal{G}_u$, only optimal GF-strategies are considered at each game in normal form $\mathsf{F}^{\mathsf{Opt}}(q)$ at states $q \in Q_u^\mathsf{A}$. Furthermore, any GF-strategy is optimal in a game in normal form of value 0 (which is the case of the game in normal forms of states in Q_0). In addition, for all $u \in V_\mathsf{A}^\mathcal{G} \setminus \{0\}$ and for all $\rho \in Q^+$, we have $\mathsf{s}_\mathsf{A}^\rho$ and s_A^u that coincide on $\mathcal{C}_u$. Therefore, $\mathsf{s}_\mathsf{A}^\rho$ is subgame almost-surely winning in $\mathcal{G}_u$. $\square$

We now have all the ingredients to prove Theorem 3.17.

Proof. By Lemma 3.10, item a. implies item b., item b. straightforwardly implies item c., item c. implies item d. by Lemma 3.18 and item d. straightforwardly implies item e.

Let us now show that item e. implies item a. By Lemma 3.19, for all positive values $u \in V_{\mathsf{A}}^{\mathcal{G}} \setminus \{0\}$, all states in $\mathcal{G}_u$ have positive values. It follows, by Corollary 3.7, that there exists a subgame almost-surely winning strategy in every game $\mathcal{G}_u$ for $u \in V_{\mathsf{A}}^{\mathcal{G}} \setminus \{0\}$. We then obtain a subgame optimal strategy by gluing these strategies together, given by Lemma 3.21.

Consider now the transfer of memory. Consider a finite arbitrary memory skeleton M and assume that the payoff function f is $(\mathsf{Opt}(\{\mathsf{F}(q) \mid q \in Q\}), \mathsf{M})-$ SAW. Let $\mathsf{K}' := \mathsf{col}[Q_{\mathsf{ns}}] \subseteq \mathsf{K}$ be the finite set of colors appearing in $\mathcal{C}$. Let $u \in V_{\mathsf{A}}^{\mathcal{G}} \setminus \{0\}$. By Lemma 3.20, there exists an action map λ_u such that, for all $m \in \mathsf{Reach}(\mathsf{M}, \mathsf{K}')$, the strategy implemented by the action map λ_u and the memory skeleton M^m is subgame almost-surely winning strategy in $\mathcal{G}_u$. We then define the action map $\lambda : Q \times \mathsf{K} \to \sum_{q \in Q} \Sigma_{\mathsf{A}}^q$ such that, for all $q \in Q$ and $k \in \mathsf{K}$, we have $\lambda(q, k) := \lambda_{\chi_{\mathcal{G}}[\mathsf{A}](q)}(q, k) \in \Sigma_{\mathsf{A}}^q$. Clearly, if we denote by $\mathsf{s_A}$ the strategy implemented by λ and M, the strategy $\mathsf{s_A}$ satisfies the condition of Lemma 3.21, it is therefore subgame optimal in $\mathcal{G}$. □

Finally, we conclude this section by giving a Corollary of Theorem 3.17. Specifically, we consider standard finite game forms, possibly turn-based ones. In fact, such a set of game forms is stable by application of the Opt operator from Definition 3.11. This is formally stated below in Proposition 3.22.

Proposition 3.22. *Consider a set of outcomes* O*. Let* $\mathsf{Std}_f(\mathsf{O})$ *(resp.* $\mathsf{TB}_f(\mathsf{O})$*) denote the set of standard finite game forms (resp. standard finite turn-based game forms) on* O*. Then,* $\mathsf{Opt}(\mathsf{Std}_f(\mathsf{O})) = \mathsf{Std}_f(\mathsf{O})$ *and* $\mathsf{Opt}(\mathsf{TB}_f(\mathsf{O})) = \mathsf{TB}_f(\mathsf{O})$*.*

Proof. Consider any standard finite game form $\mathcal{F} = \langle \Sigma_{\mathsf{A}}, \Sigma_{\mathsf{B}}, \mathsf{O}, \varrho \rangle \in \mathsf{Std}_f(\mathsf{O})$. Let $n := |\mathsf{Act_A}|$ and $k := |\mathsf{Act_B}|$. Consider any valuation of the outcomes $v : \mathsf{O} \to [0, 1]$. Consider the game in normal form $\langle \mathcal{F}, v \rangle$.

There exists a finite set $D_{\mathsf{A}} \subseteq \mathcal{D}(\mathsf{Act_A}) \subseteq \mathsf{Opt_A}(\langle \mathcal{F}, v \rangle)$ of optimal strategies such that the optimal strategies in $\langle \mathcal{F}, v \rangle$ are exactly the convex combinations of strategies in D_{A}. This is a well known result, argued for instance in [61]. The idea is to write a system of finitely many inequalities whose set of solutions is exactly the set of optimal GF-strategies $\mathsf{Opt_A}(\langle \mathcal{F}, v \rangle)$. Consider the set in $\mathbb{R}^n$ of vectors whose sum of components is equal to 1. We can express the set of optimal strategies $\mathsf{Opt_A}(\langle \mathcal{F}, v \rangle)$ as the solution of a system of inequalities. First, with n inequalities we can consider only non-negative values. Furthermore, with another k inequalities – specifying that the weighted sum in each column is at least $u = \mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{A}]$ – we have that the solutions to the system of inequalities are exactly the vectors of values corresponding to the optimal strategies in the game in normal form $\langle \mathcal{F}, v \rangle$. The result then follows from standard system of inequalities arguments as the space of solutions is in fact a polytope. Therefore, $\mathsf{Opt}(\mathcal{F}, v) = \langle D_{\mathsf{A}}, \mathsf{Act}_{\mathsf{B}}, \mathsf{O}, \varrho \rangle_s \in \mathsf{Std}_f(\mathsf{O})$.

If in addition, the game form $\mathcal{F}$ is turn-based, then so is the game form $\mathsf{Opt}(\mathcal{F}, v)$. □

We deduce the corollary below.

Corollary 3.23. *Consider a set of colors* K *and a PI upward well-founded payoff function* $f : \mathsf{K}^\omega \to [0, 1]$. *Assume that there is a memory skeleton* M *such that the payoff function* f *is* $(\mathsf{Std}_f(\mathsf{O}), \mathsf{M})$-$\mathsf{SAW}$ *(resp.* $(\mathsf{TB}_f(\mathsf{O}), \mathsf{M})$-$\mathsf{SAW}$*). Then, in all standard finite concurrent (resp. turn-based) games* $\mathcal{G} = \langle \mathcal{C}, f \rangle$*, subgame optimal strategies, when they exist, can be found among* M-*implementable strategies.*

To obtain a simpler statement, let us write this corollary when f is win/lose, M is of size 1 (i.e. we consider positional strategies).

Corollary 3.24. *Consider a set of colors* K *and a PI objective* $W \subseteq \mathsf{K}^\omega$. *Assume that in all standard finite concurrent games* $\mathcal{G} = \langle \mathcal{C}, W \rangle$*, when there is a subgame almost-surely winning strategy, there is one that is positional. Then, in all standard finite concurrent games* $\mathcal{G} = \langle \mathcal{C}, W \rangle$*, subgame optimal strategies, when they exist, can be found among positional strategies.*

For an application of this corollary, see Proposition 5.8 in Chapter 5.

Finally, we conclude this section by mentioning that the slicing technique — i.e. considering different values slices, and then glue together strategies from different value slices — was already used in the context of concurrent games in [50]. The authors focus on parity objectives and establish a memory transfer result from limit-sure winning (i.e. almost-optimal for the value 1) to almost-optimal strategies. As an application, they show that, for co-Büchi objectives, since positional strategies are sufficient to win limit-surely, then they also are to play almost-optimally. Their construction made heavy use of the specific nature of the parity objectives. Furthermore, the paper contains complexity results, on which we do not focus in this dissertation.

3.3 Subgame optimal strategies in standard games

In this section, we focus on standard finite games. Recall, this means that, in the games we consider, there are finitely many states and at all local interactions, both players have finitely many actions. Furthermore, although we will not make use of this fact, since standard finite game forms are valuable by Lemma 1.14, standard finite games have a value by Theorem 2.3.

3.3.1 . Application to finite turn-based games

The aim of Section 3.2 was to extend an already existing result — from [58] — on turn-based games to the context of arbitrary concurrent games, with PI upward well-founded payoff functions. This required an adaptation of the assumptions. However, it is in fact possible to retrieve the original result on turn-based games proved in [58] from Theorem 3.17 in a fairly straightforward manner. Specifically, in [58], it is shown — among other results — that there are always optimal strategies in finite turn-based games[10] with PI objectives [58, Theorem 4.3] and that the amount of memory sufficient to be almost-surely winning is also sufficient to be optimal [58, Theorem 4.5]. Note that we do not prove exactly the same result since we show the existence of subgame optimal strategies and we transfer the amount of memory from what is sufficient to be subgame almost-surely winning. On another note, other results — unrelated to the questions considered in this chapter — are shown in [58], in particular the authors have provided several algorithmic results, see for instance [58, Theorem 4.4, Theorem 5.1].

We state formally below the existence of subgame optimal strategies in finite turn-based games, the transfer of memory can then be deduced from Corollary 3.23.

Corollary 3.25. *In all finite turn-based games $\mathcal{G} = \langle \mathcal{C}, f \rangle$ with a PI upward well-founded payoff function f, both players have a deterministic subgame optimal strategy.*

The proof of this corollary is actually quite simple by applying Theorem 3.17 and showing that item c. always holds in finite turn-based games. This last part amounts to showing that locally optimal strategies achieve the same values as all strategies in finite turn-based games. This was already noticed in [62, Section 4.1.1][11], and it can be proved straightforwardly by using Theorem 3.1. In addition, we would like to mention that the result for Player A still holds even if we assume that in the local interactions belonging to Player B, she has infinitely many actions. This also works symmetrically for player B.

Proof. We prove the result for Player A, but the proof is similar for Player B. Let $q \in Q$ and consider the finite set of actions $\mathsf{Act}^q_\mathsf{A}$ available to Player A in the game form $\mathsf{F}(q)$. We let

$$\eta_q := \min_{a \in \mathsf{Act}^q_\mathsf{A} \setminus \mathsf{Opt}_\mathsf{A}(\langle \mathsf{F}(q), \chi_\mathcal{G} \rangle)} \chi_\mathcal{G}(q) - \mathsf{out}[\langle \mathsf{F}(q), \chi_\mathcal{G} \rangle][a] > 0$$

be the minimum of how much a sub-optimal action at state q deviates from an optimal action. We let $\eta := \min_{q \in Q} \eta_q > 0$. Then, consider any $0 < \varepsilon < \eta$.

[10] Recall Definition 1.19: we assume that all local interactions are finite

[11] Note that Gimbert is an author of both [58] cited above and [62], though these two papers seem unrelated.

Since the game $\mathcal{G}$ is turn-based, it is supremized by the collection of the sets of deterministic GF-strategies. Hence, by Theorem 3.1, since the game is turn-based, there is a Player-A deterministic strategy $\mathsf{s_A}$ that is subgame ε-optimal. Hence, by Lemma 3.10 and Proposition 3.9, we have $\mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_{\mathcal{G}} \rangle](\mathsf{s_A}(\rho)) \geq \chi_{\mathcal{G}}[\mathsf{s_A}](\rho_{\mathsf{lt}}) \geq \chi_{\mathcal{G}}(\rho_{\mathsf{lt}}) - \varepsilon = \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_{\mathcal{G}} \rangle] - \varepsilon$. Since $\mathsf{s_A}$ is deterministic and by definition η and ε, it must be that $\mathsf{s_A}(\rho) \in \mathsf{Act}_{\mathsf{A}}^{\rho_{\mathsf{lt}}} \cap \mathsf{Opt_A}(\langle \mathsf{F}(q), \chi_{\mathcal{G}} \rangle)$. That is, the strategy $\mathsf{s_A}$ is locally optimal. Hence, item c. of Theorem 3.1 holds, and therefore subgame optimal strategies exist in $\mathcal{G}$ for Player A. □

Finally, we apply Corollary 3.23 to a specific PI upward well-founded function such that each win/lose objective obtained from it is a parity objective (this corresponds to the notion of priority game, studied for instance in [28]). This function will be measurable and $\mathsf{Std}_f(\mathsf{O})$-$\mathsf{PSAW}$. This comes from the fact that in finite turn-based games with parity objectives, there are always positional optimal strategies for both players [27, 28]. Note that this result is already known, see [28, Lemma 9]

Definition 3.16. *Consider a finite set of colors $\mathsf{K} \subseteq \mathbb{N}$ and a map $g : \mathsf{K} \to [0,1]$. We let $f_{\mathsf{Par}}(\mathsf{K}, g) : \mathsf{K}^\omega \to [0,1]$ be such that, for all $\rho \in \mathsf{K}^\omega$, we have:*

$$f_{\mathsf{Par}}(\mathsf{K}, g)(\rho) := g(\max \mathsf{InfOft}(\rho)) \in [0,1]$$

where the notation $\mathsf{InfOft}(\rho)$ was introduced in Definition 1.25 and refers to the set of colors seen infinitely often in ρ.

Proposition 3.26 (Proof 3.6.6)**.** *For all finite sets of colors $\mathsf{K} \subseteq \mathbb{N}$ and maps $g : \mathsf{K} \to [0,1]$, the function $f_{\mathsf{Par}}(\mathsf{K}, g) : \mathsf{K}^\omega \to [0,1]$ is measurable, PI upward well-founded and $\mathsf{Std}_f(\mathsf{O})$-$\mathsf{PSAW}$ (for both players).*

Corollary 3.27. *Consider a finite set of colors $\mathsf{K} \subseteq \mathbb{N}$ and a map $g : \mathsf{K} \to [0,1]$. In all finite turn-based games with $f_{\mathsf{Par}}(\mathsf{K}, g) : \mathsf{K}^\omega \to [0,1]$ as payoff function, both players have positional optimal strategies.*

Proof. Corollary 3.25 ensures that both players have subgame optimal strategies. Furthermore, Corollary 3.23 along with Proposition 3.26 ensure that such a subgame optimal strategy can be chosen positional. □

3.3.2 . When optimality implies subgame optimality

In this subsection, we focus on when the existence of optimal strategies implies the existence of subgame optimal strategies in standard concurrent games. This is not always the case as exemplified in the game of Figure 3.2. The goal of this subsection is not to consider the kind of payoff functions for which this holds but rather to come up with a structural condition on the optimal strategy considered to ensure this transfer. By structural condition, we mean a condition that does not depend on the payoff function considered, only on the arena.

Let us give the intuition behind the structural condition we consider. Consider again the co-Büchi game of Figure 3.2. Recall that the optimal strategy we described first plays the top row with increasing probability and the middle row with decreasing probability and then, once Player B plays the second column, switches to a positional strategy playing the bottom row with positive, yet small enough probability. Note that switching strategy is essential. Indeed, if Player A does not switch, Player B could at some point opt for the middle column and see indefinitely the state q_1' with very high probability. In fact, what happens in that case is rather counter-intuitive: once Player B switches, there is infinitely often a positive probability to reach the stopping state of value 1. However, the probability to ever reach this outcome can be arbitrarily small, if Player B waits long enough before playing the middle column. This happens because the probability ε_k to visit that outcome goes (fast) to 0 when k goes to ∞. In fact, such an optimal strategy is not "positively bounded" in the sense that it may prescribe positive and yet arbitrarily small probabilities.

In this subsection, we consider positively bounded strategies, i.e. strategies for which there is a positive $\delta > 0$ such that any positive probability is at least δ.

Definition 3.17 (Positively bounded strategy)**.** *Let $\mathcal{C}$ be a concurrent arena. A Player-A strategy $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ is* positively bounded *if there is some $\delta > 0$ such that, for all $\rho \in Q^+$ and $a \in \mathsf{Act}_\mathsf{A}^{\rho_{\mathsf{lt}}}$, we have $\mathsf{s}_\mathsf{A}(\rho)(q) \in [0] \cup [\delta, 1]$.*

Interestingly, if we assume that there is an optimal strategy that is positively bounded, then there is a subgame optimal strategy (that is also positively bounded).

Theorem 3.28. *Consider a standard finite concurrent PI game $\mathcal{G}$ and, for all $q \in Q$, a subset of GF-strategies $\Lambda_q \subseteq \mathcal{D}(\mathsf{Act}_\mathsf{A}^q)$. Let $\Lambda = (\Lambda_q)_{q \in Q}$. Then, both assertions below are equivalent:*

a. Player-A has a positively bounded subgame optimal strategy generated by Λ;

b. Player-A has a positively bounded optimal strategy generated by Λ.

Note that, as for Theorem 3.17, the equivalence is stated in terms of existence of strategies, not on the strategies themselves. Interestingly, the proof of Theorem 3.28 above uses the notion of reset strategies, as in Section 3.1. We give a proof sketch here. The complete proof is quite technical and can be found in Appendix 3.6.7.

Proof Sketch. Consider an optimal positively bounded strategy $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ generated by Λ. We build a subgame optimal strategy $\mathsf{s}_\mathsf{A}' \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ in the following way: for all $\rho \in Q^+$, if the residual strategy $\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho)}$ is optimal from ρ_{lt}, then

$\mathsf{s}'_\mathsf{A}(\rho) := \mathsf{s}_\mathsf{A}(\rho)$, otherwise $\mathsf{s}'_\mathsf{A}(\rho) := \mathsf{s}_\mathsf{A}(\rho_{\mathsf{lt}})$ (i.e. we reset the strategy). Straightforwardly, the strategy s'_A is positively bounded and is generated by Λ. We want to apply Corollary 3.14 to prove that it is subgame optimal. One can see that it is locally optimal (by the criterion chosen for resetting the strategy and by Lemma 3.10). Consider now some $\rho \in Q^+$ and a state $q \in Q$. Assume that the residual strategy $\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho)}$ is optimal from ρ_{lt} but that the residual strategy $\mathsf{s}_\mathsf{A}^{\rho}$ is not from q. Then, similarly to why local optimality is necessary for subgame optimality (recall Lemma 3.10 cited above), one can show that a Player B action $b \in \mathsf{Act}_\mathsf{B}^{\rho_{\mathsf{lt}}}$ leading to q from ρ with positive probability is such that $\chi_\mathcal{G}(\rho_{\mathsf{lt}}) < \mathsf{out}[\langle \mathsf{F}(q), \chi_\mathcal{G} \rangle](\mathsf{s}_\mathsf{A}(\rho), b)$. Hence, there is a positive probability from ρ, if Player B opts for the action b, to reach a state of value different from $u = \chi_\mathcal{G}(q)$. And if this happens infinitely often, a state of value different from u will be reached almost-surely[12]. Thus, if a value slice is never left, almost-surely, the strategy s'_A only resets finitely often.

Consider now some $\rho \in Q^+$, a Player-B strategy $\mathsf{s}_\mathsf{B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$ and a value $u \in V^\mathcal{G} \setminus \{0\}$. From what we argued above, the probability of the event $Q^* \cdot (Q_u)^\omega$ (resp. $\{f \geq u\} \cap Q^* \cdot (Q_u)^\omega$) is the same if we intersect it with the fact that the strategy s'_A only resets finitely often. Furthermore, if the strategy does not reset anymore from some point on, and all states have the same value $u > 0$, then the strategy is, somehow, subgame optimal. It follows that the probability of $\{f \geq u\}$ is 1 by Theorem 3.12. We can then conclude by applying Corollary 3.14. □

3.4 Reduction to turn-based games: finite-choice strategies

In this section, we focus on how to transfer already existing results on turn-based games to standard concurrent games. Note that it is different from what we did in Subsection 3.3.1 since, here we do not prove results on turn-based games but rather use already existing ones. We establish such transfers in the second subsection (and also in Chapter 6), whereas the first subsection, which is quite heavy on notations, gives the necessary definitions and lemmas to prove these results.

3.4.1 . Sequentialization of standard concurrent games

To use what already exists in turn-based games, we define how to sequentialize a concurrent game into a turn-based game. That is, we make Player A play first, and then Player B respond — therefore, she has more information

[12]This holds because the strategy s_A is positively bounded: the probability to see a state of different value is bounded below by the product of the constant δ of Definition 3.17 and the smallest positive probability distribution over states in local interactions.

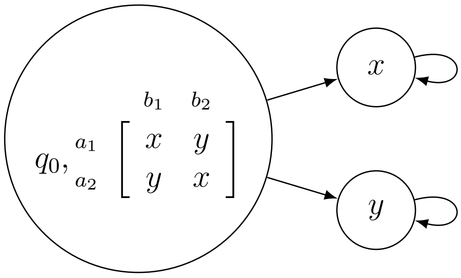

Figure 3.10: A concurrent arena.

when playing than in the original concurrent arena. This transformation is defined on standard games that need not be finite. Note that we first introduced the notion of sequentialization of standard concurrent games in [38], along with the notions of parallelization of sequentialization of strategies that we will consider later in this subsection. We will come back to that paper, and to the results of this section, in Chapter 6.

Definition 3.18 (Sequentialization of concurrent games)**.** *Consider a standard concurrent arena $\mathcal{C}$, a collection $\Lambda = (\Lambda_q)_{q \in Q} \subseteq \prod_{q \in Q} \Sigma_{\mathsf{A}}^q$ of sets of Player-A GF-strategies, another set of colors K' and a function $\eta : \mathsf{K} \to \mathsf{K}'$. The turn-based arena $\mathcal{C}(\Lambda, \eta)$ that is the sequentialization of $\mathcal{C}$ w.r.t. Λ and η is defined by $\mathcal{C}(\Lambda, \eta) := \langle Q_A \cup Q_{\mathsf{B}}, \mathsf{F}^\Lambda, \mathsf{K}^\eta, \mathsf{col}^\eta \rangle$ where:*

- $Q_{\mathsf{A}} := Q$ *and* $Q_{\mathsf{B}} := \cup_{q \in Q} \cup_{\sigma_{\mathsf{A}} \in \Lambda_q} (q, \sigma_{\mathsf{A}})$*;*
- *for all Player-A states $q \in Q_{\mathsf{A}}$, $\mathsf{F}^\Lambda(q) := \langle \Lambda_q, \{*\}, Q_{\mathsf{B}}, \mathsf{Next}_q^\Lambda \rangle_{\mathsf{s}}$ where for all $\sigma_{\mathsf{A}} \in \Lambda_q$, we have $\mathsf{Next}_q^\Lambda(\sigma_{\mathsf{A}})((q, \sigma_{\mathsf{A}})) := 1$. Note that the strategies available to Player A at such a state is equal to $\mathcal{D}(\Lambda_q)$.*
- *for all Player-B states $(q, \sigma_{\mathsf{A}}) \in Q_{\mathsf{B}}$, $\mathsf{F}^\Lambda(q) := \langle \{*\}, \mathsf{Act}_{\mathsf{B}}^q, Q_{\mathsf{A}}, \mathbb{E}(\varrho_q(\sigma_{\mathsf{A}}, \cdot)) \rangle_{\mathsf{s}}$.*
- $\mathsf{K}^\eta := \mathsf{K} \uplus \mathsf{K}'$*, col^η coincides with col on $Q = Q_{\mathsf{A}}$ and, for all $(q, \sigma_{\mathsf{A}}) \in Q_{\mathsf{B}}$, we have $\mathsf{col}^\eta((q, \sigma_{\mathsf{A}})) := \eta \circ \mathsf{col}(q)$.*

Consider now a game $\mathcal{G} = \langle \mathcal{C}, f \rangle$. The sequentialization $f_\eta : (\mathsf{K}^\eta)^\omega \to [0, 1]$ of the payoff function f w.r.t. η is such that, for all $\rho \in \mathsf{K}^\omega$:

$$f_\eta(\rho) := \begin{cases} 0 & \text{if } \rho \notin (\mathsf{K} \cup \mathsf{K}')^* \cdot (\mathsf{K} \cdot \mathsf{K}')^\omega \\ f_{\mathsf{K},\mathsf{K}'}(\rho') & \text{otherwise, for } \rho' \text{ the longest suffix of } \rho \in (\mathsf{K} \cdot \mathsf{K}')^\omega \end{cases}$$

where $f_{\mathsf{K},\mathsf{K}'}$ comes from Definition 1.8. We denote by $\mathcal{G}(\Lambda, \eta) := \langle \mathcal{C}(\Lambda, \eta), f_\eta \rangle$ the sequentialization of the game $\mathcal{G}$ w.r.t. Λ and η.

Example 3.2. *We have depicted in Figure 3.11 two possible sequentializations of the concurrent arena depicted in Figure 3.10. In both cases, we make Player A (she owns square-shaped states) play first, and then Player B respond*

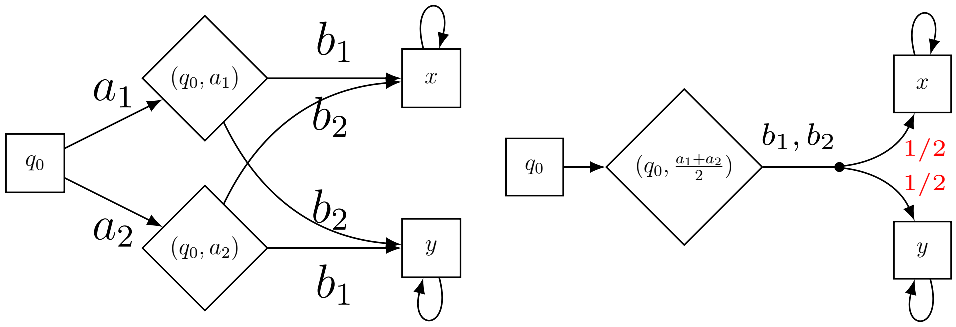

Figure 3.11: The sequentialization of the arena of Figure 3.10 with $\Lambda_{q_0} = \{a_1, a_2\}$ on the left and with $\Lambda_{q_0} = \{\frac{a_1+a_2}{2}\}$ on the right.

(she owns diamond-shaped states). Note that, in the sequentialization on the right of Figure 3.11, the choices of Player B *lead to the same outcome since Player* A *plays both rows with uniform probability. Hence, regardless of what Player* B *does, both states x and y are reached with uniform probability.*

We would like to relate what happens in a concurrent game and its sequentialized version. To do so, we translate strategies back and forth between the two games. Let us first define the parallelization of strategies, that is the transfer of strategies from the sequentialized version $\mathcal{C}(\Lambda, \eta)$ of a concurrent arena $\mathcal{C}$ back to that concurrent arena. In the following, we will consider two cases: the parallelization of Player-A deterministic strategies — along with a parallelization of Player-B arbitrary strategies, w.r.t. a Player-A deterministic strategy — and the parallelization of Player-A finite-memory strategies — which will not induce a parallelization of Player-B strategies. We first focus on the case of Player-A deterministic strategies, we will consider the case of finite-memory Player-A strategies at the end of this subsection. We define how to extend a finite path in the concurrent arena $\mathcal{C}$ into a path in its sequentialized version $\mathcal{C}(\Lambda, \eta)$, given such a Player-$\mathsf{A}$ deterministic strategy $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$. Recall that such a strategy s_A is such that, for all $\rho \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^* \cdot Q_\mathsf{A}$, we have $\mathsf{s}_\mathsf{A}(\rho) \in \Lambda_{\rho_{\mathsf{lt}}}$, which allows to define such an extension of finite paths.

Definition 3.19 (Parallelization of strategies w.r.t. a deterministic Player-A strategy). *Consider a standard concurrent arena $\mathcal{C}$, a collection $\Lambda = (\Lambda_q)_{q \in Q} \in \prod_{q \in Q} \Sigma_\mathsf{A}^q$ of sets of Player-A GF-strategies and $\eta : \mathsf{K} \to \mathsf{K}'$ for some set K'. Consider a Player-A deterministic strategy $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$. We define the function $\theta^\mathsf{A}(\mathsf{s}_\mathsf{A}) : Q^+ \to (Q_\mathsf{A} \cdot Q_\mathsf{B})^* \cdot Q_\mathsf{A}$ inductively such that, for all $\rho \in Q^+$, we have $\theta^\mathsf{A}(\mathsf{s}_\mathsf{A})(\rho)_{\mathsf{lt}} = \rho_{\mathsf{lt}}$. Specifically, for all $q \in Q$, we set $\theta^\mathsf{A}(\mathsf{s}_\mathsf{A})(q) := q$. Furthermore, for all $\rho \cdot q \in Q^+$, we set:*

$$\theta^\mathsf{A}(\mathsf{s}_\mathsf{A})(\rho \cdot q) := \theta^\mathsf{A}(\mathsf{s}_\mathsf{A})(\rho) \cdot (\rho_{\mathsf{lt}}, \mathsf{s}_\mathsf{A}(\theta^\mathsf{A}(\mathsf{s}_\mathsf{A})(\rho))) \cdot q$$

We then define $\theta^{\mathsf{B}}(\mathsf{s_A}) : Q^+ \to (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^+$ *by, for all* $\rho \in Q^+$: $\theta^{\mathsf{B}}(\mathsf{s_A})(\rho) := \theta^{\mathsf{A}}(\mathsf{s_A})(\rho) \cdot (\rho_{\mathsf{lt}}, \mathsf{s_A}(\theta^{\mathsf{A}}(\mathsf{s_A})(\rho)))$.

*Consider now a deterministic Player-*A *strategy* $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}(\Lambda,\eta)}$ *(resp. and an arbitrary Player-*B *strategy* $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}(\Lambda,\eta)}$*). We let* $\mathsf{Pr}_{\mathsf{A}}^{\Lambda}(\mathsf{s_A}) \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ *(resp.* $\mathsf{Pr}_{\mathsf{B}}^{\Lambda}(\mathsf{s_A}, \mathsf{s_B}) \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$*) be the* parallelization *of the strategy* $\mathsf{s_A}$ *(resp. of the strategy* $\mathsf{s_B}$ *w.r.t.* $\mathsf{s_A}$*) such that, for all* $\rho \in Q^+$*, we have:*

$$\mathsf{Pr}_{\mathsf{A}}^{\Lambda}(\mathsf{s_A})(\rho) := \mathsf{s_A}(\theta^{\mathsf{A}}(\mathsf{s_A})(\rho)) \in \Lambda_{\rho_{\mathsf{lt}}} \subseteq \mathcal{D}(\mathsf{Act}_{\mathsf{A}}^{\rho_{\mathsf{lt}}})$$
$$\mathsf{Pr}_{\mathsf{B}}^{\Lambda}(\mathsf{s_A}, \mathsf{s_B})(\rho) := \mathsf{s_B}(\theta^{\mathsf{B}}(\mathsf{s_A})(\rho)) \in \mathcal{D}(\mathsf{Act}_{\mathsf{B}}^{\rho_{\mathsf{lt}}})$$

We can then relate the expected value of the payoff functions in a concurrent arena and its sequentialized version w.r.t. to the above-defined parallelization of strategies.

Lemma 3.29 (Proof 3.6.8)**.** *Consider a standard concurrent game* $\mathcal{G}$*, a collection* $\Lambda = (\Lambda_q)_{q \in Q} \in \prod_{q \in Q} \Sigma_{\mathsf{A}}^q$ *of sets of Player-*A GF*-strategies and* $\eta : \mathsf{K} \to \mathsf{K}'$ *for some set* K'*. For all Player-*A *deterministic strategies* $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}(\Lambda,\eta)}$*, Player-*B *strategies* $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}(\Lambda,\eta)}$ *in* $\mathcal{C}(\Lambda, \eta)$ *and states* $q \in Q = Q_{\mathsf{A}}$*, we have:*

$$\mathbb{E}_{\mathcal{C},q}^{\mathsf{Pr}_{\mathsf{A}}^{\Lambda}(\mathsf{s_A}), \mathsf{Pr}_{\mathsf{B}}^{\Lambda}(\mathsf{s_A}, \mathsf{s_B})}[(f_{\mathcal{C}})^q] = \mathbb{E}_{\mathcal{C}(\Lambda,\eta),q}^{\mathsf{s_A},\mathsf{s_B}}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q]$$

Let us now define a way to translate strategies from a concurrent arena to its sequentialized version, which is called the sequentialization of strategies. In this direction, it suffices to consider a projection of the paths $(Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^+$ into paths in $Q_{\mathsf{A}}^+ = Q^+$. Note that, when doing so for Player B, there is loss of information since she no longer knows what Player A has played before making her move — which is the situation in the concurrent setting.

Definition 3.20 (Sequentialization of strategies)**.** *Consider a standard concurrent arena* $\mathcal{C}$*, a collection* $\Lambda = (\Lambda_q)_{q \in Q} \in \prod_{q \in Q} \Sigma_{\mathsf{A}}^q$ *of sets of Player-*A GF*-strategies and* $\eta : \mathsf{K} \to \mathsf{K}'$ *for some set* K'*. Consider a Player* $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$ *and a Player-*C *strategy* $\mathsf{t_C} \in \mathsf{S}_{\mathsf{C}}^{\mathcal{C}}$*. We let* $\mathsf{s_C}(\mathsf{t_C}) \in \mathsf{S}_{\mathsf{C}}^{\mathcal{C}(\Lambda,\eta)}$ *be such that, for all* $\rho \in (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^* \cdot Q^{\mathsf{C}}$ *(where* $Q^{\mathsf{A}} := Q_{\mathsf{A}}$ *and* $Q^{\mathsf{B}} := \epsilon$*), we have:*

$$\mathsf{s_C}(\mathsf{t_C})(\rho) := \mathsf{s_C} \circ \phi_{Q_{\mathsf{A}},Q_{\mathsf{B}}}(\rho)$$

The strategy $\mathsf{s_C}(\mathsf{t_C}) \in \mathsf{S}_{\mathsf{C}}^{\mathcal{C}(\Lambda,\eta)}$ *is defined arbitrarily, in a deterministic way, on any other path.*

Interestingly, the sequentialization and parallelization of strategies relate.

Lemma 3.30 (Proof 3.6.9)**.** *Consider a standard concurrent arena* $\mathcal{C}$*, a collection* $\Lambda = (\Lambda_q)_{q \in Q} \in \prod_{q \in Q} \Sigma_{\mathsf{A}}^q$ *of sets of Player-*A GF*-strategies and* $\eta : \mathsf{K} \to \mathsf{K}'$ *for some set* K'*. For all Player-*A *strategies* $\mathsf{t_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ *generated by* Λ*, the strategy* $\mathsf{s_A}(\mathsf{t_A}) \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}(\Lambda,\eta)}$ *is deterministic and we have:*

$$\mathsf{Pr}_{\mathsf{A}}^{\Lambda}(\mathsf{s_A}(\mathsf{t_A})) = \mathsf{t_A}$$

Furthermore, for all Player-B *strategies* $\mathsf{t_B} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$, *we have* $\mathsf{s_B}(\mathsf{t_B}) \in \mathsf{S}^{\mathcal{C}(\Lambda,\eta)}_{\mathsf{B}}$ *and for all Player-*A *deterministic strategies* $\mathsf{s_A} \in \mathsf{S}^{\mathcal{C}(\Lambda,\eta)}_{\mathsf{A}}$, *we have:*

$$\mathsf{Pr}^{\Lambda}_{\mathsf{B}}(\mathsf{s_A}, \mathsf{s_B}(\mathsf{t_B})) = \mathsf{t_B}$$

We can now compare the change in values of the games after this sequentialization. Let us first consider Player B. After this sequentialization, Player B has more information when playing than before since she knows the GF-strategy played by Player A in the current state of the game. Hence, the Player-B value of the game, after sequentialization, has not increased (i.e. it has not worsen, from Player B's point of view), as stated in the lemma below.

Proposition 3.31. *For all standard concurrent games* $\mathcal{G}$, *collections* $\Lambda = (\Lambda_q)_{q \in Q} \in \prod_{q \in Q} \Sigma^q_{\mathsf{A}}$ *of sets of Player-*A GF*-strategies and* $\eta : \mathsf{K} \to \mathsf{K}'$ *for some set* K', *we have for all* $q \in Q = Q_{\mathsf{A}}$, $\chi_{\mathcal{G}}[\mathsf{B}](q) \geq \chi_{\mathcal{G}(\Lambda,\eta)}[\mathsf{B}](q)$.

Proof. Consider any Player-B strategy $\mathsf{t_B} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$ and the Player-B strategy $\mathsf{s_B}(\mathsf{t_B}) \in \mathsf{S}^{\mathcal{C}(\Lambda,\eta)}_{\mathsf{B}}$ in the turn-based arena $\mathcal{C}(\Lambda, \eta)$. Let us show that $\chi_{\mathcal{G}}(\mathsf{t_B})(q) \geq \chi_{\mathcal{G}(\Lambda,\eta)}(\mathsf{s_B}(\mathsf{t_B}))(q)$. Consider any Player-$\mathsf{A}$ deterministic strategy $\mathsf{s_A} \in \mathsf{S}^{\mathcal{C}(\Lambda,\eta)}_{\mathsf{A}}$. Let $\mathsf{t_A} := \mathsf{Pr}^{\Lambda}_{\mathsf{A}}(\mathsf{s_A}) \in \mathsf{S}^{\mathcal{C}}_{\mathsf{A}}$ be a Player-A strategy in the arena $\mathcal{C}$. By Lemma 3.29, we have:

$$\mathbb{E}^{\mathsf{t_A},\mathsf{Pr}^{\Lambda}_{\mathsf{B}}(\mathsf{s_A},\mathsf{s_B}(\mathsf{t_B}))}_{\mathcal{C},q}[(f_{\mathcal{C}})^q] = \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}(\mathsf{t_B})}_{\mathcal{C}(\Lambda,\eta),q}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q]$$

By Lemma 3.30, we have $\mathsf{Pr}^{\Lambda}_{\mathsf{B}}(\mathsf{s_A}, \mathsf{s_B}(\mathsf{t_B})) = \mathsf{t_B}$. Hence:

$$\chi_{\mathcal{G}}(\mathsf{t_B})(q) \geq \mathbb{E}^{\mathsf{t_A},\mathsf{t_B}}_{\mathcal{C},q}[(f_{\mathcal{C}})^q] = \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}(\mathsf{t_B})}_{\mathcal{C}(\Lambda,\eta),q}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q]$$

As this holds for all Player-A deterministic strategies $\mathsf{s_A} \in \mathsf{S}^{\mathcal{C}(\Lambda,\eta)}_{\mathsf{A}}$ and since deterministic strategies achieve the same values that all strategies in turn-based games by Corollary 2.17, it follows that $\chi_{\mathcal{G}}(\mathsf{t_B})(q) \geq \chi_{\mathcal{G}(\Lambda,\eta)}(\mathsf{s_B}(\mathsf{t_B}))(q) \geq \chi_{\mathcal{G}(\Lambda,\eta)}[\mathsf{B}](q)$. As this holds for all Player-B strategies $\mathsf{t_B} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$, it follows that $\chi_{\mathcal{G}}[\mathsf{B}](q) \geq \chi_{\mathcal{G}(\Lambda,\eta)}[\mathsf{B}](q)$. □

The case of Player A is not exactly symmetrical. Indeed, she cannot achieve the same value in $\mathcal{G}(\Lambda, \eta)$ than in $\mathcal{G}$ because she has less available strategies, and Player B knows what she played before playing. Hence, the Player-A value of the game, after sequentialization, has not increased. However, this value in the sequentialization is at least the supremum of the values of strategies generated by Λ.

Proposition 3.32. *Consider a standard game* $\mathcal{G}$, *a collection* $\Lambda = (\Lambda_q)_{q \in Q} \in \prod_{q \in Q} \Sigma^q_{\mathsf{A}}$ *of sets of Player-*A GF*-strategies and* $\eta : \mathsf{K} \to \mathsf{K}'$ *for some set* K'. *Let* $q \in Q = Q_{\mathsf{A}}$. *We denote by* $\mathsf{S}^{\mathcal{C}}_{\mathsf{A}}(\Lambda)$ *the set of Player-*A *strategies generated by*

Λ *in the arena* $\mathcal{C}$. *For all deterministic Player-*A *strategies* $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$, *the Player-*$\mathsf{A}$ *strategy* $\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}) \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ *is generated by* Λ *and is such that:*

$$\chi_{\mathcal{G}(\Lambda,\eta)}(\mathsf{s_A})(q) \leq \chi_\mathcal{G}(\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}))(q)$$

In fact, we have:

$$\chi_{\mathcal{G}(\Lambda,\eta)}[\mathsf{A}](q) = \sup_{\mathsf{t_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda)} \chi_\mathcal{G}[\mathsf{t_A}](q) \leq \chi_\mathcal{G}[\mathsf{A}](q)$$

Proof. Consider such a Player-A deterministic strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$ and the Player-A strategy $\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}) \in \mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda)$ in the concurrent arena $\mathcal{C}$. Note that $\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}) \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$ straightforwardly from the definition. Let us show that $\chi_{\mathcal{G}(\Lambda,\eta)}(\mathsf{s_A})(q) \leq \chi_\mathcal{G}(\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}))(q)$. Consider any Player-$\mathsf{B}$ strategy $\mathsf{t_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$. Let $\mathsf{s}'_\mathsf{B} := \mathsf{s_B}(\mathsf{t_B}) \in \mathsf{S}_\mathsf{B}^{\mathcal{C}(\Lambda,\eta)}$ be a Player-B strategy in the arena $\mathcal{C}(\Lambda,\eta)$. By Lemma 3.29, we have:

$$\mathbb{E}_{\mathcal{C},q}^{\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}),\mathsf{Pr}_\mathsf{B}^\Lambda(\mathsf{s_A},\mathsf{s}'_\mathsf{B})}[(f_\mathcal{C})^q] = \mathbb{E}_{\mathcal{C}(\Lambda,\eta),q}^{\mathsf{s_A},\mathsf{s}'_\mathsf{B}}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q]$$

Furthermore, we have $\mathsf{s}'_\mathsf{B} = \mathsf{s_B}(\mathsf{t_B})$, hence by Lemma 3.30, we have $\mathsf{Pr}_\mathsf{B}^\Lambda(\mathsf{s_A},\mathsf{s}'_\mathsf{B}) = \mathsf{t_B}$. Hence:

$$\mathbb{E}_{\mathcal{C},q}^{\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}),\mathsf{t_B}}[(f_\mathcal{C})^q] = \mathbb{E}_{\mathcal{C}(\Lambda,\eta),q}^{\mathsf{s_A},\mathsf{s}'_\mathsf{B}}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q] \geq \chi_{\mathcal{G}(\Lambda,\eta)}(\mathsf{s_A})(q)$$

As this holds for all Player-B strategies $\mathsf{t_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$, it follows that $\chi_{\mathcal{G}(\Lambda,\eta)}(\mathsf{s_A})(q) \leq \chi_\mathcal{G}(\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}))(q) \leq \sup_{\mathsf{t_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda)} \chi_\mathcal{G}[\mathsf{t_A}](q)$, since $\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}) \in \mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda)$.

Furthermore, since this holds for all Player-A deterministic strategies $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$ and since deterministic strategies achieve the same values that all strategies in turn-based games by Corollary 2.17, it follows that $\chi_{\mathcal{G}(\Lambda,\eta)}[\mathsf{A}](q) \leq \sup_{\mathsf{t_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda)} \chi_\mathcal{G}[\mathsf{t_A}](q)$.

Consider now a Player-A strategy $\mathsf{t_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ generated by Λ (i.e. $\mathsf{t_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda)$). Consider then the Player-A deterministic strategy $\mathsf{s_A}(\mathsf{t_A}) \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$ in the turn-based arena $\mathcal{C}(\Lambda,\eta)$. Let us show that $\chi_\mathcal{G}(\mathsf{t_A})(q) \leq \chi_{\mathcal{G}(\Lambda,\eta)}(\mathsf{s_A}(\mathsf{t_A}))(q)$. Consider any Player-$\mathsf{B}$ strategy $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^{\mathcal{C}(\Lambda,\eta)}$. Let $\mathsf{t_B} := \mathsf{Pr}_\mathsf{B}^\Lambda(\mathsf{s_A}(\mathsf{t_A}),\mathsf{s_B}) \in \mathsf{S}_\mathsf{B}^\mathcal{C}$ be a Player-B strategy in the arena $\mathcal{C}$. By Lemma 3.29, we have:

$$\mathbb{E}_{\mathcal{C},q}^{\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}(\mathsf{t_A})),\mathsf{t_B}}[(f_\mathcal{C})^q] = \mathbb{E}_{\mathcal{C}(\Lambda,\eta),q}^{\mathsf{s_A}(\mathsf{t_A}),\mathsf{s_B}}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q]$$

By Lemma 3.30, we have $\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}(\mathsf{t_A})) = \mathsf{t_A}$. Hence:

$$\chi_\mathcal{G}(\mathsf{t_A})(q) \leq \mathbb{E}_{\mathcal{C},q}^{\mathsf{t_A},\mathsf{t_B}}[(f_\mathcal{C})^q] = \mathbb{E}_{\mathcal{C}(\Lambda,\eta),q}^{\mathsf{s_A}(\mathsf{t_A}),\mathsf{s_B}}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q]$$

As this holds for all Player-B strategies $\mathsf{s_B} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$, it follows that $\chi_\mathcal{G}(\mathsf{t_A})(q) \leq \chi_{\mathcal{G}(\Lambda,\eta)}(\mathsf{s_A}(\mathsf{t_A}))(q) \leq \chi_{\mathcal{G}(\Lambda,\eta)}[\mathsf{A}](q)$. As this holds for all Player-A strategies $\mathsf{t_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda)$ generated by Λ, it follows that $\sup_{\mathsf{t_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda)} \chi_\mathcal{G}[\mathsf{t_A}](q) \leq \chi_{\mathcal{G}(\Lambda,\eta)}[\mathsf{A}](q)$.

Since $\mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda) \subseteq \mathsf{S}_\mathsf{A}^\mathcal{C}$ and $\chi_\mathcal{G}[\mathsf{A}](q) = \sup_{\mathsf{t_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}} \chi_\mathcal{G}[\mathsf{t_A}](q)$, it follows that we also have $\sup_{\mathsf{t_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda)} \chi_\mathcal{G}[\mathsf{t_A}](q) \leq \chi_\mathcal{G}[\mathsf{A}](q)$. $\square$

Example 3.3. *Let us illustrate these lemmas on the sequentializations of Figure 3.11. Assume that the objective of Player* A *is to reach* x*, while Player* B *wants to avoid it. In that case, the value of the original concurrent game of Figure 3.10 is* $\frac{1}{2}$ *and both players have optimal strategies: it suffices to play uniformly over their pair of actions. Now, consider the sequentialization on the left of Figure 3.11. Here, the value of the game is 0. Indeed, since Player* B *has the information of the action chosen by Player* A*, she can enforce going to state* y *surely. Note that, in this game, it is useless for Player* A *to play non-deterministic strategies. Consider now the sequentialization on the right of Figure 3.11. In that game, the value is* $\frac{1}{2}$ *since regardless of what the players do, the states* x *and* y *will be reached with probability* $\frac{1}{2}$*. Note that the optimal Player-*A *strategy in the game of Figure 3.11 that consists in playing both rows with probability* $\frac{1}{2}$ *is generated by* $\Lambda_{q_0} = \{\frac{a_1+a_2}{2}\}$*, which has induced the sequentialization on the right.*

Finite-memory strategies. Since we ultimately want to transfer results from turn-based games to standard concurrent games, we want to be able to parrallelize finite-memory strategies — recall Definition 1.36 — ideally while keeping the same amount of memory state. Given some set of colors K' and some $\eta : \mathsf{K} \to \mathsf{K}'$, consider a Player-$\mathsf{A}$ finite-memory strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$ (on the set of color $\mathsf{K} \cup \mathsf{K}'$). One can notice that, by definition of the arena $\mathcal{C}(\Lambda, \eta)$, given a state $q \in Q_\mathsf{A}$ of color $\mathsf{col}(q) = \mathsf{col}^\eta(q) := k \in \mathsf{K}$, then the color of all Player-B states reachable from q is the same, and is equal to $\eta(k) \in \mathsf{K}'$. The parallelization of Player-A finite-memory strategies $\mathsf{s_A}$ therefore only amounts to properly handling the memory update.

Definition 3.21 (Parallelization of Player-A finite-memory strategies)**.** *Consider a standard concurrent arena* $\mathcal{C}$*, a collection* $\Lambda = (\Lambda_q)_{q \in Q} \in \prod_{q \in Q} \Sigma_\mathsf{A}^q$ *of sets of Player-*A GF*-strategies and* $\eta : \mathsf{K} \to \mathsf{K}'$ *for some set* K'*. Consider a Player-*A *finite-memory strategy* $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$ *that is implemented by a memory skeleton* $\mathsf{M} = \langle M, m_{\text{init}}, \mu \rangle$ *on* $\mathsf{K} \cup \mathsf{K}'$ *and an action map* $\lambda : M \times (Q_\mathsf{A} \cup Q_\mathsf{B}) \to \cup_{q \in Q} \Sigma_\mathsf{A}(\mathsf{F}^\Lambda(q))$*.*

We denote by $\mathsf{Pr}_\mathsf{A}^\eta(\mathsf{M})$ *the memory skeleton* $\mathsf{Pr}_\mathsf{A}^\eta(\mathsf{M}) := \langle M, m_{\text{init}}, \mathsf{Pr}_\mathsf{A}^\eta(\mu) \rangle$ *on* K *such that, for all* $m \in M$ *and* $k \in \mathsf{K}$*, we have* $\mathsf{Pr}_\mathsf{A}^\eta(\mu)(m, k) := \mu(\mu(m, k), \eta(k)) \in M$*. Then, we denote by* $\mathsf{Pr}_\mathsf{A}^\Lambda(\lambda) : M \times Q \to \cup_{q \in Q} \Sigma_\mathsf{A}(\mathsf{F}(q))$ *the action map defined by, for all* $m \in M$ *and* $q \in Q$*, we have* $\mathsf{Pr}_\mathsf{A}^\Lambda(\lambda)(m, q) := \lambda(m, q) \in \mathcal{D}(\Lambda_q) \subseteq \mathcal{D}(\mathsf{Act}_\mathsf{A}^q)$*. We denote by* $\mathsf{Pr}_\mathsf{A}^{\eta,\Lambda}(\mathsf{s_A}) \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ *the Player-*A *strategy implemented by* $\mathsf{Pr}_\mathsf{A}^\eta(\mathsf{M})$ *and* $\mathsf{Pr}_\mathsf{A}^\Lambda(\lambda)$*.*

In fact, one has to check that the above definition is well-defined. Indeed, a Player-A finite-memory strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$ could be implemented with different memory skeletons and action maps. One has to check that regardless of the pair implementing the strategy on which is done the parallelization, the resulting Player-A strategy in $\mathsf{S}_\mathsf{A}^\mathcal{C}$ is the same. This is done in the lemma below.

Lemma 3.33 (Proof 3.6.10). *Consider a standard concurrent arena $\mathcal{C}$, a collection $\Lambda = (\Lambda_q)_{q \in Q} \in \prod_{q \in Q} \Sigma^q_{\mathsf{A}}$ of sets of Player-A GF-strategies and $\eta : \mathsf{K} \to \mathsf{K}'$ for some set K'. Consider a Player-A strategy $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$. For all memory skeletons M on K' and action maps $\lambda : M \times (Q_\mathsf{A} \cup Q_\mathsf{B}) \to \cup_{q \in Q_\mathsf{A} \cup Q_\mathsf{B}} \Sigma_\mathsf{A}(\mathsf{F}^\Lambda(q))$ implementing s_A, the strategy in $\mathsf{S}_\mathsf{A}^{\mathcal{C}}$ implemented by $\mathsf{Pr}^\eta_\mathsf{A}(\mathsf{M})$ and $\mathsf{Pr}^\Lambda_\mathsf{A}(\lambda)$ is the same.*

Then, letting $\mathsf{t}_\mathsf{A} := \mathsf{Pr}^{\eta,\Lambda}_\mathsf{A}(\mathsf{s}_\mathsf{A})$, for all $\pi \in Q^*$, there is some $\rho \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^*$ such that $\phi_{Q_\mathsf{A},Q_\mathsf{B}}(\rho) = \pi$ and such that $\mathsf{t}^\pi_\mathsf{A} = \mathsf{Pr}^{\eta,\Lambda}_\mathsf{A}(\mathsf{s}^\rho_\mathsf{A})$.

Similarly to what is done in Lemma 3.29, we can relate the expected values of the payoff functions in a concurrent arena and its sequentialized version w.r.t. to the parallelization of Player-A finite-memory strategies and the sequentialization of Player-B strategies.

Lemma 3.34 (Proof 3.6.11). *Consider a standard concurrent game $\mathcal{G}$, a collection $\Lambda = (\Lambda_q)_{q \in Q} \in \prod_{q \in Q} \Sigma^q_{\mathsf{A}}$ of sets of Player-A GF-strategies and $\eta : \mathsf{K} \to \mathsf{K}'$ for some set K'. For all Player-A finite-memory strategies $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$ in $\mathcal{C}(\Lambda,\eta)$, Player-$\mathsf{B}$ strategies $\mathsf{t}_\mathsf{B} \in \mathsf{S}_\mathsf{B}^{\mathcal{C}}$ in $\mathcal{C}$ and states $q \in Q$, we have:*

$$\mathbb{E}^{\mathsf{Pr}^{\eta,\Lambda}_\mathsf{A}(\mathsf{s}_\mathsf{A}),\mathsf{t}_\mathsf{B}}_{\mathcal{C},q}[(f_\mathcal{C})^q] = \mathbb{E}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}(\mathsf{t}_\mathsf{B})}_{\mathcal{C}(\Lambda,\eta),q}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q]$$

Interestingly, we can then deduce that the value of the parallelization of a finite-memory strategy is at least the value of that finite-memory strategy.

Proposition 3.35. *Consider a standard concurrent game $\mathcal{G}$, a collection $\Lambda = (\Lambda_q)_{q \in Q} \in \prod_{q \in Q} \Sigma^q_{\mathsf{A}}$ of sets of Player-A GF-strategies and $\eta : \mathsf{K} \to \mathsf{K}'$ for some set K'. For all Player-A finite-memory strategies $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$ in $\mathcal{C}(\Lambda,\eta)$ and states $q \in Q = Q_\mathsf{A}$, we have:*

$$\chi_{\mathcal{G}(\Lambda,\eta)}(\mathsf{s}_\mathsf{A})(q) \leq \chi_\mathcal{G}(\mathsf{Pr}^{\eta,\Lambda}_\mathsf{A}(\mathsf{s}_\mathsf{A}))(q)$$

Proof. Consider any Player-B strategy $\mathsf{t}_\mathsf{B} \in \mathsf{S}^\mathcal{C}_\mathsf{B}$. By Lemma 3.34, we have $\chi_{\mathcal{G}(\Lambda,\eta)}(\mathsf{s}_\mathsf{A})(q) \leq \mathbb{E}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}(\mathsf{t}_\mathsf{B})}_{\mathcal{C}(\Lambda,\eta),q}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q] = \mathbb{E}^{\mathsf{Pr}^{\eta,\Lambda}_\mathsf{A}(\mathsf{s}_\mathsf{A}),\mathsf{t}_\mathsf{B}}_{\mathcal{C},q}[(f_\mathcal{C})^q]$. As this for all Player-B strategies $\mathsf{t}_\mathsf{B} \in \mathsf{S}^\mathcal{C}_\mathsf{B}$, it follows that $\chi_{\mathcal{G}(\Lambda,\eta)}(\mathsf{s}_\mathsf{A})(q) \leq \chi_\mathcal{G}(\mathsf{Pr}^{\eta,\Lambda}_\mathsf{A}(\mathsf{s}_\mathsf{A}))(q)$. □

3.4.2 . Finite-choice strategies

Recall that the goal of this section is to retrieve results already existing in turn-based games in the context of concurrent games. We are especially interested in results existing in finite turn-games, since infinite stochastic turn-based games are hard to handle. To be able to transfer results from finite turn-based games, we focus on a special type of Player-A strategy. To gain an intuition on what the strategies we will consider are, let us consider the

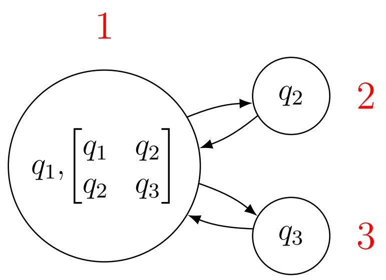

Figure 3.12: A parity game.

parity game of Figure 3.12. Note that this game was already described in [47]. We will also come back to that game in Chapter 5 and prove formally, in Proposition 5.14, the informal statements we make below. This is a parity game where the objective of Player A is to see q_2 infinitely often, while seeing q_3 only finitely often. In this game there is a subgame almost-surely winning strategy for Player A. However, one can realize that all positional Player-A strategies have value 0. Indeed, if a Player-A positional strategy $\mathsf{s_A}$ plays the bottom row with positive probability, then Player-B can positionally play the right column and ensure seeing q_3 infinitely often almost-surely. Furthermore, if $\mathsf{s_A}$ does not play the bottom row with positive probability, Player B can positionally play the left column with probability 1 and ensure looping indefinitely on q_1 without ever seeing q_2.

Let us now informally describe a Player-A subgame almost-surely winning strategy. Such a strategy could play the top row with probability $1 - \varepsilon_k$ and the bottom row with probability $\varepsilon_k > 0$ with ε_k going to 0 when k goes to ∞, where k denotes the number of times the states q_2 and q_3 are seen. Then, considering any Player-B strategy, the probability to see q_3 infinitely often is 0 if ε_k goes to 0 sufficiently fast. Furthermore, the probability to ever loop indefinitely on q_1 without ever seeing q_2 is also 0 thanks to the fact that k counts the number of times a state that is not q_1 is seen. Indeed, as long as the game loops on q_1, ε_k does not change and therefore there is probability (at least) $\varepsilon_k > 0$ to see q_2. That is, not seeing q_2 anymore does not happen, almost-surely.

What happens in this parity game of Figure 3.12 is frustrating since, although there are subgame optimal strategies, such subgame optimal strategies prescribe infinitely many different probability distributions at q_1 and cannot be found among positional optimal. This is strikingly different from the situation in finite turn-based games where there are always positional (subgame) optimal strategies [27, 28].

In fact, the issue in concurrent games lies exactly in the fact that achieving a value is that complicated — i.e. that it requires such a convoluted strategy that plays infinitely many different probability distributions. We introduce the opposite notion, that is what we call finite-choice strategies. These are strategies that, at each state, may play only among a finite set of GF-strategies.

Definition 3.22 (Finite-choice strategy)**.** *Let $\mathcal{C}$ be a concurrent arena. A Player-A strategy $\mathsf{s_A} \in \mathsf{S_A^{\mathcal{C}}}$ has* finite choice *if, for all $q \in Q$, there is a* finite *set $\Sigma_q(\mathsf{s_A}) \subseteq \Sigma_{\mathsf{A}}^q$ such that, for all $\rho \in Q^*$, we have $\mathsf{s_A}(\rho \cdot q) \in \Sigma_q(\mathsf{s_A})$. Otherwise, the strategy has* infinite choice*. The definition is analogous for Player B.*

Remark 3.1. *Note that finite-memory — and in particular positional — strategies have finite choice. Therefore, all infinite-memory strategies are not finite-choice strategies. It is also the case for deterministic strategies (since Player A has finitely many actions). In addition, in finite-state arenas, finite-choice strategies are positively-bounded.*

In fact, when finite-choice strategies achieve a value in a finite concurrent standard game with a specific objective, then we can use the already existing results in turn-based games via sequentialization. However, when doing so we add intermediate states with a color given by $\mathsf{col} : Q \to \mathsf{K}$ and $\eta : \mathsf{K} \to \mathsf{K}'$. Let us define pairs of payoff functions that can be made equal after the sequentialization.

Definition 3.23 (Payoff functions equal up to adequate interleaving)**.** *Consider two non-empty sets of colors K and K' along with two PI payoff functions $f : \mathsf{K}^\omega \to [0,1]$ and $g : (\mathsf{K} \cup \mathsf{K}')^\omega \to [0,1]$. We say that f and g are* equal up to adequate interleaving *if there is a map $\eta : \mathsf{K} \to \mathsf{K}'$ and an affine increasing function $\psi : [0,1] \to [0,1]$, such that, for all $\rho \in \mathsf{K}^\omega$:*

$$f(\rho) = \psi \circ g(\rho_0 \cdot \eta(\rho_0) \cdot \rho_1 \cdot \eta(\rho_1) \cdots)$$

Given a memory skeleton M, we extend the notion of being $(\mathsf{TB}_f(\mathsf{O}), \mathsf{M})$-$\mathsf{SAW}$ (recall subgame almost-surely winnable, where $\mathsf{TB}_f(\mathsf{O})$ refers to the set of finite turn-based game forms) to payoff functions after the sequentialization by using the above definition. This is done below in Definition 3.24.

Definition 3.24 (Seq-$(\mathsf{TB}_f(\mathsf{O}), \mathsf{M})$-SAW payoff function)**.** *Consider two non-empty sets of colors K and K' and a memory skeleton $\mathsf{M} = \langle M, m_{\mathrm{init}}, \mu \rangle$ on $\mathsf{K} \cup \mathsf{K}'$. A PI payoff function $f : \mathsf{K}^\omega \to [0,1]$ is said to be* M-subgame almost-surely winnable after sequentialization *(*Seq-M-SAW *for short) if there is a PI upward well-founded payoff function $g : (\mathsf{K} \cup \mathsf{K}')^\omega \to [0,1]$ that is $(\mathsf{TB}_f(\mathsf{O}), \mathsf{M})$-*SAW *such that f and g are equal up to adequate interleaving. When $|M| = 1$, the payoff function f is said to be* positionally subgame almost-surely winnable after sequentialization *(*Seq-PSAW *for short). This last notion does not depend on the set of colors K'.*

Let us apply this definition to parity objectives.

Proposition 3.36. *For all sets of colors* $\mathsf{K} := [\![m, n]\!] \subseteq \mathbb{N}$ *for some* $m \leq n \in \mathbb{N}$*, the parity objective* $\mathsf{Parity}_\mathsf{K}$ *— seen as the payoff function* $f := \mathbb{1}_{\mathsf{Parity}_\mathsf{K}}$ *— is* Seq-PSAW.

Proof. We let $\mathsf{K}' := \mathsf{K}$, $\eta : \mathsf{K} \to \mathsf{K}$ be such that $\eta[\mathsf{K}] := \{m\}$ and $g := f = \mathbb{1}_{\mathsf{Parity}_\mathsf{K}}$. Since positional optimal strategies always exist in finite turn-based parity games [27, 28], g is $(\mathsf{TB}_f(\mathsf{O}))$-PSAW. Furthermore, since the parity objective only considers the highest color seen infinitely often, for all $\rho \in \mathsf{K}^\omega$, we have $f(\rho) = g(\rho_0 \cdot m \cdot \rho_1 \cdot m \cdots)$ since $m = \min \mathsf{K}$. Hence, the pair (f, g) is equal up to adequate interleaving. □

Before stating the main result of this section, we need to define below the notion of B-finite standard concurrent game.

Definition 3.25 (B-finite standard concurrent game)**.** *Consider a standard concurrent game* $\mathcal{G}$*. We say that it is* B-finite *if the set of states* Q *is finite and, for all* $q \in Q$*, the set of Player-*B *actions* $\mathsf{Act}^q_\mathsf{B}$ *at state* q *is finite.*

We can finally state the memory transfer from turn-based games to standard concurrent games w.r.t. finite-choice strategies.

Theorem 3.37. *Let* K, K' *be two arbitrary sets of colors,* $\mathsf{M} = \langle M, m_{\text{init}}, \mu \rangle$ *be a finite memory skeleton on* $\mathsf{K} \cup \mathsf{K}'$ *and* $f : \mathsf{K}^\omega \to [0, 1]$ *be a PI payoff function that is* Seq-$(\mathsf{TB}_f(\mathsf{O}), \mathsf{M})$-SAW.

Then, for all B*-finite standard concurrent games* $\mathcal{G} = \langle \mathcal{C}, f \rangle$*, for all Player-*A *finite-choice strategies* $\mathsf{s_A} \in \mathsf{S}^\mathcal{C}_\mathsf{A}$*, there is a Player-*A *strategy* $\mathsf{t_A} \in \mathsf{S}^\mathcal{C}_\mathsf{A}$ *implementable by a memory skeleton using as many memory states as* M *such that, for all* $\pi \in (Q_{\mathsf{ns}})^+$*, we have* $\chi_\mathcal{G}[\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\pi)}](\pi_{\mathsf{lt}}) \leq \chi_\mathcal{G}[\mathsf{t}_\mathsf{A}^{\mathsf{tl}(\pi)}](\pi_{\mathsf{lt}})$.

Proof. First, as the payoff function $f : \mathsf{K}^\omega \to [0, 1]$ is Seq-$(\mathsf{TB}_f(\mathsf{O}), \mathsf{M})$-SAW, let us consider a PI upward well-founded function $g : (\mathsf{K} \cup \mathsf{K}')^\omega \to [0, 1]$ that is $(\mathsf{TB}_f(\mathsf{O}), \mathsf{M})$-SAW such that f and g are equal up to adequate interleaving. Let us also consider an affine increasing (and therefore invertible) function $\psi : [0, 1] \to [0, 1]$ and $\eta : \mathsf{K} \to \mathsf{K}'$ from Definition 3.23.

Consider now such a finite-choice strategy $\mathsf{s_A} \in \mathsf{S}^\mathcal{C}_\mathsf{A}$. For all $q \in Q$, we let $\Lambda_q := \{\mathsf{s_A}(\rho \cdot q) \mid \rho \in Q^*\} \subseteq \mathcal{D}(\mathsf{Act}^q_\mathsf{A})$ be a finite set — since the strategy $\mathsf{s_A}$ has finite choice — of Player-A GF-strategies at state q. Let $\Lambda := (\Lambda_q)_{q \in Q}$. Let us consider the turn-based game $\mathcal{G}(\Lambda, \eta)$. It is finite by definition of Λ. Furthermore, by Proposition 3.32, for all $q \in Q$, we have $\chi_{\mathcal{G}(\Lambda,\eta)}[\mathsf{A}](q) = \sup_{\mathsf{t_A} \in \mathsf{S}^\mathcal{C}_\mathsf{A}(\Lambda)} \chi_\mathcal{G}[\mathsf{t_A}](q)$. In particular, for all $q \in Q$ and $\rho \in Q^+$, we have $\chi_{\mathcal{G}(\Lambda,\eta)}[\mathsf{A}](q) \geq \chi_\mathcal{G}[\mathsf{s}^\rho_\mathsf{A}](q)$ since $\mathsf{s}^\rho_\mathsf{A}$ is generated by Λ.

Let us now consider the game $\mathcal{G}'(\Lambda, \eta) := \langle \mathcal{C}(\Lambda, \eta)', g \rangle$ where the arena $\mathcal{C}(\Lambda, \eta)'$ is obtained from the arena $\mathcal{C}(\Lambda, \eta)'$ by changing the values of the stopping states from u in $\mathcal{C}(\Lambda, \eta)$ to $\psi^{-1}(u)$ in $\mathcal{C}(\Lambda, \eta)'$[13]. Otherwise, the arena

[13]Note that this transformation may induce stopping states of value more than 1. However, all arguments still hold in that case.

is unchanged. We denote by $\mathsf{SeqAlt} \subseteq (Q_\mathsf{A} \cdot Q_\mathsf{B})^+$ the set:

$$\mathsf{SeqAlt} := \{\epsilon\} \cup \{q_0 \cdot (q_0, \sigma_{q_0}) \cdot q_1 \cdot (q_1, \sigma_{q_1}) \cdots q_n \cdot (q_n, \sigma_{q_n}) \mid \forall 0 \leq i \leq n,\ q_i \in Q_\mathsf{A},\ \sigma_{q_i} \in \Lambda_{q_i}\}$$

By definition of the arenas $\mathcal{C}(\Lambda, \eta)$ and $\mathcal{C}(\Lambda, \eta)'$, all finite paths that are not in SeqAlt (up to omitting the last state of the path) have probability 0 to occur in the arenas $\mathcal{C}(\Lambda, \eta)$ and $\mathcal{C}(\Lambda, \eta)'$ regardless of the strategies of the players. Furthermore, for any infinite path $\rho \in \mathsf{SeqAlt}^\omega$, we have $\psi^{-1} \circ f_\eta(\rho) = g(\rho)$ with ψ^{-1} an affine function, which therefore commutes with the expected value: for all states $q \in Q$ and pair of strategies $(x_\mathsf{A}, x_\mathsf{B}) \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)} \times \mathsf{S}_\mathsf{B}^{\mathcal{C}(\Lambda,\eta)}$, we have:

$$\psi^{-1}(\mathbb{E}^{x_\mathsf{A};x_\mathsf{B}}_{\mathcal{C}(\Lambda,\eta),q}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q]) = \mathbb{E}^{x_\mathsf{A};x_\mathsf{B}}_{\mathcal{C}(\Lambda,\eta)',q}[(g_{\mathcal{C}(\Lambda,\eta)'})^q]$$

Hence, since ψ is increasing and continuous — since it is affine, we can deduce that for all states $q \in Q$, for all Player-A strategies $x_\mathsf{A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$, we have $\psi^{-1}(\chi_{\mathcal{G}(\Lambda,\eta)}[x_\mathsf{A}](q)) = \chi_{\mathcal{G}(\Lambda,\eta)'}[x_\mathsf{A}](q)$ and $\psi^{-1}(\chi_{\mathcal{G}(\Lambda,\eta)}[\mathsf{A}](q)) = \chi_{\mathcal{G}(\Lambda,\eta)'}[\mathsf{A}](q)$.

Furthermore, since the PI function g is $(\mathsf{TB}_f(\mathsf{O}), \mathsf{M})$-SAW, by Corollary 3.25, there is a subgame almost-surely winning strategy in $\mathcal{G}(\Lambda, \eta)'$ and by Corollary 3.23, it can be found among M-implementable strategies. Therefore, we consider a Player-A strategy $x_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)'} = \mathsf{S}_\mathsf{A}^{\mathcal{C}(\Lambda,\eta)}$ that is subgame optimal in $\mathcal{G}(\Lambda, \eta)'$ and M-implementable. Consider then its parrallelization $\mathsf{t}_\mathsf{A} := \mathsf{Pr}_\mathsf{A}^{\Lambda,\eta}(x_\mathsf{A}) \in \mathsf{S}_\mathsf{A}^\mathcal{C}$. By definition (recall Definition 3.21), it is implementable with a finite memory skeleton with as many memory states as M.

Consider now any finite path $\pi \in (Q_\mathsf{ns})^+$. By (the second part of) Proposition 3.33, there is a finite path $\rho \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^*$ such that $\mathsf{t}_\mathsf{A}^{\mathsf{tl}(\pi)} = \mathsf{Pr}_\mathsf{A}^{\Lambda,\eta}(x_\mathsf{A}^\rho)$ such that $\phi_{Q_\mathsf{A},Q_\mathsf{B}}(\rho) = \pi$, and therefore ρ does not visit any stopping state in $\mathcal{C}(\Lambda, \eta)$, or equivalently in $\mathcal{C}(\Lambda, \eta)'$. Furthermore, the strategy x_A is subgame optimal in the game $\mathcal{G}(\Lambda, \eta)'$ with a PI upward well founded payoff function. Hence, $\psi^{-1}(\chi_{\mathcal{G}(\Lambda,\eta)}[x_\mathsf{A}^\rho](\pi_\mathsf{lt})) = \chi_{\mathcal{G}(\Lambda,\eta)'}[x_\mathsf{A}^\rho](\pi_\mathsf{lt}) = \chi_{\mathcal{G}(\Lambda,\eta)'}[\mathsf{A}](\pi_\mathsf{lt}) = \psi^{-1}(\chi_{\mathcal{G}(\Lambda,\eta)}[\mathsf{A}](\pi_\mathsf{lt}))$. Furthermore, we have $\chi_{\mathcal{G}(\Lambda,\eta)}[\mathsf{A}](\pi_\mathsf{lt}) \geq \chi_\mathcal{G}[\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\pi)}](\pi_\mathsf{lt})$. In addition, by Proposition 3.35, we have $\chi_{\mathcal{G}(\mathcal{C},\eta)}[x_\mathsf{A}^\rho](\pi_\mathsf{lt}) \leq \chi_\mathcal{G}[\mathsf{Pr}_\mathsf{A}^{\Lambda,\eta}(x_\mathsf{A}^\rho)](\pi_\mathsf{lt}) = \chi_\mathcal{G}[\mathsf{t}_\mathsf{A}^{\mathsf{tl}(\pi)}](\pi_\mathsf{lt})$. Overall, we do obtain, $\chi_\mathcal{G}[\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\pi)}](\pi_\mathsf{lt}) \leq \chi_\mathcal{G}[\mathsf{t}_\mathsf{A}^{\mathsf{tl}(\pi)}](\pi_\mathsf{lt})$. □

We obtain a simpler statement when applying to the special case of parity objectives.

Corollary 3.38. *Consider any* B*-finite standard concurrent parity game. For all Player-*A *finite-choice strategies* $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$*, there is a Player-*$\mathsf{A}$ *positional strategy* t_A*such that, for all* $\rho \in Q^+$*, we have* $\chi_\mathcal{G}[\mathsf{s}_\mathsf{A}](\rho) \leq \chi_\mathcal{G}[\mathsf{t}_\mathsf{A}](\rho)$*. Hence, if* s_A *is subgame optimal, so is* t_A*.*

Proof. This is a direct consequence of Theorem 3.37 and Proposition 3.36. □

3.5 Discussion and open question

In this chapter, we have established various results on subgame optimal strategies. As mentioned at the beginning of this part, we believe that Theorem 3.12 and its Corollaries 3.14 and 3.16 are important results on concurrent games and are essential to this dissertation. We also believe that Theorem 3.17 gives significant insight on why concurrent concurrent games behave so much more badly than turn-based games. Indeed, as discussed in Page 137, this theorem identifies exactly the reason why there does not always exist subgame optimal strategies in concurrent games: when restricting to locally optimal strategies, the value of some states may drop to 0.

That is not to say that the other results we have shown in this chapter have no interest. In particular, an important notion we have introduced in this chapter is the notion of finite-choice strategies, with the main result proved on finite-choice strategies being Theorem 3.37. Roughly, Theorem 3.37 states that if a finite-choice strategy achieves a value in a standard finite concurrent game with a payoff function f, then a simple strategy can achieve the same value; where simple means what is required to be optimal in turn-based games with f as payoff function. This holds for various objectives. As stated in Corollary 3.38, for a parity objective, simple means positional. Hence, in a standard finite concurrent parity game, if there is a subgame optimal strategy that is finite-choice, then there is one that is positional. The question then is: can the finite choice assumption be weakened? We know that it cannot be dropped entirely, since, as exemplified by the game depicted in Figure 3.12, subgame optimal strategies may require infinite choice. However, we believe that it may hold if finite choice is replaced by positively bounded.

Open Question 3.1. *Does it hold that in all standard finite concurrent parity games, if there is a subgame optimal strategy that is positively bounded, then there is one that is positional?*

The reason why we think Open Question 3.1 could be answered positively is because the parity objective is a qualitative objective, in the sense that what matters is only what is seen infinitely often, regardless of the frequency (contrary to a mean-payoff objective). In addition, with positively bounded strategies, what occurs infinitely often in the game is what occurs infinitely often in the support of the strategy. Therefore, it seems that what matters with a positively bounded strategy is not the exact probability distribution played, but rather the support of this distribution — though this statement should be taken cautiously. With standard finite local interactions, there are only finitely many different supports; and therefore it may be possible to use the same kind of arguments we used to prove Theorem 3.37.

3.6 Appendix

3.6.1 . Proof of Lemma 3.2

We consider a game $\mathcal{G} = \langle \mathcal{C}, g \rangle$ and we let $f := g_{\mathcal{C}}$.

To prove this lemma, we will need the notion of covering formally defined below in Definition 3.26.

Definition 3.26 (Covering). *For all $n \in \mathbb{N}$, an n-*covering *is a non-empty subset $A \subseteq \Omega_{\mathcal{C}}^{\leq n}$ of finite non-empty paths such that $\uplus_{\pi \in A} \mathsf{Cyl}(\pi) = Q^\omega$ (i.e. the union is disjoint).*

The probability of any covering, given two strategies, is 1. Furthermore, the expected value of any Player-A strategy does not decrease over coverings, against all strategies. This is stated formally in the lemma below.

Lemma 3.39. *Consider any Player-A strategy $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ and Player-B strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$. For all $n \in \mathbb{N}$, n-coverings $A \subseteq \Omega_{\mathcal{C}}^{\leq n}$ and any finite path $\rho \in (Q_{\mathsf{ns}})^+$, we have:*

$$\sum_{\pi \in A} \mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}(\pi) = 1$$

and

$$\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s_A}](\rho_{\mathsf{lt}}) \leq \sum_{\pi \in A} \mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}(\pi) \cdot \chi_{\mathcal{G}^{\mathsf{tl}(\rho \cdot \pi)}}[\mathsf{s}_{\mathsf{A}}^{\mathsf{tl}(\rho_{\mathsf{lt}} \cdot \pi)}]((\rho \cdot \pi)_{\mathsf{lt}})$$

Proof. Let us prove this lemma by induction on $n \in \mathbb{N}$. This straightforwardly holds for $n = 0$ since in that case $A = \{\epsilon\}$. Let us show it for $n = 1$. That is, let us consider some $\rho \in (Q_{\mathsf{ns}})^+$ and a 1-covering $\{\epsilon\} \neq A \subseteq \Omega_{\mathcal{C}}^{\leq 1}$ and a Player-B strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$. Since A is a covering, it must be that $A = Q$. Hence, we have $\sum_{q \in Q} \mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}(q) = \sum_{q \in Q} \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), q \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) = 1$. Now assume towards a contradiction that:

$$\varepsilon := \chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s}_{\mathsf{A}}^{\mathsf{tl}(\rho)}](\rho_{\mathsf{lt}}) - \sum_{q \in Q} \mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}(q) \cdot \chi_{\mathcal{G}^{\rho}}[\mathsf{s}_{\mathsf{A}}^{\rho}](q) > 0$$

Consider any Player-B strategy $\mathsf{s} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$ such that $\mathsf{s}(\rho_{\mathsf{lt}}) = \mathsf{s_B}(\rho)$ and such that, for all $q \in Q$, we have $\mathbb{E}^{\mathsf{s}_\mathsf{A}^\rho, \mathsf{s}^{\rho_{\mathsf{lt}}}}_{\mathcal{C},q}[f^{\rho \cdot q}] \leq \chi_{\mathcal{G}^\rho}[\mathsf{s}_\mathsf{A}^\rho](q) + \frac{\varepsilon}{2}$. Then, we have:

$$\begin{aligned}
\mathbb{E}^{\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho)}, \mathsf{s}}_{\mathcal{C},\rho_{\mathsf{lt}}}[f^\rho] &= \sum_{q \in Q} \mathbb{P}^{\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho)}, \mathsf{s}}_{\mathcal{C},\rho_{\mathsf{lt}}}(q) \cdot \mathbb{E}^{\mathsf{s}_\mathsf{A}^\rho, \mathsf{s}^{\rho_{\mathsf{lt}}}}_{\mathcal{C},q}[f^{\rho \cdot q}] \\
&= \sum_{q \in Q} \mathbb{P}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}(q) \cdot \mathbb{E}^{\mathsf{s}_\mathsf{A}^\rho, \mathsf{s}^{\rho_{\mathsf{lt}}}}_{\mathcal{C},q}[f^{\rho \cdot q}] \\
&\leq \sum_{q \in Q} \mathbb{P}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}(q) \cdot (\chi_{\mathcal{G}^\rho}[\mathsf{s}_\mathsf{A}^\rho](q) + \frac{\varepsilon}{2}) \\
&= \sum_{q \in Q} \mathbb{P}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}(q) \cdot \chi_{\mathcal{G}^\rho}[\mathsf{s}_\mathsf{A}^\rho](q) + \frac{\varepsilon}{2} \\
&= \chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho)}](\rho_{\mathsf{lt}}) - \varepsilon + \frac{\varepsilon}{2} < \chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho)}](\rho_{\mathsf{lt}})
\end{aligned}$$

This in contradiction with the definition of $\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho)}](\rho_{\mathsf{lt}})$. In fact, we have $\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho)}](\rho_{\mathsf{lt}}) \leq \sum_{q \in Q} \mathbb{P}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}(q) \cdot \chi_{\mathcal{G}^\rho}[\mathsf{s}_\mathsf{A}^\rho](q)$, and therefore the property holds for $n = 1$.

Assume now that it holds for all $k \leq n$ for some $n \geq 1$. Consider an $n+1$-covering $A \subseteq \Omega_\mathcal{C}^{\leq n+1}$ and some $\rho \in (Q_{\mathsf{ns}})^+$. Let $X_n := A \cap \Omega_\mathcal{C}^{\leq n} \subseteq A$ and $X_{n+1} := A \setminus X_n$. We let $Y_n := \{\pi_{\leq n-1} \in (Q_{\mathsf{ns}})^n \mid \pi \in X_{n+1}\}$. In fact, the set $A_n := X_n \cup Y_n \subseteq \Omega_\mathcal{C}^{\leq n}$ is an n-covering. Indeed, for all $\rho \in Q^\omega$, either there is some $i \leq n-1$ such that $\rho_{\leq i} \in A$ and therefore $\rho_{\leq i} \in X_n$. Or, since A is a covering, we have $\rho_{\leq n} \in X_{n+1}$, and in that case $\rho_{\leq n-1} \in Y_n$. In any case, we have $\rho \in \cup_{\pi \in A_n} \mathsf{Cyl}(\pi)$. Furthermore, consider some $\pi \neq \pi' \in A_n$. If $\pi, \pi' \in X_n \subseteq A$, we have $\mathsf{Cyl}(\pi) \cap \mathsf{Cyl}(\pi') = \emptyset$ since A is a covering. If $\pi, \pi' \in Y_n$, by definition of Y_n we have $|\pi| = |\pi'|$ and therefore $\mathsf{Cyl}(\pi) \cap \mathsf{Cyl}(\pi') = \emptyset$. Assume now that $\pi \in X_n \subseteq A$ and $\pi' \in Y_n$. Then, $|\pi| \leq |\pi'|$. Furthermore, there is some $q \in Q$ such that we have $\pi' \cdot q \in A$. Hence, it cannot be that $\pi \sqsubseteq \pi'$ since A is a covering. Therefore, we also have $\mathsf{Cyl}(\pi) \cap \mathsf{Cyl}(\pi') = \emptyset$. We can conclude that the set A_n is a covering.

In addition, for all $\pi \in Y_n$ and for all $q \in Q$, we have $\pi \cdot q \in A$ (since A is a covering and there is no prefix of π in A). Furthermore, Q is a 1-covering, therefore we have:

$$\sum_{q \in Q} \mathbb{P}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho \cdot \pi}(q) = 1$$

and

$$\chi_{\mathcal{G}^{\mathsf{tl}(\rho \cdot \pi)}}[\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho \cdot \pi)}](\pi_{\mathsf{lt}}) \leq \sum_{q \in Q} \mathbb{P}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho \cdot \pi}(q) \cdot \chi_{\mathcal{G}^{\rho \cdot \pi}}[\mathsf{s}_\mathsf{A}^{\rho \cdot \pi}](q)$$

Now, for all $\pi \in A \cup A_n$, we let $v(\pi) := \chi_{\mathcal{G}^{\mathsf{tl}(\rho \cdot \pi)}}[\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho \cdot \pi)}](\rho \cdot \pi_{\mathsf{lt}}) \in [0, 1]$. The above equation therefore rewrites, for all $\pi \in Y_n$:

$$v(\pi) \leq \sum_{q \in Q} \mathbb{P}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho \cdot \pi}(q) \cdot v(\pi \cdot q)$$

By applying our induction hypothesis to A_n, we obtain:

$$\begin{aligned}
\sum_{\pi \in A} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) &= \sum_{\pi \in X_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) + \sum_{\pi \in X_{n+1}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \\
&= \sum_{\pi \in X_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) + \sum_{\pi \in Y_n} \sum_{q \in Q} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi \cdot q) \\
&= \sum_{\pi \in X_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) + \sum_{\pi \in Y_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot \left(\sum_{q \in Q} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho \cdot \pi}(q) \right) \\
&= \sum_{\pi \in X_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) + \sum_{\pi \in Y_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \\
&= \sum_{\pi \in A_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) = 1
\end{aligned}$$

Furthermore:

$$\begin{aligned}
\sum_{\pi \in A} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot v(\pi) &= \sum_{\pi \in X_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot v(\pi) + \sum_{\pi \in X_{n+1}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot v(\pi) \\
&= \sum_{\pi \in X_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot v(\pi) + \sum_{\pi \in Y_n} \sum_{q \in Q} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi \cdot q) \cdot v(\pi \cdot q) \\
&= \sum_{\pi \in X_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot v(\pi) + \sum_{\pi \in Y_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot \left(\sum_{q \in Q} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho \cdot \pi}(q) \cdot v(\pi \cdot q) \right) \\
&\geq \sum_{\pi \in X_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot v(\pi) + \sum_{\pi \in Y_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi \cdot q) \cdot v(\pi) \\
&= \sum_{\pi \in A_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot v(\pi) = \sum_{\pi \in A_n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot \chi_{\mathcal{G}^{\mathsf{tl}(\rho \cdot \pi)}}[\mathsf{s_A}^{\mathsf{tl}(\rho \cdot \pi)}](\rho \cdot \pi_{\mathsf{lt}}) \\
&\geq \chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s_A}^{\mathsf{tl}(\rho)}](\rho_{\mathsf{lt}})
\end{aligned}$$

Hence, the inductive hypothesis also holds at index $n+1$. The lemma follows. □

We can now proceed to the proof of Lemma 3.2. The proof contains three parts, that are indicated in bold.

Proof. Let $\varepsilon > 0$ and $\mathsf{s_A} := \mathsf{s}_{\varepsilon,\mathsf{Rst}}$. Consider any finite path $\rho \in \Omega^+_{\mathcal{C}}$ and a Player-B strategy $\mathsf{s_B} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$. If $\rho_{\mathsf{lt}} \in Q_{\mathsf{s}}$, then straightforwardly, the Player-A strategy $\mathsf{s_A}$ is optimal from ρ. Assume now that $\rho \in (Q_{\mathsf{ns}})^+$. For all $\pi \in \Omega^*_{\mathcal{C}}$, we let $\mathsf{NbU}(\pi) := |\{\pi' \in \Omega^*_{\mathcal{C}} \mid \pi' \sqsubset \pi,\ \mathsf{U}_\varepsilon(\rho \cdot \pi') = \rho \cdot \pi'\}|$.

For all $\theta \in \Omega^+_{\mathcal{C}}$, we let $v(\theta) := \chi_{\mathcal{G}^{\mathsf{tl}(\rho \cdot \theta)}}[\mathsf{s}^{\mathsf{Pl}(\mathsf{tl}(\rho \cdot \theta))}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho \cdot \theta))}](\theta_{\mathsf{lt}})$. This expression is complicated but it expresses something simple: consider the last update $x := \mathsf{U}_\varepsilon(\mathsf{tl}(\rho \cdot \theta))$ of the path $\rho \cdot \theta$ disregarding if there is an update at $\rho \cdot \theta$. This value $v(\theta)$ is in fact equal to the value of the residual strategy $\mathsf{s}_{\varepsilon,x}$ after the history $\mathsf{Pl}(\rho \cdot \theta)$.

First part: Let us show by induction on $k \geq 1$ the property $\mathcal{P}(k)$: for all k-coverings $\{\epsilon\} \neq A \subseteq \Omega_{\mathcal{C}}^{\leq k}$, we have:

$$\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{A}](\rho_{\mathsf{lt}}) - 2 \cdot \varepsilon \leq \sum_{\pi \in A} \mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}(\pi) \cdot (v(\pi) - \mathsf{NbU}(\pi) \cdot \varepsilon) \tag{3.1}$$

Let us show $\mathcal{P}(1)$. Let $\{\epsilon\} \neq A \subseteq \Omega_{\mathcal{C}}^{\leq 1}$ be a 1-covering. If there is an update at ρ (i.e. $\mathsf{U}_\varepsilon(\rho) := \rho$), then for all $q \in Q$, we have $\mathsf{NbU}(q) = 1$ and the strategy at ρ is ε-optimal. That is, we have $\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s}_{\varepsilon,\rho}](\rho_{\mathsf{lt}}) \geq \chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{A}](\rho_{\mathsf{lt}}) - \varepsilon$, by definition of the strategy $\mathsf{s}_{\varepsilon,\rho}$. Hence, by Lemma 3.39, we have:

$$\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s}_{\varepsilon,\rho}](\rho_{\mathsf{lt}}) \leq \sum_{q \in Q} \mathbb{P}_{\mathcal{C},\rho_{\mathsf{lt}}}^{\mathsf{s}_{\varepsilon,\rho},\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho)}}(q) \cdot \chi_{\mathcal{G}^{\rho}}[\mathsf{s}_{\varepsilon,\rho}^{\rho_{\mathsf{lt}}}](q) = \sum_{q \in Q} \mathbb{P}_{\mathcal{C},\rho_{\mathsf{lt}}}^{\mathsf{s}_{\varepsilon,\rho},\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho)}}(q) \cdot v(q)$$

Furthermore, since $\mathsf{s_A}(\rho) = \mathsf{s}_{\varepsilon,\rho}(\rho_{\mathsf{lt}})$ and $\mathsf{s_B}(\rho) = \mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho)}(\rho_{\mathsf{lt}})$, it follows that for all $q \in Q$, we have $\mathbb{P}_{\mathcal{C},\rho_{\mathsf{lt}}}^{\mathsf{s}_{\varepsilon,\rho},\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho)}}(q) = \mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}(q)$. Equation 3.1 follows. Similarly, if there is no update at ρ (i.e. if $\mathsf{U}_\varepsilon(\rho) \neq \rho$), we have, for all $q \in Q$, $\mathsf{NbU}(q) = 0$ and, by definition of the update function U_ε, $\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho))}^{\mathsf{Pl}(\mathsf{tl}(\rho))}](\rho_{\mathsf{lt}}) \geq \chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{A}](\rho_{\mathsf{lt}}) - 2 \cdot \varepsilon$. Again, by Lemma 3.39 and as for the previous case, we have:

$$\begin{aligned}
\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho))}^{\mathsf{Pl}(\mathsf{tl}(\rho))}](\rho_{\mathsf{lt}}) &\leq \sum_{q \in Q} \mathbb{P}_{\mathcal{C},\rho_{\mathsf{lt}}}^{\mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho))}^{\mathsf{Pl}(\mathsf{tl}(\rho))},\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho)}}(q) \cdot \chi_{\mathcal{G}^{\rho}}[\mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho))}^{\mathsf{Pl}(\rho)}](q) \\
&= \sum_{q \in Q} \mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}(q) \cdot v(q)
\end{aligned}$$

Indeed, since there was no update at ρ, we have $\mathsf{U}_\varepsilon(\mathsf{tl}(\rho)) = \mathsf{U}_\varepsilon(\rho)$ and $\mathsf{Pl}(\rho) = \mathsf{Pl}(\mathsf{tl}(\rho) \cdot \rho_{\mathsf{lt}})$ and therefore $\mathsf{s_A}(\rho) = \mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho))}(\mathsf{Pl}(\mathsf{tl}(\rho) \cdot \rho_{\mathsf{lt}}))$. Hence, Equation 3.1 follows.

In any case, the property $\mathcal{P}(1)$ holds.

Assume now that $\mathcal{P}(k)$ holds for some $k \geq 1$. Consider a $k+1$-covering $\{\epsilon\} \neq A \subseteq \Omega_{\mathcal{C}}^{\leq k+1}$. The covering we define to apply our induction hypothesis is similar to the one used in the proof of Lemma 3.39. Let $X_k := A \cap Q^{\leq k} \subseteq A$ and $X_{k+1} := A \setminus X_k$. We let $Y_k := \{\pi_{\leq k-1} \in (Q_{\mathsf{ns}})^k \mid \pi \in X_{k+1}\}$.

In fact, the set $A_k := X_k \cup Y_k \subseteq \Omega_{\mathcal{C}}^{\leq k}$ is a k-covering. Indeed, for all $\rho \in Q^\omega$, either there is some $i \leq k-1$ such that $\rho_{\leq i} \in A$ and therefore $\rho_{\leq i} \in X_k$. Or, since A is a covering, we have $\rho_{\leq k} \in X_{k+1}$, and in that case $\rho_{\leq k-1} \in Y_k$. In any case, we have $\rho \in \cup_{\pi \in A_k} \mathsf{Cyl}(\pi)$. Furthermore, consider some $\pi \neq \pi' \in A_k$. If $\pi, \pi' \in X_k \subseteq A$, we have $\mathsf{Cyl}(\pi) \cap \mathsf{Cyl}(\pi') = \emptyset$ since A is a covering. If $\pi, \pi' \in Y_k$, by definition of Y_k we have $|\pi| = |\pi'|$ and therefore $\mathsf{Cyl}(\pi) \cap \mathsf{Cyl}(\pi') = \emptyset$. Assume now that $\pi \in X_k \subseteq A$ and $\pi' \in Y_k$. Then, $|\pi| \leq |\pi'|$. Furthermore, there is some $q \in Q$ such that we have $\pi' \cdot q \in A$. Hence, it cannot be that $\pi \sqsubseteq \pi'$ since A is a covering. Therefore, we also have $\mathsf{Cyl}(\pi) \cap \mathsf{Cyl}(\pi') = \emptyset$. We can conclude that the set A_k is a k-covering.

We can therefore apply $\mathcal{P}(k)$ to it. We have:

$$\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{A}](\rho_{\mathsf{lt}}) - 2 \cdot \varepsilon \leq \sum_{\pi \in A_k} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot (v(\pi) - \mathsf{NbU}(\pi) \cdot \varepsilon)$$

Let $\pi \in A_k$. Let $A_\pi := \{\theta \in Q^* \mid \pi \cdot \theta \in A\}$ be the set of finite paths leading from π to a path in A. Let us show that:

$$v(\pi) \leq \sum_{\theta \in A_\pi} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho\cdot\pi}(\theta) \cdot (v(\pi \cdot \theta) - (\mathsf{NbU}(\pi \cdot \theta) - \mathsf{NbU}(\pi)) \cdot \varepsilon) \tag{3.2}$$

There are two possibilities:

- Either $\pi \in X_k$, in which case $A_\pi = \{\epsilon\}$ and Equation 3.2 straightforwardly holds.
- Or, $\pi \in Y_k$, in which case $A_\pi = Q$, since A is a covering and no prefix of π is in A. There are again two possibilities:
  - Either there is an update at $\rho \cdot \pi \in \Omega^+_{\mathcal{C}}$. That is, we have $\mathsf{U}_\varepsilon(\rho \cdot \pi) = \rho \cdot \pi$. This implies that, for all $q \in Q$, $\mathsf{NbU}(\rho \cdot \pi \cdot q) = \mathsf{NbU}(\rho \cdot \pi) + 1$. In addition, since there is an update at $\rho \cdot \pi$, it means that:

    $$v(\pi) = \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{s}^{\mathsf{Pl}(\mathsf{tl}(\rho\cdot\pi))}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho\cdot\pi))}](\pi_{\mathsf{lt}}) < \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{A}](\pi_{\mathsf{lt}}) - 2 \cdot \varepsilon$$

    Furthermore, we have $\chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{A}](\pi_{\mathsf{lt}}) - \varepsilon \leq \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{s}_{\varepsilon,\rho\cdot\pi}](\pi_{\mathsf{lt}})$, by definition of the strategy $\mathsf{s}_{\varepsilon,\rho\cdot\pi}$. In addition, by Lemma 3.39, we have:

    $$\chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{s}_{\varepsilon,\rho\cdot\pi}](\pi_{\mathsf{lt}}) \leq \sum_{q \in Q} \mathbb{P}^{\mathsf{s}_{\varepsilon,\rho\cdot\pi},\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}_{\mathcal{C},\pi_{\mathsf{lt}}}(q) \cdot \chi_{\mathcal{G}^{\rho\cdot\pi}}[\mathsf{s}^{\pi_{\mathsf{lt}}}_{\varepsilon,\rho\cdot\pi}](q)$$

    with $\chi_{\mathcal{G}^{\rho\cdot\pi}}[\mathsf{s}^{\pi_{\mathsf{lt}}}_{\varepsilon,\rho\cdot\pi}](q) = v(\pi \cdot q)$. Since we have $\mathsf{s_A}(\rho\cdot\pi) = \mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\rho\cdot\pi)}(\mathsf{Pl}(\rho\cdot\pi)) = \mathsf{s}_{\varepsilon,\rho\cdot\pi}(\pi_{\mathsf{lt}})$, it follows that, for all $q \in Q$, we have $\mathbb{P}^{\mathsf{s}_{\varepsilon,\rho\cdot\pi},\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}_{\mathcal{C},\pi_{\mathsf{lt}}}(q) = \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho\cdot\pi}(q)$. Overall, we have:

    $$\begin{aligned} v(\pi) &< \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{A}](\pi_{\mathsf{lt}}) - 2 \cdot \varepsilon \leq \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{s}_{\varepsilon,\rho\cdot\pi}](\pi_{\mathsf{lt}}) - \varepsilon \\ &\leq \sum_{q \in Q} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho\cdot\pi}(q) \cdot \chi_{\mathcal{G}^{\rho\cdot\pi}}[\mathsf{s}^{\pi_{\mathsf{lt}}}_{\varepsilon,\rho\cdot\pi}](q) - \varepsilon = \sum_{q \in Q} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho\cdot\pi}(q) \cdot (v(\pi \cdot q) - \varepsilon) \end{aligned}$$

    We obtain Equation 3.2 since, for all $q \in Q$, we have $\mathsf{NbU}(\pi \cdot q) = \mathsf{NbU}(\pi) + 1$.
  - Or, there is no update at $\rho \cdot \pi$, that is $\mathsf{U}_\varepsilon(\rho \cdot \pi) \neq \rho \cdot \pi$. In that case, for all $q \in Q$, we have $\mathsf{NbU}(\pi \cdot q) = \mathsf{NbU}(\pi)$. By Lemma 3.39, we have:

    $$\begin{aligned} v(\pi) &= \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{s}^{\mathsf{Pl}(\mathsf{tl}(\rho\cdot\pi))}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho\cdot\pi))}](\pi_{\mathsf{lt}}) \\ &\leq \sum_{q \in Q} \mathbb{P}^{\mathsf{s}^{\mathsf{Pl}(\mathsf{tl}(\rho\cdot\pi))}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho\cdot\pi))},\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}_{\mathcal{C},\pi_{\mathsf{lt}}}(q) \cdot \chi_{\mathcal{G}^{\rho\cdot\pi}}[\mathsf{s}^{\mathsf{Pl}(\mathsf{tl}(\rho\cdot\pi))\cdot\pi_{\mathsf{lt}}}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho\cdot\pi))}](q) \end{aligned}$$

Furthermore, we have $\mathsf{U}_\varepsilon(\mathsf{tl}(\rho\cdot\pi)) = \mathsf{U}_\varepsilon(\rho\cdot\pi)$ and therefore $\mathsf{Pl}(\mathsf{tl}(\rho\cdot\pi))\cdot\pi_{\mathsf{lt}} = \mathsf{Pl}(\rho\cdot\pi)$. Hence, for all $q \in Q$, we have $\chi_{\mathcal{G}^{\rho\cdot\pi}}[\mathsf{s}^{\mathsf{Pl}(\mathsf{tl}(\rho\cdot\pi))\cdot\pi_{\mathsf{lt}}}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho\cdot\pi))}](q) = \chi_{\mathcal{G}^{\rho\cdot\pi}}[\mathsf{s}^{\mathsf{Pl}(\rho\cdot\pi)}_{\varepsilon,\mathsf{U}_\varepsilon(\rho\cdot\pi)}](q) = v(\pi\cdot q)$. In addition, we have:

$$\mathsf{s}_\mathsf{A}(\rho\cdot\pi) = \mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\rho\cdot\pi)}(\mathsf{Pl}(\rho\cdot\pi)) = \mathsf{s}^{\mathsf{Pl}(\mathsf{tl}(\rho\cdot\pi))}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho\cdot\pi))}(\pi_{\mathsf{lt}})$$

Hence, for all $q \in Q$, we have $\mathbb{P}^{\mathsf{s}^{\mathsf{Pl}(\mathsf{tl}(\rho\cdot\pi))}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho\cdot\pi))},\mathsf{s}^{\mathsf{tl}(\rho\cdot\pi)}_\mathsf{B}}_{\mathcal{C},\pi_{\mathsf{lt}}}(q) = \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho\cdot\pi}(q)$. That is:

$$v(\pi) \leq \sum_{q\in Q} \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho\cdot\pi}(q)\cdot v(\pi\cdot q)$$

We obtain Equation 3.2 since, for all $q \in Q$, $\mathsf{NbU}(\pi\cdot q) = \mathsf{NbU}(\pi)$.

We have established Equation 3.2 for all $\pi \in A_k$. We can deduce that, for all $\pi \in A_k$:

$$v(\pi) - \mathsf{NbU}(\pi)\cdot\varepsilon \leq \sum_{\theta\in A_\pi} \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho\cdot\pi}(\theta)\cdot(v(\pi\cdot\theta) - \mathsf{NbU}(\pi\cdot\theta)\cdot\varepsilon)$$

Hence:

$$\begin{aligned}\sum_{\pi\in A_k} \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}(\pi)\cdot(v(\pi) - \mathsf{NbU}(\pi)\cdot\varepsilon) &\leq \sum_{\pi\in A_k}\sum_{\theta\in A_\pi} \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}(\pi\cdot\theta)\cdot(v(\pi\cdot\theta) - \mathsf{NbU}(\pi\cdot\theta)\cdot\varepsilon)\\ &= \sum_{\pi\in A} \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}(\pi)\cdot(v(\pi) - \mathsf{NbU}(\pi)\cdot\varepsilon)\end{aligned}$$

Overall, with our induction hypothesis, we obtain:

$$\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{A}](\rho_{\mathsf{lt}}) - 2\cdot\varepsilon \leq \sum_{\pi\in A} \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}(\pi)\cdot(v(\pi) - \mathsf{NbU}(\pi)\cdot\varepsilon)$$

That is, we obtain Equation 3.1, and the property $\mathcal{P}(k+1)$ follows. Therefore, the property $\mathcal{P}(n)$ holds for all $n \in \mathbb{N}$.

Second part: For all $\pi \in \Omega^+_\mathcal{C}$, we let $\mathsf{UAft}(\pi) := \{\theta \in Q^+ \mid \pi\cdot\theta \in \Omega^+_\mathcal{C},\ \mathsf{U}_\varepsilon(\rho\cdot\pi\cdot\theta) = \rho\cdot\pi\cdot\theta\}$ be the finite set of paths for which there is an update after π. Note that, if $\pi_{\mathsf{lt}} \in Q_\mathsf{s}$, we have $\mathsf{UAft}(\pi) = \emptyset$. Let us show the equation below:

$$v(\pi) \leq \mathbb{E}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho\cdot\pi}[f^{\rho\cdot\pi}\cdot\mathbb{1}_{(Q^+\setminus\mathsf{UAft}(\pi))^\omega}] + \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho\cdot\pi}[\mathsf{UAft}(\pi)] \tag{3.3}$$

Let $\pi \in \Omega^+_\mathcal{C}$. If $\pi_{\mathsf{lt}} \in Q_\mathsf{s}$, we have $v(\pi) = \mathsf{val}(\pi_{\mathsf{lt}}) = \mathbb{E}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho\cdot\pi}[f^{\rho\cdot\pi}\cdot\mathbb{1}_{(Q^+\setminus\mathsf{UAft}(\pi))^\omega}]$ with $\mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho\cdot\pi}[\mathsf{UAft}(\pi)] = 0$. Hence, the equation holds. Assume now that $\pi_{\mathsf{lt}} \notin Q_\mathsf{s}$, i.e. $\pi \in (Q_{\mathsf{ns}})^+$. Let $\mathsf{s}_\pi := \mathsf{s}^{\mathsf{tl}(\mathsf{Pl}(\rho\cdot\pi))}_{\varepsilon,\mathsf{U}_\varepsilon(\rho\cdot\pi)}$. For all $\theta \in \Omega^*_\mathcal{C} \setminus \mathsf{UAft}(\pi)\cdot Q^*$, we have $\mathsf{U}_\varepsilon(\rho\cdot\pi\cdot\theta) = \mathsf{U}_\varepsilon(\rho\cdot\pi)$ and $\mathsf{Pl}(\rho\cdot\pi\cdot\theta) = \mathsf{Pl}(\rho\cdot\pi)\cdot\theta$. Hence,

$\mathsf{s}_\pi(\pi_{\mathsf{lt}} \cdot \theta) = \mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\rho\cdot\pi)}(\mathsf{Pl}(\rho \cdot \pi \cdot \theta)) = \mathsf{s}_\mathsf{A}(\rho \cdot \pi \cdot \theta) = \mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho\cdot\pi)}(\pi_{\mathsf{lt}} \cdot \theta)$. Therefore, we have:

$$\mathbb{E}_{\mathcal{C},\pi_{\mathsf{lt}}}^{\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho\cdot\pi)},\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}[f^{\rho\cdot\pi} \cdot \mathbb{1}_{(Q^*\setminus\mathsf{UAft}(\pi))^\omega}] = \mathbb{E}_{\mathcal{C},\pi_{\mathsf{lt}}}^{\mathsf{s}_\pi,\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}[f^{\rho\cdot\pi} \cdot \mathbb{1}_{(Q^*\setminus\mathsf{UAft}(\pi))^\omega}]$$

and

$$\mathbb{P}_{\mathcal{C},\pi_{\mathsf{lt}}}^{\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho\cdot\pi)},\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{UAft}(\pi)] = \mathbb{P}_{\mathcal{C},\pi_{\mathsf{lt}}}^{\mathsf{s}_\pi,\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{UAft}(\pi)]$$

Hence:

$$\begin{aligned}
\chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{s}_\pi](\pi_{\mathsf{lt}}) &\leq \mathbb{E}_{\mathcal{C},\pi_{\mathsf{lt}}}^{\mathsf{s}_\pi,\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}[f^{\rho\cdot\pi}] \\
&= \mathbb{E}_{\mathcal{C},\pi_{\mathsf{lt}}}^{\mathsf{s}_\pi,\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}[f^{\rho\cdot\pi} \cdot \mathbb{1}_{(Q^*\setminus\mathsf{UAft}(\pi))^\omega}] + \mathbb{E}_{\mathcal{C},\pi_{\mathsf{lt}}}^{\mathsf{s}_\pi,\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}[f^{\rho\cdot\pi} \cdot \mathbb{1}_{\mathsf{UAft}(\pi)}] \\
&\leq \mathbb{E}_{\mathcal{C},\pi_{\mathsf{lt}}}^{\mathsf{s}_\pi,\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}[f^{\rho\cdot\pi} \cdot \mathbb{1}_{(Q^*\setminus\mathsf{UAft}(\pi))^\omega}] + \mathbb{P}_{\mathcal{C},\pi_{\mathsf{lt}}}^{\mathsf{s}_\pi,\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{UAft}(\pi)] \\
&= \mathbb{E}_{\mathcal{C},\pi_{\mathsf{lt}}}^{\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho\cdot\pi)},\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}[f^{\rho\cdot\pi} \cdot \mathbb{1}_{(Q^*\setminus\mathsf{UAft}(\pi))^\omega}] + \mathbb{P}_{\mathcal{C},\pi_{\mathsf{lt}}}^{\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\rho\cdot\pi)},\mathsf{s}_\mathsf{B}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{UAft}(\pi)] \\
&= \mathbb{E}_{\mathcal{C},\rho\cdot\pi}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}[f^{\rho\cdot\pi} \cdot \mathbb{1}_{(Q^+\setminus\mathsf{UAft}(\pi))^\omega}] + \mathbb{P}_{\mathcal{C},\rho\cdot\pi}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}[\mathsf{UAft}(\pi)]
\end{aligned}$$

Now, as before, there are two cases:

- Either, there is an update at $\rho \cdot \pi$. That is, we have $\mathsf{U}_\varepsilon(\rho \cdot \pi) = \rho \cdot \pi$. This means that $v(\pi) = \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho\cdot\pi))}^{\mathsf{Pl}(\mathsf{tl}(\rho\cdot\pi))}](\pi_{\mathsf{lt}}) < \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{A}](\pi_{\mathsf{lt}}) - 2 \cdot \varepsilon$. Furthermore, by definition of the strategy $\mathsf{s}_\pi = \mathsf{s}_{\varepsilon,\rho\cdot\pi}$, we have $\chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{A}](\pi_{\mathsf{lt}}) - \varepsilon \leq \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{s}_\pi](\pi_{\mathsf{lt}})$. Overall, we obtain:

  $$v(\pi) \leq \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{s}_\pi](\pi_{\mathsf{lt}}) \leq \mathbb{E}_{\mathcal{C},\rho\cdot\pi}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}[f^{\rho\cdot\pi} \cdot \mathbb{1}_{(Q^+\setminus\mathsf{UAft}(\pi))^\omega}] + \mathbb{P}_{\mathcal{C},\rho\cdot\pi}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}[\mathsf{UAft}(\pi)]$$

- Or, there is no update at π, $\mathsf{U}_\varepsilon(\mathsf{tl}(\rho \cdot \pi)) = \mathsf{U}_\varepsilon(\rho \cdot \pi)$ and therefore $\mathsf{Pl}(\mathsf{tl}(\rho\cdot\pi)) \cdot \pi_{\mathsf{lt}} = \mathsf{Pl}(\rho\cdot\pi)$. Since we also have $\mathsf{tl}(\mathsf{Pl}(\rho\cdot\pi)) \cdot \pi_{\mathsf{lt}} = \mathsf{Pl}(\rho\cdot\pi)$, it follows that $\mathsf{Pl}(\mathsf{tl}(\rho \cdot \pi)) = \mathsf{tl}(\mathsf{Pl}(\rho \cdot \pi))$. Therefore, $\mathsf{s}_\pi = \mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\rho\cdot\pi)}^{\mathsf{tl}(\mathsf{Pl}(\rho\cdot\pi))} = \mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho\cdot\pi))}^{\mathsf{Pl}(\mathsf{tl}(\rho\cdot\pi))}$. Hence, $v(\pi) = \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{s}_{\varepsilon,\mathsf{U}_\varepsilon(\mathsf{tl}(\rho\cdot\pi))}^{\mathsf{Pl}(\mathsf{tl}(\rho\cdot\pi))}](\pi_{\mathsf{lt}}) = \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{s}_\pi](\pi_{\mathsf{lt}})$. Hence, we have:

  $$v(\pi) = \chi_{\mathcal{G}^{\mathsf{tl}(\rho\cdot\pi)}}[\mathsf{s}_\pi](\pi_{\mathsf{lt}}) \leq \mathbb{E}_{\mathcal{C},\rho\cdot\pi}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}[f^{\rho\cdot\pi} \cdot \mathbb{1}_{(Q^+\setminus\mathsf{UAft}(\pi))^\omega}] + \mathbb{P}_{\mathcal{C},\rho\cdot\pi}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}[\mathsf{UAft}(\pi)]$$

This proves that Equation 3.3 holds for all $\pi \in Q^+$.

Third part: Since, for all $\pi \in \Omega_\mathcal{C}^+$, $v(\pi) \in [0, 1]$ it follows directly from Equation 3.1 that:

$$\lim_{n\to\infty} \mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}[\mathsf{NbU} \geq n] = 0 \tag{3.4}$$

Let $\delta > 0$. We let $\mathsf{NoMrU}_\delta \subseteq \Omega_\mathcal{C}^+$ be the set of finite paths such that with probability at most δ that is another update afterwards. That is: $\mathsf{NoMrU}_\delta := \{\pi \in \Omega_\mathcal{C}^+ \mid \mathbb{P}_{\mathcal{C},\rho\cdot\pi}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}[\mathsf{UAft}(\pi)] \leq \delta\}$. In fact, $\mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}[(Q^*\setminus\mathsf{NoMrU}_\delta)^\omega] = 0$. Indeed, for all $n \in \mathbb{N}$, we have $\mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}[(Q^*\setminus\mathsf{NoMrU}_\delta)^\omega \cap (\mathsf{NbU} \geq n)] = \mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}[(Q^*\setminus\mathsf{NoMrU}_\delta)^\omega]$

since, given $(Q^* \setminus \mathsf{NoMrU}_\delta)^\omega$, there is infinitely often a probability greater than $\delta > 0$ that there is one more update. We can then deduce from Equation 3.4 that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[(Q^* \setminus \mathsf{NoMrU}_\delta)^\omega] = 0$.

Consider now some $n_\delta \in \mathbb{N}$ such that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[\mathsf{NoMrU}_\delta \cap \Omega^{\leq n_\delta}_{\mathcal{C}}] \geq \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[\mathsf{NoMrU}_\delta] - \delta = 1 - \delta$. Let $A \subseteq \mathsf{NoMrU}_\delta \cap \Omega^{\leq n_\delta}_{\mathcal{C}}$ denote the set of finite paths of $\mathsf{NoMrU}_\delta \cap \Omega^{\leq n_\delta}_{\mathcal{C}}$ with no prefix in NoMrU_δ. By definition, we have $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[\mathsf{NoMrU}_\delta \cap \Omega^{\leq n_\delta}_{\mathcal{C}}] = \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[A] \geq 1 - \delta$. We also let $X \subseteq \Omega^{n_\delta}_{\mathcal{C}}$ be the set of finite paths in $\Omega^{n_\delta}_{\mathcal{C}}$ with no prefix in A: $X := \Omega^{n_\delta}_{\mathcal{C}} \setminus (\cup_{\pi \in A} \pi \cdot Q^*)$. By construction, the set $X \cup A$ is a n_δ-covering. We can therefore apply Equation 3.1 to it:

$$\begin{aligned}
\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{A}](\rho_{\mathsf{lt}}) - 2 \cdot \varepsilon &\leq \sum_{\pi \in X \cup A} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot (v(\pi) - \mathsf{NbU}(\pi) \cdot \varepsilon) \\
&\leq \sum_{\pi \in A} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot v(\pi) + (1 - \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[A]) \\
&\leq \sum_{\pi \in A} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot v(\pi) + \delta
\end{aligned}$$

Furthermore, for all $\pi \in A$, we have, by Equation 3.3:

$$v(\pi) \leq \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho\cdot\pi}[f^{\rho\cdot\pi} \cdot \mathbb{1}_{(Q^+ \setminus \mathsf{UAft}(\pi))^\omega}] + \delta \leq \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho\cdot\pi}[f^{\rho\cdot\pi}] + \delta$$

Hence:

$$\begin{aligned}
\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{A}](\rho_{\mathsf{lt}}) - 2 \cdot \varepsilon &\leq \sum_{\pi \in A} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot v(\pi) + \delta \\
&\leq \sum_{\pi \in A} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(\pi) \cdot (\mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho\cdot\pi}[f^{\rho\cdot\pi}] + \delta) + \delta \\
&\leq \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[f^\rho] + 2 \cdot \delta
\end{aligned}$$

As this holds for all $\delta > 0$, it follows that $\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{A}](\rho_{\mathsf{lt}}) - 2 \cdot \varepsilon \leq \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[f^\rho]$. As this holds for all Player-B strategies $\mathsf{s_B} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$ and finite paths $\rho \in (Q_{\mathsf{ns}})^+$, it follows that the Player-A strategy $\mathsf{s_A}$ is subgame 2ε-optimal. □

3.6.2 . Proof of Lemma 3.5

Proof. We prove the result for upward well-founded functions, it is symmetrical for downward well-founded ones. Assume that f is upward well-founded. Let $x \in (0,1]$. For all $n \in \mathbb{N}$, we let $\varepsilon_n := \frac{1}{2^{n+1}} > 0$. If, for all $n \in \mathbb{N}$, we have $[x - \varepsilon_n, x) \cap E \neq \emptyset$, it follows that can build an infinite ascending chain in E, which is not possible by assumption. Hence, there is some $n \in \mathbb{N}$ such that $[x - \varepsilon_n, x) \cap E = \emptyset$.

Let us now prove the other direction. Assume towards a contradiction that there is an infinite ascending chain $(x_n)_{n \in \mathbb{N}}$ in E. Let $x := \sup_{n \in \mathbb{N}} x_n \in [0,1]$. Since the chain is ascending, we have $x > x_n$ for all $n \in \mathbb{N}$. By assumption, there is some $\varepsilon > 0$ such that $[x - \varepsilon, x) \cap E = \emptyset$. That is, for all $n \in \mathbb{N}$, we have

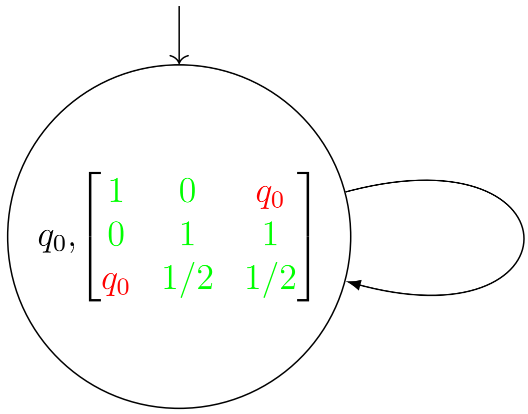

Figure 3.13: A reachability game.

$x_n \leq x - \varepsilon < x$. This is in contradiction with the fact that $x = \sup_{n \in \mathbb{N}} x_n \in [0, 1]$.

□

3.6.3 . Optimal strategy that is locally optimal but not subgame optimal

Consider the game of Figure 3.13: it is a reachability game, that is if it loops indefinitely on q_0, the value is 0. The value of the state q_0 is $1/2$, it is achieved by a Player A positional strategy playing the two top rows with probability $1/2$ and by a Player B positional strategy playing the two leftmost columns with probability $1/2$.

However, denoting a_1, a_2 and a_3 the three actions available to Player A at state q_0 from top to bottom, consider the following Player-A strategy $\mathsf{s_A}$: $\mathsf{s_A}(q_0)(a_1) = \mathsf{s_A}(q_0)(a_2) := 1/2$ and $\mathsf{s_A}(q_0^{n+1})(a_3) := 1$ for all $n \geq 1$. Then, this strategy is locally optimal and it is optimal. Indeed, if the game loops at least once on q_0, then there was the same probability to loop on q_0 and to reach outcome 1. Hence, the mean of the values is at least $1/2$ which is the value of the state q_0. However, it is not subgame optimal since after the game loops once on q_0, then Player B can ensure value 0 by playing indefinitely the left column with probability 1.

3.6.4 . Proof of Proposition 3.9 and Proposition 3.10

In fact, we first prove Proposition 3.10 and then use it to prove Proposition 3.9.

We prove Proposition 3.10.

Proof. Recall Definition 3.3, the valuation $\chi_{\mathcal{G}}[\mathsf{s_A}]^\rho : Q \to [0, 1]$ is such that, for all $q \in Q$, we have $\chi_{\mathcal{G}}[\mathsf{s_A}]^\rho(q) = \chi_{\mathcal{G}^\rho}[\mathsf{s_A^\rho}](q)$. Let $\varepsilon > 0$. Consider a Player-B strategy $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$ such that the GF-strategy $\mathsf{s_B}(\rho) \in \Sigma_\mathsf{B}^{\rho_{\mathsf{lt}}}$ is such that:

$$\mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{s_A}]^\rho \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{s_A}]^\rho \rangle](\mathsf{s_A}(\rho)) + \frac{\varepsilon}{2}$$

and for all $q \in Q$, we have

$$\mathbb{E}_{\mathcal{C},\rho\cdot q}^{\mathsf{s_A},\mathsf{s_B}}[(f_\mathcal{C})^{\rho\cdot q}] \leq \chi_{\mathcal{G}^\rho}[\mathsf{s}_\mathsf{A}^\rho](q) + \frac{\varepsilon}{2} = \chi_\mathcal{G}[\mathsf{s_A}](\rho \cdot q) + \frac{\varepsilon}{2}$$

We have:

$$\begin{aligned}
\mathbb{E}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}[(f_\mathcal{C})^{\rho}] &= \sum_{q\in Q} \mathbb{P}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}(q) \cdot \mathbb{E}_{\mathcal{C},\rho\cdot q}^{\mathsf{s_A},\mathsf{s_B}}[(f_\mathcal{C})^{\rho\cdot q}] \\
&= \sum_{q\in Q} \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), q \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \cdot \mathbb{E}_{\mathcal{C},\rho\cdot q}^{\mathsf{s_A},\mathsf{s_B}}[(f_\mathcal{C})^{\rho\cdot q}] \\
&\leq \sum_{q\in Q} \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), q \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \cdot (\chi_\mathcal{G}[\mathsf{s_A}](\rho \cdot q) + \frac{\varepsilon}{2}) \\
&= \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{s_A}]^\rho \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) + \frac{\varepsilon}{2} \\
&\leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{s_A}]^\rho \rangle](\mathsf{s_A}(\rho)) + \varepsilon
\end{aligned}$$

Since this holds for all $\varepsilon > 0$, it follows that $\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s_A}](\rho_{\mathsf{lt}}) \leq \mathbb{E}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}[(f_\mathcal{C})^{\rho}] \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{s_A}]^\rho \rangle](\mathsf{s_A}(\rho))$. Since $\chi_\mathcal{G}[\mathsf{s_A}]^\rho \leq \chi_\mathcal{G}[\mathsf{A}]^\rho$ by definition of the value, it follows that $\mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{s_A}]^\rho \rangle](\mathsf{s_A}(\rho)) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}]^\rho \rangle](\mathsf{s_A}(\rho))$.

If $\mathsf{s_A}$ is subgame optimal, for all $\rho \in (Q_{\mathsf{ns}})^+$, we have $\chi_\mathcal{G}[\mathsf{A}](\rho) = \chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s_A}](\rho_{\mathsf{lt}}) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}]^\rho \rangle](\mathsf{s_A}(\rho))$, and therefore the Player-A strategy is locally optimal. □

We can now prove Proposition 3.9.

Proof. By Proposition 3.10, for all Player-A strategies $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$, we have $\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s_A}](\rho_{\mathsf{lt}}) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}]^\rho \rangle](\mathsf{s_A}(\rho)) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}]^\rho \rangle](\mathsf{A})$. Therefore, $\chi_\mathcal{G}[\mathsf{A}](\rho) = \chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{A}](\rho_{\mathsf{lt}}) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}]^\rho \rangle](\mathsf{A})$.

Now, let $\varepsilon > 0$. Consider a Player-A strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ such that the GF-strategy $\mathsf{s_A}(\rho) \in \Sigma_\mathsf{A}^{\rho_{\mathsf{lt}}}$ is such that:

$$\mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}]^\rho \rangle](\mathsf{s_A}(\rho)) \geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}]^\rho \rangle](\mathsf{A}) - \frac{\varepsilon}{2}$$

and for all $q \in Q$, we have

$$\chi_{\mathcal{G}^\rho}[\mathsf{s_A}](q) \geq \chi_{\mathcal{G}^\rho}[\mathsf{A}](q) - \frac{\varepsilon}{2} = \chi_\mathcal{G}[\mathsf{A}](\rho \cdot q) - \frac{\varepsilon}{2}$$

For all Player-B strategies $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$, we have:

$$\begin{aligned}
\mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[(f_\mathcal{C})^\rho] &= \sum_{q \in Q} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}(q) \cdot \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho\cdot q}[(f_\mathcal{C})^{\rho\cdot q}] \\
&= \sum_{q \in Q} \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), q\rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \cdot \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho\cdot q}[(f_\mathcal{C})^{\rho\cdot q}] \\
&\geq \sum_{q \in Q} \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), q\rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \cdot (\chi_\mathcal{G}[\mathsf{A}](\rho \cdot q) - \frac{\varepsilon}{2}) \\
&= \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}]^\rho\rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) + \frac{\varepsilon}{2} \\
&\geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}]^\rho\rangle](\mathsf{s_A}) - \varepsilon \\
&\geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}]^\rho\rangle](\mathsf{A}) - \varepsilon
\end{aligned}$$

As this holds for all Player-B strategies $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$, it follows that $\chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s_A}](\rho_{\mathsf{lt}}) \geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}]^\rho\rangle](\mathsf{s_A}(\rho)) - \varepsilon$. As this holds for $\varepsilon > 0$, it follows that $\chi_\mathcal{G}[\mathsf{A}](\rho) = \chi_{\mathcal{G}^{\mathsf{tl}(\rho)}}[\mathsf{s_A}](\rho_{\mathsf{lt}}) \geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}]^\rho\rangle](\mathsf{A})$. □

3.6.5 . Proof of Lemma 3.11

Proof. Coonsider a Player-B strategy $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$ and some $\pi \in \Omega_\mathcal{C}^+$. Let $\rho \in Q^*$. If $\rho \in Q^* \cdot Q_\mathsf{s} \cdot Q^*$, then the inequality straightforwardly holds. Assume now that $\rho \in (Q_{\mathsf{ns}})^*$. We have, by Lemma 1.10 and Definition 1.28:

$$\begin{aligned}
(v_\mathsf{s})^\pi(\rho) = v(\pi \cdot \rho) &\leq \mathsf{val}[\langle \mathsf{F}((\pi \cdot \rho)_{\mathsf{lt}}), v^{\pi\cdot\rho}\rangle](\mathsf{s_A}(\pi \cdot \rho)) \\
&\leq \mathsf{out}[\langle \mathsf{F}((\pi \cdot \rho)_{\mathsf{lt}}), v^{\pi\cdot\rho}\rangle](\mathsf{s_A}(\pi \cdot \rho), \mathsf{s_B}(\pi \cdot \rho)) \\
&= \sum_{q \in Q} \mathsf{out}[\langle \mathsf{F}((\pi \cdot \rho)_{\mathsf{lt}}), q\rangle](\mathsf{s_A}(\pi \cdot \rho), \mathsf{s_B}(\pi \cdot \rho)) \cdot v^{\pi\cdot\rho}(q) \\
&= \sum_{q \in Q} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi\cdot\rho}(q) \cdot v(\pi \cdot \rho \cdot q) \\
&= \sum_{q \in Q} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi\cdot\rho}(q) \cdot (v_\mathsf{s})^\pi(\rho \cdot q)
\end{aligned}$$

That is, the valuation $(v_\mathsf{s})^\pi$ is non-decreasing from π. □

3.6.6 . Proof of Proposition 3.26

We would like to mention that the transformation of priority games into parity games that we used in this proof was already introduced in [63, Corollary 3.8].

Proof. Let $E := f_{\mathsf{Par}}(\mathsf{K}, g)[\mathsf{K}^\omega]$. The function $f_{\mathsf{Par}}(\mathsf{K}, g)$ is straightforwardly upward well-founded since E is finite. Furthermore, it is PI since it only depends on the set of colors seen infinitely often. Now, consider some $\alpha \in [0, 1]$. We have:

$$f_{\mathsf{Par}}(\mathsf{K}, g)^{-1}[[0, \alpha]] = \bigcup_{i \in \mathsf{K},\, g(i) \leq \alpha} \{\max \mathsf{InfOtf}(\rho) = i\} \in \mathsf{Borel}(\mathsf{K})$$

Hence, the function $f_{\mathsf{Par}}(\mathsf{K}, g)$ is measurable.

Let us now define a function $h_\alpha : \mathsf{K} \to \mathbb{N}$ such that, for all $i \in \mathsf{K}$, we have:

$$h_\alpha(i) := \begin{cases} 2 \cdot i & \text{if } g(i) \geq \alpha \\ 2 \cdot i + 1 & \text{otherwise} \end{cases}$$

Then, we have, for all $\rho \in \mathsf{K}^\omega$:

$$f_{\mathsf{Par}}(\mathsf{K}, g)(\rho) \geq \alpha \Leftrightarrow h_\alpha(\mathsf{InfOtf}(\rho)) \in \mathsf{Parity}_{h_\alpha[\mathsf{K}]}$$

Indeed, let $x := \max \mathsf{InfOtf}(\rho)$. First, note that the function h_α is monotone (i.e. for all $i, j \in \mathsf{K}$, we have $i < j$ if and only if $h_\alpha(i) < h_\alpha(j)$). Therefore, $h_\alpha(x) = \max \mathsf{InfOtf}(\rho)$. Hence, we have $f_{\mathsf{Par}}(\mathsf{K}, g)(\rho) \geq \alpha$ iff $g(x) \geq \alpha$ iff $h_\alpha(x)$ is even iff $h_\alpha(\mathsf{InfOtf}(\rho)) \in \mathsf{Parity}_{h_\alpha[\mathsf{K}]}$.

Hence, in finite turn-based games on the set of colors K, playing optimally for the objective $\mathsf{Parity}_{h_\alpha[\mathsf{K}]}$ where all colors $i \in \mathsf{K}$ are replaced by the color $h_\alpha(i)$ is also playing optimally for the objective $\{f \geq \alpha\}$. Therefore both players have positional optimal strategies (this also holds for the objective $\{f \leq \alpha\}$). Hence, the function $f_{\mathsf{Par}}(\mathsf{K}, g) : \mathsf{K}^\omega \to [0, 1]$ is $\mathsf{Std}_f(\mathsf{O})$-PSAW (for both players) since in finite turn-based games with parity objectives, there are always positional optimal strategies [27, 28]. □

3.6.7 . Proof of Theorem 3.28

Proof. Let us denote by $\mathsf{s}_{\mathsf{pb}} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ an optimal positively bounded Player-A strategy generated by Λ. Let us define a Player-A subgame optimal strategy $\mathsf{s}_{\mathsf{Rst}} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$. To do so, we define a map on finite paths $\mathsf{Rst} : Q^+ \to Q^+$ such that for all $q \in Q$, we let $\mathsf{Rst}(q) := q$ and for all $\rho \cdot q \in Q^+$, we let:

$$\mathsf{Rst}(\rho \cdot q) := \begin{cases} \mathsf{Rst}(\rho) \cdot q & \text{if } \chi_{\mathcal{G}}[\mathsf{s}_{\mathsf{pb}}^{\mathsf{Rst}(\rho)}](q) = \chi_{\mathcal{G}}(q) \\ q & \text{otherwise} \end{cases}$$

Note that, this the game is PI, both functions $\chi_{\mathcal{G}} : Q \to [0, 1]$ and $\chi_{\mathcal{G}^\rho} : Q \to [0, 1]$ are the same. Informally, the map Rst resets whenever the strategy s_{pb} is not optimal anymore. We can now define the strategy $\mathsf{s}_{\mathsf{Rst}}$ in the following way, for all $\rho \in Q^+$:

$$\mathsf{s}_{\mathsf{Rst}}(\rho) := \mathsf{s}_{\mathsf{pb}}(\mathsf{Rst}(\rho)) \in \Sigma_{\mathsf{A}}(\rho_{\mathsf{lt}})$$

Since the strategy s_{pb} is positively bounded and generated by Λ, it follows that the strategy $\mathsf{s}_{\mathsf{Rst}}$ also is. Let us show that it is subgame optimal by applying Corollary 3.14.

Let $\rho \in Q^+$. We have $\rho_{\mathsf{lt}} = \mathsf{Rst}(\rho)_{\mathsf{lt}}$. Furthermore, the strategy $\mathsf{s}_{\mathsf{pb}}^{\mathsf{tl}(\mathsf{Rst}(\rho))}$ is optimal from $\mathsf{Rst}(\rho)_{\mathsf{lt}}$. Hence, by Lemma 3.10 — and since the game is PI —

we have:

$$\begin{aligned}\chi_{\mathcal{G}}(\rho_{\mathsf{lt}}) &= \chi_{\mathcal{G}}[\mathsf{s}_{\mathsf{pb}}^{\mathsf{tl}(\mathsf{Rst}(\rho))}](\rho_{\mathsf{lt}}) \\ &\leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_{\mathcal{G}}\rangle](\mathsf{s}_{\mathsf{pb}}(\mathsf{Rst}(\rho))) \\ &= \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_{\mathcal{G}}\rangle](\mathsf{s}_{\mathsf{Rst}}(\rho))\end{aligned}$$

In fact, the strategy $\mathsf{s}_{\mathsf{Rst}}$ is locally optimal.

Let us now show that it ensures the second property of Corollary 3.14. Let $\rho \in (Q_{\mathsf{ns}})^+$ and consider a Player-B deterministic strategy $\mathsf{s}_{\mathsf{B}} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$. For all $\pi \in Q^*$, we denote by $\mathsf{s}_{\mathsf{A}}^{\pi}$ the residual strategy $\mathsf{s}_{\mathsf{Rst}}^{\mathsf{tl}(\rho\cdot\pi)}$ and by $\mathsf{s}_{\mathsf{B}}^{\pi}$ the residual strategy $\mathsf{s}_{\mathsf{B}}^{\mathsf{tl}(\rho\cdot\pi)}$. We also denote $\mathbb{P}^{\mathsf{s}_{\mathsf{Rst}},\mathsf{s}_{\mathsf{B}}}_{\mathcal{C},\rho\cdot\pi}$ by $\mathbb{P}^{\pi}$ — when $\pi = \epsilon$, we omit it. Consider some value $u \in V^{\mathcal{G}} \setminus \{0\}$. We introduce two notations:

- we denote by $\mathsf{MayEx}_u \subseteq Q^+$ the set of finite paths ending in Q_u with a positive probability to exit this value slice: $\mathsf{MayEx}_u := \{\pi \in Q^* \cdot Q_u \mid \mathbb{P}^{\pi}[Q \setminus Q_u] > 0\}$.
- we also denote by $\mathsf{Deviate} \subseteq Q^+$ the set of finite paths where the strategy s_{pb} is not optimal: $\mathsf{Deviate} := \{\pi \in Q^+ \mid \mathsf{Rst}(\rho \cdot \pi) = \pi_{\mathsf{lt}}\}$.

Let us show the three following facts:

(a). $\mathbb{P}[Q^* \cdot (Q_u)^\omega \cap (Q^* \cdot \mathsf{MayEx}_u)^\omega] = 0$;

(b). $\mathbb{P}[Q^* \cdot (Q_u)^\omega \cap Q^* \cdot (Q \setminus \mathsf{MayEx}_u)^\omega] \leq \mathbb{P}^{\mathsf{s},\mathsf{s}_{\mathsf{B}}}[Q^* \cdot (Q_u)^\omega \cap Q^* \cdot (Q \setminus \mathsf{Deviate})^\omega]$;

(c). $\mathbb{P}[Q^* \cdot (Q_u \setminus \mathsf{Deviate})^\omega] = \mathbb{P}[\{f \geq u\} \cap Q^* \cdot (Q_u \setminus \mathsf{Deviate})^\omega]$.

If we assume that all these facts hold, then we obtain:

$$\begin{aligned}\mathbb{P}[Q^* \cdot (Q_u)^\omega] &= \mathbb{P}[Q^* \cdot (Q_u)^\omega \cap Q^* \cdot (Q \setminus \mathsf{MayEx}_u)^\omega] && \text{by fact (a)} \\ &\leq \mathbb{P}[Q^* \cdot (Q_u)^\omega \cap Q^* \cdot (Q \setminus \mathsf{Deviate})^\omega] && \text{by fact (b)} \\ &= \mathbb{P}[Q^* \cdot (Q_u \setminus \mathsf{Deviate})^\omega] \\ &= \mathbb{P}[\{f \geq u\} \cap Q^* \cdot (Q_u \setminus \mathsf{Deviate})^\omega] && \text{by fact (c)} \\ &\leq \mathbb{P}[\{f \geq u\} \cap Q^* \cdot (Q_u)^\omega] \\ &\leq \mathbb{P}[Q^* \cdot (Q_u)^\omega]\end{aligned}$$

In fact, all these inequalities are equalities. We can then apply Corollary 3.14 to conclude. Let us now show all these facts one by one.

(a). Consider some $\pi \in \mathsf{MayEx}_u$. We have $\mathbb{P}^{\pi}[Q \setminus Q_u] > 0$. Let $b := \mathsf{s}_{\mathsf{B}}(\rho \cdot \pi) \in \mathsf{Act}_{\mathsf{B}}^{\pi_{\mathsf{lt}}}$ (recall that s_{B} is a deterministic strategy) and let $A_{Q \setminus Q_u} := \{a \in \mathsf{Act}_{\mathsf{A}}^{\pi_{\mathsf{lt}}} \mid \varrho_{\pi_{\mathsf{lt}}}(a,b)[Q \setminus Q_u] > 0\}$. Then, $\mathsf{s}_{\mathsf{A}}^{\pi}(\pi_{\mathsf{lt}})[A_{Q \setminus Q_u}] > 0$ hence $\mathsf{s}_{\mathsf{A}}^{\pi}(\pi_{\mathsf{lt}})[A_{Q \setminus Q_u}] \geq c$ for some fixed $c > 0$ (since $\mathsf{s}_{\mathsf{Rst}}$ is positively bounded). We let:

$$x := \min_{q \in Q} \min_{(a,b) \in \mathsf{Act}_{\mathsf{A}}^q \times \mathsf{Act}_{\mathsf{B}}^q} \min_{q' \in \mathsf{Sp}(\varrho_q(a,b))} \varrho_q(a,b)(q') > 0$$

We have $\mathbb{P}^{\pi}[Q \setminus Q_u] = \mathsf{out}[\langle \mathsf{F}(\pi_{\mathsf{lt}}), \mathbb{1}_{Q \setminus Q_u} \rangle](\mathsf{s}_{\mathsf{A}}^{\pi}(\pi_{\mathsf{lt}}), b) \geq c \cdot x$. In fact, this holds for all $\pi \in \mathsf{MayEx}_u$. Hence, for all $\pi \in Q^*$, we have $\mathbb{P}^{\pi}[(Q_u)^\omega \mid (Q^* \cdot \mathsf{MayEx}_u)^\omega] \leq \lim_{n \to \infty} (1 - c \cdot x)^n = 0$. It follows that $\mathbb{P}[Q^* \cdot (Q_u)^\omega \cap (Q^* \cdot \mathsf{MayEx}_u)^\omega] = 0$.

(b). Let us show that $\mathbb{P}[Q^* \cdot (Q_u)^\omega \cap Q^* \cdot (Q \setminus \mathsf{MayEx}_u)^\omega \cap (Q^* \cdot \mathsf{Deviate})^\omega] = 0$. Let $\theta \in Q^* \cdot (Q_u)^\omega \cap Q^* \cdot (Q \setminus \mathsf{MayEx}_u)^\omega \cap (Q_{\mathsf{ns}})^\omega$. Let $n \in \mathbb{N}$ be an index such that $\theta_{\geq n} \in (Q_u \setminus \mathsf{MayEx}_u)^\omega$. Consider, assuming it exists, the least index $i \geq n+1$ such that $\theta_i \in \mathsf{Deviate}$. By definition, we have $\chi_{\mathcal{G}}[\mathsf{s}_{\mathsf{pb}}^{\mathsf{Rst}(\rho \cdot \theta_{\leq i-1})}](\theta_i) < \chi_{\mathcal{G}}(\theta_i)$ and $\chi_{\mathcal{G}}[\mathsf{s}_{\mathsf{pb}}^{\mathsf{Rst}(\rho \cdot \theta_{\leq i-2})}](\theta_{i-1}) = \chi_{\mathcal{G}}(\theta_{i-1})$. We let $\varepsilon := \chi_{\mathcal{G}}(\theta_i) - \chi_{\mathcal{G}}[\mathsf{s}_{\mathsf{pb}}^{\mathsf{Rst}(\rho \cdot \theta_{\leq i-1})}](\theta_i) > 0$. Furthermore, by Lemma 3.10 — and since the game is PI — we have, for $b := \mathsf{s}_{\mathsf{B}}(\rho_{\mathsf{lt}} \cdot \theta_{\leq i-1})$ (recall that s_{B} is deterministic):

$$\begin{aligned}
u = \chi_{\mathcal{G}}(\theta_{i-1}) &= \chi_{\mathcal{G}}[\mathsf{s}_{\mathsf{pb}}^{\mathsf{Rst}(\rho \cdot \theta_{\leq i-2})}](\theta_{i-1}) \\
&\leq \mathsf{val}[\langle \mathsf{F}(\theta_{i-1}), \chi_{\mathcal{G}}[\mathsf{s}_{\mathsf{pb}}]^{\mathsf{Rst}(\rho \cdot \theta_{\leq i-2})} \rangle](\mathsf{s}_{\mathsf{pb}}(\mathsf{Rst}(\rho \cdot \theta_{\leq i-1}))) \\
&= \mathsf{val}[\langle \mathsf{F}(\theta_{i-1}), \chi_{\mathcal{G}}[\mathsf{s}_{\mathsf{pb}}]^{\mathsf{Rst}(\rho \cdot \theta_{\leq i-2})} \rangle](\mathsf{s}_{\mathsf{A}}^{\theta_{\leq i-1}}(\theta_{i-1})) \\
&\leq \mathsf{out}[\langle \mathsf{F}(\theta_{i-1}), \chi_{\mathcal{G}}[\mathsf{s}_{\mathsf{pb}}]^{\mathsf{Rst}(\rho \cdot \theta_{\leq i-2})} \rangle](\mathsf{s}_{\mathsf{A}}^{\theta_{\leq i-1}}(\theta_{i-1}), b) \\
&= \sum_{q \in Q} \mathbb{P}^{\theta_{\leq i-1}}(q) \cdot \chi_{\mathcal{G}}[\mathsf{s}_{\mathsf{pb}}]^{\mathsf{Rst}(\rho \cdot \theta_{\leq i-2})}(q) \\
&\leq \sum_{q \in Q} \mathbb{P}^{\theta_{\leq i-1}}(q) \cdot \chi_{\mathcal{G}}(q) - \varepsilon \cdot \mathbb{P}^{\rho \cdot \theta_{\leq i-1}}(\theta_i) \\
&= \mathsf{out}[\langle \mathsf{F}(\theta_{i-1}), \chi_{\mathcal{G}} \rangle](\mathsf{s}_{\mathsf{Rst}}(\rho \cdot \theta_{\leq i-1}), b) - \varepsilon \cdot \mathbb{P}^{\theta_{\leq i-1}}(\theta_i)
\end{aligned}$$

Hence, if $\mathbb{P}^{\theta_{\leq i-1}}[\theta_i] > 0$, we have $\mathsf{out}[\langle \mathsf{F}(\theta_{i-1}), \chi_{\mathcal{G}} \rangle](\mathsf{s}_{\mathsf{A}}^{\theta_{\leq i-1}}(\theta_{i-1}), b) > u$. In that case, at $\theta_{\leq i-1}$, there is a non-zero probability to reach a state of value different from u, i.e. $\mathbb{P}^{\theta_{\leq i-1}}[Q \setminus Q_u] > 0$. That is, $\theta_{i-1} \in \mathsf{MayEx}_u$. That is a path – with a positive probability to occur – that does not visit MayEx_u does not visit $\mathsf{Deviate}$ as well. (Note that, if at some point a stopping state is seen, then $\mathsf{Deviate}$ and MayEx_u will not occur anymore). Hence, almost-surely, a path visiting MayEx_u only finitely often visits $\mathsf{Deviate}$ only finitely often. It follows that $\mathbb{P}[Q^* \cdot (Q_u)^\omega \cap Q^* \cdot (Q \setminus \mathsf{MayEx}_u)^\omega \cap (Q^* \cdot \mathsf{Deviate})^\omega] = 0$. That is:

$$\begin{aligned}
&\mathbb{P}[Q^* \cdot (Q_u)^\omega \cap Q^* \cdot (Q \setminus \mathsf{MayEx}_u)^\omega] \\
&\qquad = \mathbb{P}[Q^* \cdot (Q_u)^\omega \cap Q^* \cdot (Q \setminus \mathsf{MayEx}_u)^\omega \cap Q^* \cdot (Q \setminus \mathsf{Deviate})^\omega] \\
&\qquad \leq \mathbb{P}[Q^* \cdot (Q_u)^\omega \cap Q^* \cdot (Q \setminus \mathsf{Deviate})^\omega]
\end{aligned}$$

(c). Consider any $\pi \in (Q_{\mathsf{ns}})^*$ such that $\pi_{\mathsf{lt}} \in Q_u$ and $\pi \notin \mathsf{Deviate}$. We let $Q_{\mathsf{ch}} \subseteq (Q_{\mathsf{ns}})^+$ be such that $Q_{\mathsf{ch}} := \{\theta \in (Q_{\mathsf{ns}})^+,\ \theta_{\mathsf{lt}} \notin Q_u,\ \text{ or } \pi \cdot \theta \in \mathsf{Deviate}\}$. Let us now define a new game $\mathcal{G}_\pi^{\mathsf{stop}} = \langle \mathcal{C}_\pi^{\mathsf{stop}}, f \rangle$ that behaves exactly like $\mathcal{G}$ (from $\rho\cdot\pi$) as long as we stay outside of Q_{ch}, while any path $\pi \in Q_{\mathsf{ch}}$ is replaced by a stopping state of value $\chi_{\mathcal{G}}[\mathsf{A}](\pi)$. Clearly, for all $\pi \in Q_{\mathsf{ch}}$, we have $\chi_{\mathcal{G}}[\mathsf{A}](\rho \cdot \pi \cdot \theta) = \chi_{\mathcal{G}_\pi^{\mathsf{stop}}}[\mathsf{A}](\theta)$. It follows that all finite paths in $Q^+ \setminus Q_{\mathsf{ch}}$ have the same value u than in the game $\mathcal{G}$. Consider now the strategy $\mathsf{s}_{\mathsf{Rst}}$ in that game (note that it is uniquely defined on Q_{ch}). Clearly, it is optimal from all states in Q_{ch}. Furthermore, by definition of $\mathsf{Deviate}$, the strategy $\mathsf{s}_{\mathsf{Rst}}$ is also optimal after every finite path in $Q^+ \setminus Q_{\mathsf{ch}}$. That is, it is subgame optimal in the game $\mathcal{G}_\pi^{\mathsf{stop}}$. Hence, by Theorem 3.12, in the arena $\mathcal{C}_\pi^{\mathsf{stop}}$, it satisfies that against all Player-B, the probability that $f_{\mathcal{C}}$ is at least u given that the game stays in $Q^+ \setminus Q_{\mathsf{ch}}$ is 1. It follows that: $\mathbb{P}_\pi[(Q_u \setminus \mathsf{Deviate})^\omega] = \mathbb{P}_\pi[\{f \geq u\} \cap (Q_u \setminus \mathsf{Deviate})^\omega]$. Since this holds for all such $\pi \in (Q_{\mathsf{ns}})^*$ such that $\pi_{\mathsf{lt}} \in Q_u$ and $\pi \notin \mathsf{Deviate}$, it follows that: $\mathbb{P}[Q^* \cdot (Q_u \setminus \mathsf{Deviate})^\omega] = \mathbb{P}[Q^* \cdot \{f \geq u\} \cap (Q_u \setminus \mathsf{Deviate})^\omega]$.

Note that we can indeed consider only Player-B deterministic strategies since, once the Player-A strategy is fixed, we obtain an MDP where Player B plays alone. Hence, ε-optimal Player-B strategies can be found among deterministic strategies, by Corollary 2.17. □

3.6.8 . Proof of Lemma 3.29

We first show the lemma below.

Lemma 3.40. *Consider a standard concurrent game $\mathcal{G}$, a collection $\Lambda = (\Lambda_q)_{q\in Q} \in \prod_{q\in Q} \Sigma_{\mathsf{A}}^q$ of sets of Player-A GF-strategies and $\eta : \mathsf{K} \to \mathsf{K}'$ for some set K'. For all pairs of strategies $\mathsf{s}_{\mathsf{A}} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}(\Lambda,\eta)}$ and $\mathsf{s}_{\mathsf{B}} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}(\Lambda,\eta)}$, denoting $\mathcal{T}^{\mathcal{C}(\Lambda,\eta)} := \mathcal{T}_{\mathcal{C}(\Lambda,\eta)}^{\mathsf{s}_{\mathsf{A}},\mathsf{s}_{\mathsf{B}}}$, we have:*

$$\mathbb{E}^{\mathcal{C}(\Lambda,\eta)}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q] = \mathbb{E}^{\mathcal{C}(\Lambda,\eta)}[((f_{\mathcal{C}})_{Q_{\mathsf{A}},Q_{\mathsf{B}}})^q]$$

Proof. The equality straightforwardly holds if $q \in Q_{\mathsf{s}}$. Assume now that $q \in Q_{\mathsf{ns}}$. Since the stochastic tree $\mathcal{T}^{\mathcal{C}(\Lambda,\eta)}$ is $(Q_{\mathsf{A}}, Q_{\mathsf{B}})$-alternating, we have:

$$\mathbb{E}^{\mathcal{C}(\Lambda,\eta)}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q] = \mathbb{E}^{\mathcal{C}(\Lambda,\eta)}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q \cdot \mathbb{1}_{(Q_{\mathsf{B}}\cdot Q_{\mathsf{A}})^\omega}]$$

and

$$\mathbb{E}^{\mathcal{C}(\Lambda,\eta)}[((f_{\mathcal{C}})_{Q_{\mathsf{A}},Q_{\mathsf{B}}})^q] = \mathbb{E}^{\mathcal{C}(\Lambda,\eta)}[((f_{\mathcal{C}})_{Q_{\mathsf{A}},Q_{\mathsf{B}}})^q \cdot \mathbb{1}_{(Q_{\mathsf{B}}\cdot Q_{\mathsf{A}})^\omega}]$$

Furthermore, consider some $\rho \in (Q_{\mathsf{B}} \cdot Q_{\mathsf{A}})^\omega$. If ρ ever reaches a stopping state, denoting $q_{\mathsf{s}} \in Q_{\mathsf{s}} \in Q_{\mathsf{A}}$ (since there is no stopping states in Q_{B}) the first one reached, we have:

$$((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q(\rho) = (f_\eta)_{\mathcal{C}(\Lambda,\eta)}(q\cdot\rho) = \mathsf{val}(q_{\mathsf{s}}) = (f_{\mathcal{C}})_{Q_{\mathsf{A}},Q_{\mathsf{B}}}(q\cdot\rho) = ((f_{\mathcal{C}})_{Q_{\mathsf{A}},Q_{\mathsf{B}}})^q(\rho)$$

Assume now that ρ never reaches any stopping state. We have:

$$\begin{aligned}((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q(\rho) &= (f_\eta)_{\mathcal{C}(\Lambda,\eta)}(q\cdot\rho) = f_\eta \circ (\mathsf{col}^\eta)(q\cdot\rho) \\ &= f_\eta(\mathsf{col}(q)\cdot\eta\circ\mathsf{col}(\rho_0)\cdot\mathsf{col}(\rho_1)\cdot\eta\circ\mathsf{col}(\rho_2)\cdots) \\ &= f(\mathsf{col}(q)\cdot\mathsf{col}(\rho_1)\cdot\mathsf{col}(\rho_3)\cdots) \\ &= f\circ\mathsf{col}^\omega(q\cdot\rho_1\cdot\rho_3\cdots) = f_{\mathcal{C}}\circ\phi_{Q_\mathsf{A},Q_\mathsf{B}}(q\cdot\rho) \\ &= (f_{\mathcal{C}})^q_{Q_\mathsf{A},Q_\mathsf{B}}(\rho)\end{aligned}$$

That is, the functions $((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q$ and $((f_{\mathcal{C}})_{Q_\mathsf{A},Q_\mathsf{B}})^q$ coincide on $(Q_\mathsf{B}\cdot Q_\mathsf{A})^\omega$. Hence, we have

$$\mathbb{E}^{\mathcal{C}(\Lambda,\eta)}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q] = \mathbb{E}^{\mathcal{C}(\Lambda,\eta)}[((f_{\mathcal{C}})_{Q_\mathsf{A},Q_\mathsf{B}})^q]$$

□

We can now proceed to the proof of Lemma 3.29.

Proof. The equality straightforwardly holds if $q \in Q_\mathsf{s}$. Assume now that $q \in Q_\mathsf{ns}$. We want to apply Lemma 1.7. We let $\mathcal{T}^{\mathcal{C}} := \mathcal{T}_{\mathcal{C},q}^{\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s}_\mathsf{A}),\mathsf{Pr}_\mathsf{B}^\Lambda(\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B})}$ and $\mathcal{T}^{\mathcal{C}(\Lambda,\eta)} := \mathcal{T}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C}(\Lambda,\eta),q}$.

For all $\pi \in (Q_\mathsf{A})^*$, we have:

$$\mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}[\cup_{\pi'\in\mathsf{T}(\pi)}\mathsf{Cyl}(\pi')] = \sum_{\pi'\in\mathsf{T}(\pi)\cap(Q_\mathsf{B}\cdot Q_\mathsf{A})^*} \mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}(\pi')$$

Let us show by induction on $n \in \mathbb{N}$ the following property: for all $\pi \in (Q_\mathsf{A})^{\leq n}$, we have:

$$\mathbb{P}_q^{\mathcal{C}}(\pi) = \sum_{\pi'\in\mathsf{T}(\pi)\cap(Q_\mathsf{B}\cdot Q_\mathsf{A})^*} \mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}(\pi') = \mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}(q^{-1}\cdot\theta^\mathsf{A}(\mathsf{s}_\mathsf{A})(q\cdot\pi))$$

where, for all $q\cdot\rho \in (Q_\mathsf{A}\cup Q_\mathsf{B})^*$, we let $q^{-1}\cdot(q\cdot\rho) := \rho$. This straightforwardly holds for $n = 0$. Assume now that this holds for some $n \in \mathbb{N}$. Let $\pi \in (Q_\mathsf{A})^{n+1}$. We have:

$$\mathsf{T}(\pi)\cap(Q_\mathsf{B}\cdot Q_\mathsf{A})^* := \{\rho\cdot q'\cdot\pi_{\mathsf{lt}} \mid \rho\in\mathsf{T}(\mathsf{tl}(\pi))\cap(Q_\mathsf{B}\cdot Q_\mathsf{A})^*,\ q'\in Q_\mathsf{B}\}$$

Hence, letting $p := \sum_{\pi' \in \mathsf{T}(\pi) \cap (Q_\mathsf{B} \cdot Q_\mathsf{A})^*} \mathbb{P}^{\mathcal{C}(\Lambda,\eta)}_q(\pi')$, we have:

$$\begin{aligned}
p = & \sum_{\rho \in \mathsf{T}(\mathsf{tl}(\pi)) \cap (Q_\mathsf{B} \cdot Q_\mathsf{A})^*} \sum_{q' \in Q_\mathsf{B}} \mathbb{P}^{\mathcal{C}(\Lambda,\eta)}_q(\rho \cdot q' \cdot \pi_{\mathsf{lt}}) \\
= & \sum_{\rho \in \mathsf{T}(\mathsf{tl}(\pi)) \cap (Q_\mathsf{B} \cdot Q_\mathsf{A})^*} \mathbb{P}^{\mathcal{C}(\Lambda,\eta)}_q(\rho) \cdot \sum_{q' \in Q_\mathsf{B}} \mathbb{P}^{\mathcal{C}(\Lambda,\eta)}_{q \cdot \rho}(q' \cdot \pi_{\mathsf{lt}}) \\
= & \mathbb{P}^{\mathcal{C}(\Lambda,\eta)}_q(q^{-1} \cdot \theta^\mathsf{A}(\mathsf{s_A})(q \cdot \mathsf{tl}(\pi))) \cdot \sum_{q' \in Q_\mathsf{B}} \mathbb{P}^{\mathcal{C}(\Lambda,\eta)}_{\theta^\mathsf{A}(\mathsf{s_A})(q \cdot \mathsf{tl}(\pi))}(q' \cdot \pi_{\mathsf{lt}}) \\
= & \mathbb{P}^{\mathcal{C}(\Lambda,\eta)}_q(q^{-1} \cdot \theta^\mathsf{A}(\mathsf{s_A})(q \cdot \mathsf{tl}(\pi))) \cdot \mathbb{P}^{\mathcal{C}(\Lambda,\eta)}_{\theta^\mathsf{A}(q \cdot \mathsf{tl}(\pi))}((\mathsf{tl}(\pi)_{\mathsf{lt}}, \mathsf{s_A}(\theta^\mathsf{A}(\mathsf{s_A})(q \cdot \mathsf{tl}(\pi)))) \cdot \pi_{\mathsf{lt}}) \\
(= & \mathbb{P}^{\mathcal{C}(\Lambda,\eta)}_q(q^{-1} \cdot \theta^\mathsf{A}(\mathsf{s_A})(q \cdot \pi))\ ) \\
= & \mathbb{P}^{\mathcal{C}(\Lambda,\eta)}_q(q^{-1} \cdot \theta^\mathsf{A}(\mathsf{s_A})(q \cdot \mathsf{tl}(\pi))) \cdot \mathsf{out}[\langle \mathsf{F}(\mathsf{tl}(\pi)_{\mathsf{lt}}), \pi_{\mathsf{lt}} \rangle](\mathsf{s_A}(\theta^\mathsf{A}(\mathsf{s_A})(q \cdot \mathsf{tl}(\pi))), \mathsf{s_B}(\theta^\mathsf{B}(\mathsf{s_A})(q \cdot \mathsf{tl}(\pi)))) \\
= & \mathbb{P}^{\mathcal{C}(\Lambda,\eta)}_q(q^{-1} \cdot \theta^\mathsf{A}(\mathsf{s_A})(q \cdot \mathsf{tl}(\pi))) \cdot \mathsf{out}[\langle \mathsf{F}(\mathsf{tl}(\pi)_{\mathsf{lt}}), \pi_{\mathsf{lt}} \rangle](\mathsf{Pr}^\Lambda_\mathsf{A}(\mathsf{s_A})(q \cdot \mathsf{tl}(\pi)), \mathsf{Pr}^\Lambda_\mathsf{A}(\mathsf{s_A}, \mathsf{s_B})(q \cdot \mathsf{tl}(\pi))) \\
= & \mathbb{P}^{\mathcal{C}}_q(\mathsf{tl}(\pi)) \cdot \mathbb{P}^{\mathcal{C}}_{q \cdot \mathsf{tl}(\pi)}(\pi_{\mathsf{lt}}) = \mathbb{P}^{\mathcal{C}}_q(\pi)
\end{aligned}$$

Thus the property holds at $n+1$. Therefore, it holds for all $n \in \mathbb{N}$. Hence, we can apply Lemma 1.7 to obtain that $\mathbb{E}^{\mathcal{C}}[(f_{\mathcal{C}})^q] = \mathbb{E}^{\mathcal{C}(\Lambda,\eta)}[((f_{\mathcal{C}})_{Q_\mathsf{A},Q_\mathsf{B}})^q]$. With Lemma 3.40, it follows that $\mathbb{E}^{\mathcal{C}}[(f_{\mathcal{C}})^q] = \mathbb{E}^{\mathcal{C}(\Lambda,\eta)}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q]$. □

3.6.9 . Proof of Lemma 3.30

Proof. First, for all $\rho \in (Q_\mathsf{A} \cdot Q_\mathsf{B})^* \cdot Q_\mathsf{A}$, we have $\mathsf{s_A}(\mathsf{t_A})(\rho) = \mathsf{t_A}(\rho_0 \cdot \rho_2 \cdots \rho_{\mathsf{lt}}) \in \Lambda_{\rho_{\mathsf{lt}}}$. Hence, the strategy $\mathsf{s_A}(\mathsf{t_A})$ is indeed deterministic. Furthermore, for all $\rho \in Q^+$, we have, for all Player-A deterministic strategies $x_\mathsf{A} \in \mathsf{S}^{\mathcal{C}(\Lambda,k_n)}_\mathsf{A}$ — straightforwardly from the definition — $\phi_{Q_\mathsf{A},Q_\mathsf{B}}(\theta^\mathsf{A}(x_\mathsf{A})(\rho)) = \phi_{Q_\mathsf{A},Q_\mathsf{B}}(\theta^\mathsf{B}(x_\mathsf{A})(\rho)) = \rho$. Hence, for all $\rho \in Q^+$,we have:

$$\mathsf{Pr}^\Lambda_\mathsf{A}(\mathsf{s_A}(\mathsf{t_A}))(\rho) = \mathsf{s_A}(\mathsf{t_A})(\theta^\mathsf{A}(\mathsf{s_A}(\mathsf{t_A}))(\rho)) = \mathsf{t_A} \circ \phi_{Q_\mathsf{A},Q_\mathsf{B}}(\theta^\mathsf{A}(\mathsf{s_A}(\mathsf{t_A}))(\rho)) = \mathsf{t_A}(\rho)$$

Similarly, for any deterministic Player-A strategy $x_\mathsf{A} \in \mathsf{S}^{\mathcal{C}(\Lambda,k_n)}_\mathsf{A}$, we have:

$$\mathsf{Pr}^\Lambda_\mathsf{B}(x_\mathsf{A}, \mathsf{s_B}(\mathsf{t_B}))(\rho) = \mathsf{s_B}(\mathsf{t_B})(\theta^\mathsf{B}(x_\mathsf{A})(\rho)) = \mathsf{t_B} \circ \phi_{Q_\mathsf{A},Q_\mathsf{B}}(\theta^\mathsf{B}(x_\mathsf{A})(\rho)) = \mathsf{t_B}(\rho)$$

We obtain the desired equalities. □

3.6.10 . Proof of Lemma 3.33

Proof. Consider any memory skeleton $\mathsf{M} = \langle M, m_{\text{init}}, \mu \rangle$ and any actions map $\lambda : M \times (Q_\mathsf{A} \cup Q_\mathsf{B}) \to \sum_{q \in Q_\mathsf{A} \cup Q_\mathsf{B}} \Sigma_\mathsf{A}(q)$ that implement the strategy $\mathsf{s_A}$.

Consider any finite $\rho \in Q^+ = Q^+_\mathsf{A}$ with $n := |\rho| \in \mathbb{N}$ and, for all $0 \leq i \leq n-2$, let $k_i := \mathsf{col}(\rho_i)$. Then, for all $0 \leq i \leq n-1$, for all $\sigma_i \in \Lambda_{\rho_i}$, we have:

$$\begin{aligned}
\mathsf{Pr}^{\eta,\Lambda}_\mathsf{A}(\mathsf{s_A})(\rho) &= \mathsf{Pr}^{\eta,\Lambda}_\mathsf{A}(\lambda)(\mathsf{Pr}^{\eta,\Lambda}_\mathsf{A}(\mu)(m_{\text{init}}, \mathsf{col}^*(\mathsf{tl}(\rho))), \rho_{\mathsf{lt}}) \\
&= \lambda(\mu^*(m_{\text{init}}, k_0 \cdot \eta(k_0) \cdots k_{n-2} \cdot \eta(k_{n-2})), \rho_{\mathsf{lt}}) \\
&= \mathsf{s_A}(\rho_0 \cdot (\rho, \sigma_0) \cdots \rho_{n-2} \cdot (\rho_{n-2}, \sigma_{n-2}) \cdot \rho_{n-1})
\end{aligned}$$

Hence, the definition of $\mathsf{Pr}_{\mathsf{A}}^{\eta,\Lambda}(\mathsf{s_A})$ does not depend on the memory skeleton M or on the action map λ chosen to implement $\mathsf{s_A}$, only on the strategy $\mathsf{s_A}$ itself.

Now, consider some $\pi \in Q^*$. If $\pi = \epsilon$, it suffices to consider $\rho := \epsilon$. Consider assume that $\pi \in Q^+$, we let $n := |\pi| \geq 1$, that is $\pi = \pi 0 \cdots \pi_{n-1}$. We consider the finite path $\rho := \pi_0(\pi_0, \sigma_0) \cdots - \pi_{n-1} \cdot (\pi_{n-1}, \sigma_{n-1}) \in (Q_{\mathsf{A}} \cdot Q_{\mathsf{B}})^+$, where for all $0 \leq i \leq n-1$, we have $\sigma_i \in \Lambda_{\pi_i}$. With this choice, we have $\phi_{Q_{\mathsf{A}},Q_{\mathsf{B}}}(\rho) = \pi$. For all $0 \leq i \leq n-1$, we let $k_i := \mathsf{col}(\pi_i)$.

Let us now consider a memory skeleton and action map that implement the strategy $\mathsf{s}_{\mathsf{A}}^{\rho}$. In fact, for all $\theta \in (Q_{\mathsf{A}} \cup Q_{\mathsf{B}})^+$, we have:

$$\begin{aligned}\mathsf{s}_{\mathsf{A}}^{\rho}(\theta) &= \mathsf{s_A}(\rho \cdot \theta) \\ &= \lambda(\mu^*(m_{\text{init}}, (\mathsf{col}^{\eta})^*(\mathsf{tl}(\rho \cdot \theta))), \theta_{\mathsf{lt}}) \\ &= \lambda(\mu^*(\mu^*(m_{\text{init}}, (\mathsf{col}^{\eta})^*(\rho)), \mathsf{tl}(\theta))), \theta_{\mathsf{lt}})\end{aligned}$$

Hence, letting $m := \mu^*(m_{\text{init}}, (\mathsf{col}^{\eta})^*(\rho)) \in M$, the memory skeleton $\langle M, m, \mu \rangle$ and the action map λ implement the strategy $\mathsf{s}_{\mathsf{A}}^{\rho}$. Consider now some $\theta \in Q^+$. Let $k := |\theta| \in \mathbb{N}$ and for all $0 \leq j \leq k-1$, we let $c_j := \mathsf{col}(\theta_j)$. Then, we have:

$$\begin{aligned}(\mathsf{Pr}_{\mathsf{A}}^{\eta,\Lambda}(\mathsf{s_A}))^{\pi}(\theta) &= \mathsf{Pr}_{\mathsf{A}}^{\eta,\Lambda}(\lambda)(\mathsf{Pr}_{\mathsf{A}}^{\eta,\Lambda}(\mu)(m_{\text{init}}, \mathsf{col}^*(\mathsf{tl}(\pi \cdot \theta))), \theta_{\mathsf{lt}}) \\ &= \lambda(\mu^*(m_{\text{init}}, k_0 \cdot \eta(k_0) \cdots k_{n-1} \cdot \eta(k_{n-1}) \cdot c_0 \cdot \eta(c_0) \cdots c_{k-2} \cdot \eta(c_{k-2})), \theta_{\mathsf{lt}}) \\ &= \lambda(\mu^*(m, c_0 \cdot \eta(c_0) \cdots c_{k-2} \cdot \eta(c_{k-2})), \theta_{\mathsf{lt}}) \\ &= \mathsf{Pr}_{\mathsf{A}}^{\eta,\Lambda}(\lambda)(\mathsf{Pr}_{\mathsf{A}}^{\eta,\Lambda}(\mu)(m, \mathsf{col}^*(\mathsf{tl}(\theta))), \theta_{\mathsf{lt}}) \\ &= (\mathsf{Pr}_{\mathsf{A}}^{\eta,\Lambda}(\mathsf{s}_{\mathsf{A}}^{\rho}))(\theta)\end{aligned}$$

We obtain the desired result.

□

3.6.11 . Proof of Lemma 3.34

Proof. This proof is quite similar to the proof of Lemma 3.29.

The equality straightforwardly holds if $q \in Q_{\mathsf{s}}$. Assume now that $q \in Q_{\mathsf{ns}}$. We want to apply Lemma 1.7. We let $\mathcal{T}^{\mathcal{C}} := \mathcal{T}^{\mathsf{Pr}_{\mathsf{A}}^{\Lambda,\eta}(\mathsf{s_A}),\mathsf{t_B}}$ and $\mathcal{T}^{\mathcal{C}(\Lambda,\eta)} := \mathcal{T}_{\mathcal{C}(\Lambda,\eta),q}^{\mathsf{s_A},\mathsf{s_B}(\mathsf{t_B})}$.

We denote by $\mathsf{SeqAlt_A} \subseteq (Q_{\mathsf{B}} \cdot Q_{\mathsf{A}})^+$ the set:

$$\begin{aligned}\mathsf{SeqAlt_A} := \{(q_0, \sigma_{q_0}) \cdot q_1 \cdot (q_1, \sigma_{q_1}) \cdots q_n \mid q_0 = q,\ \forall 1 \leq i \leq n,\ q_i \in Q_{\mathsf{A}}, \\ \forall 0 \leq i \leq n-1,\ \sigma_{q_i} \in \Lambda_{q_i}\}\end{aligned}$$

and by $\mathsf{SeqAlt_B} \subseteq (Q_{\mathsf{B}} \cdot Q_{\mathsf{A}})^* \cdot Q_{\mathsf{B}}$ the set:

$$\mathsf{SeqAlt_B} := \{\rho \cdot (\rho_{\mathsf{lt}}, \sigma) \mid \rho \in \mathsf{SeqAlt_A},\ \sigma \in \Lambda_{\rho_{\mathsf{lt}}}\}$$

For all $\pi \in Q^+$, we let $\mathsf{SeqAlt_A}(\pi) := \{\rho \in \mathsf{SeqAlt_A} \mid \phi_{Q_{\mathsf{A}},Q_{\mathsf{B}}}(\rho) = \pi\}$.

For all $\pi \in Q^*$ and $\rho \in \mathsf{T}(\pi) \cap \mathsf{SeqAlt}_\mathsf{A}$, letting $n := |\rho| \in \mathbb{N}$ and $k_{2i+1} := \mathsf{col}(\rho_{2i+1}) = \mathsf{col}(\pi_i)$ for all $i \geq 0$ such that $2i+1 < n$, we have:

$$\begin{aligned}
\mathsf{s}_\mathsf{A}(q \cdot \rho) &= \lambda(\mu^*(m_{\text{init}}, \mathsf{col}^\eta(\mathsf{tl}(q \cdot \rho))), \rho_{\mathsf{lt}}) \\
&= \lambda(\mu(m_{\text{init}}, \mathsf{col}(q) \cdot \eta(\mathsf{col}(q)) \cdot k_1 \cdot \eta(k_1) \cdots k_{n-3} \cdot \eta(k_{n-3}) \cdot k_{n-1}), \rho_{\mathsf{lt}}) \\
&= \mathsf{Pr}_\mathsf{A}^\Lambda(\lambda)(\mathsf{Pr}_\mathsf{A}^\eta(\mu)(m_{\text{init}}, \mathsf{col}(q) \cdot k_1 \cdots k_{n-3} \cdot k_{n-1}), \rho_{\mathsf{lt}}) \\
&= \mathsf{Pr}_\mathsf{A}^\Lambda(\lambda)(\mathsf{Pr}_\mathsf{A}^\eta(\mu)^*(m_{\text{init}}, \mathsf{col}(q \cdot \pi)), \pi_{\mathsf{lt}}) \\
&= \mathsf{Pr}_\mathsf{A}^{\Lambda,\eta}(\mathsf{s}_\mathsf{A})(q \cdot \pi)
\end{aligned}$$

Similarly, for all $\rho \in \mathsf{T}(\pi) \cap \mathsf{SeqAlt}_\mathsf{B}$, we have:

$$\mathsf{s}_\mathsf{B}(\mathsf{t}_\mathsf{B})(q \cdot \rho) = \mathsf{t}_\mathsf{B} \circ \phi_{Q_\mathsf{A}, Q_\mathsf{B}}(q \cdot \rho) = \mathsf{t}_\mathsf{B}(\pi)$$

Now, for all $\pi \in (Q_\mathsf{A})^*$ and by definition of the arena $\mathcal{C}(\Lambda, \eta)$, we have:

$$\mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}[\cup_{\pi' \in \mathsf{T}(\pi)} \mathsf{Cyl}(\pi')] = \sum_{\pi' \in \mathsf{T}(\pi) \cap \mathsf{SeqAlt}_\mathsf{A}} \mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}(\pi')$$

Let us show by induction on $n \in \mathbb{N}$ the property: for all $\pi \in (Q_\mathsf{A})^{\leq n}$, we have:

$$\mathbb{P}_q^{\mathcal{C}}(\pi) = \sum_{\pi' \in \mathsf{T}(\pi) \cap \mathsf{SeqAlt}_\mathsf{A}} \mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}(\pi')$$

This straightforwardly holds for $n = 0$. Assume now that this holds for some $n \in \mathbb{N}$. Let $\pi \in (Q_\mathsf{A})^{n+1}$. We have:

$$\mathsf{T}(\pi) \cap \mathsf{SeqAlt}_\mathsf{A} := \{\rho \cdot ((q \cdot \rho)_{\mathsf{lt}}, \sigma) \cdot \pi_{\mathsf{lt}} \mid \rho \in \mathsf{T}(\mathsf{tl}(\pi)) \cap \mathsf{SeqAlt}_\mathsf{A},\ \sigma \in \Lambda_{(q \cdot \rho)_{\mathsf{lt}}}\}$$

Hence, letting $p := \sum_{\pi' \in \mathsf{T}(\pi) \cap \mathsf{SeqAlt}_\mathsf{A}} \mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}(\pi')$, we have:

$$\begin{aligned}
p &= \sum_{\rho \in \mathsf{T}(\mathsf{tl}(\pi)) \cap \mathsf{SeqAlt}_\mathsf{A}} \sum_{\sigma \in \Lambda_{(q \cdot \rho)_{\mathsf{lt}}}} \mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}(\rho \cdot ((q \cdot \rho)_{\mathsf{lt}}, \sigma) \cdot \pi_{\mathsf{lt}}) \\
&= \sum_{\rho \in \mathsf{T}(\mathsf{tl}(\pi)) \cap \mathsf{SeqAlt}_\mathsf{A}} \mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}(\rho) \cdot \sum_{\sigma \in \Lambda_{(q \cdot \rho)_{\mathsf{lt}}}} \mathbb{P}_{q \cdot \rho}^{\mathcal{C}(\Lambda,\eta)}(((q \cdot \rho)_{\mathsf{lt}}, \sigma) \cdot \pi_{\mathsf{lt}}) \\
&= \sum_{\rho \in \mathsf{T}(\mathsf{tl}(\pi)) \cap \mathsf{SeqAlt}_\mathsf{A}} \mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}(\rho) \cdot \sum_{\sigma \in \Lambda_{(q \cdot \rho)_{\mathsf{lt}}}} \mathsf{s}_\mathsf{A}(q \cdot \rho)(\sigma) \cdot \mathbb{P}_{q \cdot \rho \cdot ((q \cdot \rho)_{\mathsf{lt}}, \sigma)}^{\mathcal{C}(\Lambda,\eta)}(\pi_{\mathsf{lt}}) \\
&= \sum_{\rho \in \mathsf{T}(\mathsf{tl}(\pi)) \cap \mathsf{SeqAlt}_\mathsf{A}} \mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}(\rho) \cdot \sum_{\sigma \in \Lambda_{(q \cdot \rho)_{\mathsf{lt}}}} \mathsf{s}_\mathsf{A}(q \cdot \rho)(\sigma) \\
&\times \sum_{b \in \mathsf{Act}_\mathsf{B}^{\rho_{\mathsf{lt}}}} \mathsf{s}_\mathsf{B}(\mathsf{t}_\mathsf{B})(q \cdot \rho \cdot ((q \cdot \rho)_{\mathsf{lt}}, \sigma))(b) \cdot \mathbb{E}(\varrho_{(q \cdot \rho)_{\mathsf{lt}}}(\sigma, b))(\pi_{\mathsf{lt}}) \\
&= \sum_{\rho \in \mathsf{T}(\mathsf{tl}(\pi)) \cap \mathsf{SeqAlt}_\mathsf{A}} \mathbb{P}_q^{\mathcal{C}(\Lambda,\eta)}(\rho) \cdot \mathbb{E}(\varrho_{(q \cdot \rho)_{\mathsf{lt}}}(\mathsf{Pr}_\mathsf{A}^{\Lambda,\eta}(\mathsf{s}_\mathsf{A})(q \cdot \mathsf{tl}(\pi)), \mathsf{t}_\mathsf{B}(q \cdot \mathsf{tl}(\pi))))(\pi_{\mathsf{lt}}) \\
&= \mathbb{P}_q^{\mathcal{C}}(\mathsf{tl}(\pi)) \cdot \mathbb{P}_{q \cdot \mathsf{tl}(\pi)}^{\mathcal{C}}(\pi_{\mathsf{lt}}) = \mathbb{P}_q^{\mathcal{C}}(\pi)
\end{aligned}$$

Thus the property holds at $n+1$. Therefore, it holds for all $n \in \mathbb{N}$. Hence, we can apply Lemma 1.7 to obtain that $\mathbb{E}^{\mathcal{C}}[(f_{\mathcal{C}})^q] = \mathbb{E}^{\mathcal{C}(\Lambda,\eta)}[((f_{\mathcal{C}})_{Q_\mathsf{A}, Q_\mathsf{B}})^q]$. With Lemma 3.40, it follows that $\mathbb{E}^{\mathcal{C}}[(f_{\mathcal{C}})^q] = \mathbb{E}^{\mathcal{C}(\Lambda,\eta)}[((f_\eta)_{\mathcal{C}(\Lambda,\eta)})^q]$. □

Part II

Concurrent parity games

Contrary to the previous part in which we studied concurrent games with arbitrary payoff functions (into $[0, 1]$), in this part, we focus on the existence and nature of optimal strategies in finite-state parity games. The situation in turn-based games is rather simple to describe: in such games, both players have positional optimal strategies [27, 28]. Recall that these positional strategies are also subgame optimal since parity objectives are prefix-independent. The situation in the more general setting of concurrent parity games is much more heterogeneous, depending on the exact parity objective considered, i.e. on the number of colors involved. We study different questions regarding the existence of, and the simplest type of strategies among which we can find, (subgame and/or ε-)optimal strategies. The goal of this part is to both provide new results and use them to (almost) complete the picture of how concurrent parity games behave. In particular, we intend to give an overview of how concurrent parity games behave, mostly by gathering already-existing results in the literature. However, it may be interesting as a future work to study in more detail what exact kind of strategies are necessary and sufficient to play (subgame or almost) optimally. (More specifically, we could go beyond the infinite-choice class of strategies.)

The results are summarized in Table 3.1. The rows of this table refer to the objectives considered, whereas the columns of this table (except the rightmost one) refer to a property on finite-state concurrent games with these objectives. Specifically, the two leftmost columns 1 and 2 specify which objective is considered, and what type of local interactions we are considering: Max. refers to local interactions which are maximizable w.r.t. Player A, Arb. refers to arbitrary local interactions that are not necessarily maximizable w.r.t. Player A. Then, the four middle columns $(3, 4, 5, 6)$ refer to some properties on the corresponding games. Specifically, the $\exists$ Opt. ? (3) column is a yes-or-no question about whether there always exist optimal strategies. Furthermore, the three columns $(4, 5, 6)$ refer to the nature of the "simplest" strategies which can achieve the requirements of the columns, i.e. being:

4. ε-Opt.: ε-optimal strategy for all positive $\varepsilon > 0$;

5. Optimal: optimal (recall that such strategies are optimal from every state), when it is possible;

6. SubG. Opt.: subgame optimal, when it is possible.

As one can see, the cells in these three columns are filled with either positional — recall, those are strategies that only depend on the current state of the game — or ∞-choice, recall Definition 3.22: these are strategies that, in at least some state of the game, play infinitely many different GF-strategies. Note that, if we restricted the setting to standard games, this duality positional/∞-choice would be a direct consequence of Corollary 3.38. This is not the case here since we consider general games with arbitrary local interactions. In the following, we call positive results the cells with (green) "yes" and "positional". The other results are called negative.

For $k \in \{c, n\}$, we call ^{k}-results the results with ^{k} next to them in Table 3.1. The novel results that we establish in this part are the ^{n}-results. On the other hand, the ^{c}-results are both new and straightforward consequences of ^{n}-results. All the other results of Table 3.1 are straightforward adaptations of results previously existing in the literature. Specifically, they either extend results on standard finite concurrent games to the more general setting of arbitrary finite-state concurrent games; or strengthen the results in the literature stating that no finite-memory strategy has a desirable property, by stating that neither has any finite-choice strategy.

There is one positive already existing result with standard finite game form, the cell $(\mathsf{coBuchi}, 4)$, that we did not extend into the more general case of arbitrary local interaction maximizable w.r.t. Player A, either as a positive or as a negative result. For all other positive results, including those already known when all local interactions are standard finite, we provide complete proofs that hold even with arbitrary local interactions. We also exhibit examples witnessing negative results (most of these examples are already known). Note that all negative results are witnessed by games with only one non-trivial local interaction that is standard. This local interaction is finite for most negative results. However, it is not the case for the three cells (Safety, 4), (co-Büchi, 4) and (co-Büchi, 6) where, in the non-trivial local interaction, Player A has infinitely many actions, while Player B has only finitely many. As can be read in the table, it would not have been possible to witness these results with a standard finite local interaction. Finally, the rightmost column (7) contains a reference to the theorem summarizing the results for the corresponding objective. In the summarizing proof of these theorems at the end of each section, we refer to all previously known results for the corresponding objective.

This part contains two chapters. In Chapter 4, we focus on the safety and reachability objectives whereas in Chapter 5 we deal with Büchi, co-Büchi and parity objectives.

1	GF 2	∃ Opt. ? 3	ε-Opt. 4	Optimal 5	SubG. Opt. 6	Thm. 7
Safety	Max.	Yes	Positional	Positional	Positional	4.5
	Arb.	No	∞-choice			
Reach	M./A.	No	Positional	Positional^n	Positional^c	4.12
Buchi	M./A.	No	∞-choice	Positional^n	Positional^c	5.5
coBuchi	Max.	No	Pos$*$/?	∞-choice^n	Pos^{*n}/∞-choice	5.13
	Arb.		∞-choice		∞-choice	
Parity	M./A.	No	∞-choice	∞-choice	∞-choice	5.15

Table 3.1: A table summarizing the situation in finite-state concurrent games with several objectives where the local interactions are maximizable for Player A (rows 'Max.') and arbitrary (rows 'Arb.'). When there is only one row for an objective, it means that the results are the same whether we assume that the local interactions are maximizable or not, written M./A. The results Pos$*$ hold with standard finite local interactions, but do not pertain (a priori) to arbitrary local interactions maximizable w.r.t. Player A. Finally, ^{n}-results are the novel results proved in this part while the ^{c}-results are new and are straightforward consequences of ^{n}-results.

4 - Safety and Reachability objectives

In this chapter, we focus on safety and reachability games. In [48], the authors study the safety and reachability objectives in infinite MDPs, turn-based and standard concurrent games. In particular, they specify in a more precise way than what we do in this chapter the exact quantity of memory necessary and sufficient to play optimally[1], almost-optimally, etc. On the other hand, we focus almost entirely on finite-state games and we study how the game behaves depending on the type of (non necessarily standard) game forms occurring in the game, i.e. whether they are maximizable or not.

We first focus on safety objectives. In fact, we first study the well-known notion of upper semi-continuous payoff functions, which can be seen as a generalization of safety objectives. More specifically, we characterize such payoff functions (Proposition 4.1) with subgame optimal strategies and local optimality in infinite games. We then use this result to give the complete picture of how arbitrary finite-state concurrent safety games behave (Theorem 4.5).

We then consider reachability objectives in finite-state games. We first show that, in finite-state reachability games, the Player-A value can be computed with a least fixed point, even with arbitrary game forms (Proposition 4.7). We then describe a procedure to distinguish from which states Player A has an optimal strategy, and from which states she does not. This, in turn, gives that, whenever there is an optimal Player-A strategy in finite-state reachability games, there is one that is positional (Theorem 4.11). We are then able to give the complete picture of how arbitrary finite-state concurrent reachability games behave (Theorem 4.12).

4.1 Safety objectives and upper semi-continuous payoff functions

4.1.1 . Upper semi-continuous functions

Before considering safety objectives, we start by considering those payoff functions w.r.t. which Player A has always subgame optimal strategies, when all local interactions are maximizable w.r.t. to Player A. In turns out that these payoff functions are exactly the ones for which, in all arbitrary games, a Player-A strategy is subgame optimal if and only if it is locally optimal. Therefore, let us consider the necessary and sufficient condition for a Player-A strategy to be subgame optimal stated in Corollary 3.14 (for $v := \chi_{\mathcal{G}}[\mathsf{A}]$).

[1]In that article, they consider optimal strategies that may not be optimal from all states whereas we say that a strategy is optimal if it achieves the value of the game from every state of the game.

We are looking for those payoff functions for which the second condition of Corollary 3.14 always holds. Let us first introduce below the notion of limit of a sequence of infinite sequence of colors.

Definition 4.1 (Limit of a sequence of infinite sequences of colors). *Consider a non-empty set* K*, some infinite sequence of colors* $\rho \in \mathsf{K}^\omega$ *and* $(\rho^n)_{n\in\mathbb{N}} \in (\mathsf{K}^\omega)^\mathbb{N}$*. We say that* ρ *is the* limit *of* $(\rho^n)_{n\in\mathbb{N}}$ *if, for all* $k \in \mathbb{N}$*, there is some* $n_k \in \mathbb{N}$ *such that for all* $n \geq n_k$*, we have* $\rho_{\leq k} = (\rho^n)_{\leq k}$*.*

We can now define formally the notion of payoff function we are interested in. These functions f are the ones for which the value $f(\rho)$ of any infinite path ρ that is the limit of $(\rho^n)_{n\in\mathbb{N}}$ is at least $\mathsf{limsup}(f(\rho^n))_{n\in\mathbb{N}}$. This corresponds to the known notion of upper semi-continuous payoff functions. See for instance [64, 65] for examples of use of this notion in game theory. We define it formally below in Definition 4.2.

Definition 4.2 (Upper semi-continuous payoff functions). *Consider a non-empty set of colors* K *and a payoff function* $f : \mathsf{K}^\omega \to [0,1]$*. It is* upper semi-continuous *if, for all* $\rho \in \mathsf{K}^\omega$ *that is the limit of* $(\rho^n)_{n\in\mathbb{N}} \in (\mathsf{K}^\omega)^\mathbb{N}$*, we have* $\mathsf{limsup}(f(\rho^n))_{n\in\mathbb{N}} \leq f(\rho)$*.*

Remark 4.1. *We make two remarks here. First, this notion is incomparable in strength with upward well-foundedness from Definition 3.5. Furthermore, we would recover exactly the same functions if upper semi-continuous payoff functions were defined with* liminf *instead of* limsup[2]*.*

In fact, as formally stated below, upper semi-continuous payoff functions are exactly the payoff functions for which, in all games, subgame optimal strategies are exactly locally optimal strategies. Equivalently, upper semi-continuous payoff functions are exactly the payoff functions for which Player A always has subgame optimal strategies in games with maximizable (w.r.t. Player A) games forms. In fact, both these statements remain true if we only consider MDPs instead of concurrent games. These equivalences hold only when considering games without stopping states with positives values. This constitutes a one-to-two-player lift as the result can be read as follows: if a payoff function behaves properly in all one-player game, then it also does in all two-player games. See for instance [66] for another (much stronger) one-to-two-player lift.

[2]The reason why is, given a sequence $(\rho^n)_{n\in\mathbb{N}} \in (\mathsf{K}^\omega)^\mathbb{N}$, we could extract a subsequence $(\rho^{\varphi(n)})_{n\in\mathbb{N}}$ for some increasing $\varphi : \mathbb{N} \to \mathbb{N}$ such that $\mathsf{limsup}(f(\rho^n))_{n\in\mathbb{N}} = \lim(f(\rho^{\varphi(n)}))_{n\in\mathbb{N}} = \mathsf{liminf}(f(\rho^{\varphi(n)}))_{n\in\mathbb{N}}$.

Proposition 4.1. *Consider a set of colors* K *and a payoff function* $f : \mathsf{K}^\omega \to [0,1]$*. The assertions below are equivalent, assuming that the games and MDPs we consider are without stopping states with positive value:*

a. *the function* f *is upper semi-continuous;*

b. *in all arbitrary concurrent games* $\langle \mathcal{C}, f \rangle$*, Player* A *subgame optimal strategies coincide with locally optimal strategies;*

c. *in all arbitrary MDPs* $\langle \mathcal{C}, f \rangle$*, Player* A *subgame optimal strategies coincide with locally optimal strategies;*

d. *in all arbitrary concurrent games* $\langle \mathcal{C}, f \rangle$ *maximizable w.r.t. to Player* A*, Player* A *has subgame optimal strategies;*

e. *in all arbitrary MDPs* $\langle \mathcal{C}, f \rangle$ *maximizable w.r.t. to Player* A*, Player* A *has subgame optimal strategies.*

As a side remark, the equivalences b.$\Leftrightarrow$ c. and d.$\Leftrightarrow$ e. provide a one-to-two-player lift. That is, one can infer properties on two-player games from properties on one-player games (i.e. MDPs).

Proof. Let us first show the implication $a. \Rightarrow b.$. Assume that f is upper semi-continuous. Consider a concurrent game $\mathcal{G} = \langle \mathcal{C}, f \rangle$ without stopping states of positive value (i.e. all stopping states have value 0) and a Player-A strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$. By Proposition 3.10, if $\mathsf{s_A}$ is subgame optimal, then it is also locally optimal. Assume now that $\mathsf{s_A}$ is locally optimal. We want to apply Theorem 3.12 to show that $\mathsf{s_A}$ is subgame optimal. Consider an infinite path $\rho \in Q^\omega$. If it ever reaches a stopping state q, then we have $\mathsf{limsup}_{\chi_\mathcal{G}[\mathsf{A}]}(\rho) = \mathsf{val}(q) = f_\mathcal{C}(\rho)$. Hence, we do have $f_\mathcal{C}(\rho) \geq \mathsf{limsup}_{\chi_\mathcal{G}[\mathsf{A}]}(\rho)$. Assume now that ρ does not visit any stopping state. Let $u := \mathsf{limsup}_{\chi_\mathcal{G}[\mathsf{A}]}(\rho)$. If $u = 0$, then straightforwardly $f_\mathcal{C}(\rho) \geq \mathsf{limsup}_{\chi_\mathcal{G}[\mathsf{A}]}(\rho)$. Assume now that $u > 0$. For all $n \in \mathbb{N}$, we let $i_n \in \mathbb{N}$ be such that $\chi_\mathcal{G}[\mathsf{A}](\rho_{\leq i_n}) \geq u - \frac{1}{n}$ (which exists by definition limsup). Let $\pi^n := \mathsf{col}^+(\rho_{\leq i_n}) \in \mathsf{K}^+$ be the corresponding finite sequence of colors. Since there is no stopping states of positive values and $\chi_\mathcal{G}[\mathsf{A}](\rho_{\leq i_n}) \geq u - \frac{1}{n}$, it follows that there is some $\theta^n \in \mathsf{K}^\omega$ with $\theta^n \in \mathsf{Cyl}(\pi^n)$ and $f(\theta^n) \geq u - \frac{1}{2n}$. Consider then the sequence $(\theta^n)_{n \in \mathbb{N}} \in (\mathsf{K}^\omega)^\mathbb{N}$ of infinite sequences of colors. By construction, we have $\mathsf{limsup}(f(\theta^n))_{n \in \mathbb{N}} \geq u$ and $\mathsf{col}^\omega(\rho)$ is equal to the limit of $(\theta_n)_{n \in \mathbb{N}} \in (\mathsf{K}^\omega)^\mathbb{N}$. Hence, since f is upper semi-continuous, it follows that $f_\mathcal{C}(\rho) = f \circ \mathsf{col}^\omega(\rho) \geq u$. Since this holds for all positive $u \in (0,1]$, it follows that $\mathsf{limsup}_{\chi_\mathcal{G}[\mathsf{A}]}(\rho) \leq f_\mathcal{C}(\rho)$, which holds for all $\rho \in Q^\omega$. Hence, the first and the second conditions of Theorem 3.12 are ensured. Therefore, the Player-A strategy $\mathsf{s_A}$ is subgame optimal.

Clearly, we then have $b. \Rightarrow c.$ We also have $b. \Rightarrow d.$ Indeed, consider any game$\langle \mathcal{C}, f \rangle$ maximizable w.r.t. Player A without stopping states of positive value. Then, Player A has a locally optimal strategy: it amounts to play opti-

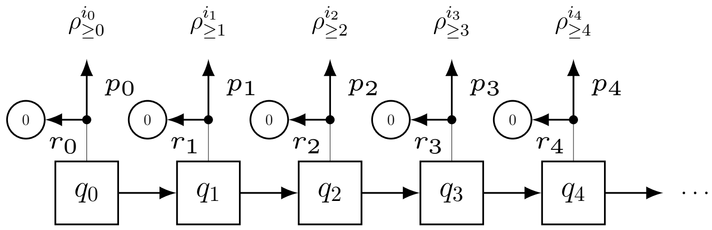

Figure 4.1: An MDP where Player A plays alone, with $r_i := 1 - p_i$.

mally at each local interaction, which is possible since they are all maximizable w.r.t. Player A. Then, by assumption *b.*, the Player-A locally optimal strategy is subgame optimal. Furthermore, we also have straightforwardly that $d. \Rightarrow e.$.

Let us now show that $c. \Rightarrow a.$ and $e. \Rightarrow a.$, the construction is the same for both points. Assume that f is not upper semi-continuous. Hence, there is some $\rho \in \mathsf{K}^\omega$ and $(\rho^n)_{n\in\mathbb{N}} \in (\mathsf{K}^\omega)^\mathbb{N}$ such that ρ is the limit of $(\rho^n)_{n\in\mathbb{N}}$ and $\mathsf{limsup}(f(\rho^n))_{n\in\mathbb{N}} > f(\rho)$. Let $\delta := \mathsf{limsup}(f(\rho^n))_{n\in\mathbb{N}} - f(\rho) > 0$. For all $n \in \mathbb{N}$, we let $i_n \in \mathbb{N}$ be such that $\rho^{i_n}_{\leq n} = \rho_{\leq n}$ and $f(\rho^{i_n}) \geq f(\rho) + \delta/2$. We then define an MDP $\Gamma = \langle \mathcal{C}, f\rangle$ on the set of colors K where Player A plays alone. In that MDP, there is an infinite chain of states $(q_n)_{n\in\mathbb{N}}$ such that, for all $n \in \mathbb{N}$, we have $\mathsf{col}(q_n) := \rho_n$. Furthermore, for all $n \in \mathbb{N}$, at state q_n, Player A has two available actions, one, called a_n^{cont}, that makes the game continue on the chain to q_{n+1} with probability 1 and another one, called a_n^{stop}, that visits with probability $p_n \in [0, 1]$ (not yet defined) the infinite path $\rho^{i_n}_{n+1} \cdot \rho^{i_n}_{n+2} \cdots \in \mathsf{K}^\omega$ with probability p_n and that goes to a sink state of value 0 with probability $1 - p_n$. The probability $p_n \in [0, 1]$ is chosen such that $p_n \cdot f(\rho^{i_n}) = f(\rho) + \frac{\delta}{2} - \frac{1}{2^n}$. An illustration of this game is given in Figure 4.1. Then, the value of all states q_n for $n \in \mathbb{N}$ is equal to $f(\rho) + \delta/2$. Indeed, for all $N \in \mathbb{N}$, from any state $q_n \in Q$, Player A can, with a deterministic strategy, play the actions a_k^{cont} for $k \leq N - 1$ steps until reaching the state q_N from which she can play the action a_N^{stop} to ensure (at least) the value $f(\rho) + \frac{\delta}{2} - \frac{1}{2^n}$. However, to be locally optimal, a Player-A strategy has to play, for all $n \in \mathbb{N}$, deterministically the actions a_n^{cont} after the sequence $q_0 \cdots q_n$. However, the value of such a strategy is equal to $f(\rho) < f(\rho) + \delta/2$. That is, no locally optimal strategy is (subgame) optimal. Hence, both *c.* and *e.* do not hold. □

4.1.2 . Safety objectives

Before considering the safety objective, we would like to mention that, in all the games we will consider later on in this chapter (i.e. the subsequent safety and reachability games), we will not consider any stopping state. The reason for that is twofold: first, since we want to study specific objectives, we do not want stopping states to interfere with how an objectives behaves — for instance, in a safety game, once the target is reached, the game should have value 0, which may not happen if a stopping state is reached subsequently. Second, stopping states can be straightforwardly implemented with Nature states and two self-looping states, one of value 1 and the other of value 0, so this is without loss of generality.

Let us now come back to the safety objective. In fact, the associated payoff function is upper semi-continuous. This is a general property ensured by win/lose payoff functions whose winning set is closed. This is stated in the lemma below.

Lemma 4.2. *Consider a non-empty set of colors* K *and win/lose payoff function* $f : \mathsf{K}^\omega \to \{0, 1\}$. *It is upper semi-continuous if and only if the set* $f^{-1}[\{1\}]$ *is closed.*

Proof. We let $W_1 := f^{-1}[\{1\}]$ and $W_0 := f^{-1}[\{0\}]$. Now, recall that the set W_1 is closed if and only if its complement W_0 is open, that is the set W_0 can be written as an arbitrary union of cylinders.

Let us assume that W_0 is open. That is, there is some set $A \subseteq \mathsf{K}^*$ such that $W_0 = \cup_{\pi \in A} \mathsf{Cyl}(\pi)$. Consider an infinite path $\rho \in \mathsf{K}^\omega$ and a sequence $(\rho^n)_{n \in \mathbb{N}} \in (\mathsf{K}^\omega)^\mathbb{N}$ such that ρ is the limit of $(\rho^n)_{n \in \mathbb{N}}$. If $f(\rho) = 1$, then we have $\mathsf{limsup}(f(\rho^n))_{n \in \mathbb{N}} \leq f(\rho)$. Assume now that $f(\rho) = 0$, that is $\rho \in W_0$. Then, there is some $\pi \in A$ such that $\rho \in \mathsf{Cyl}(\pi)$. Hence, there is some $k \in \mathbb{N}$ such that, for all $n \geq k$, we have $\rho^n \in \mathsf{Cyl}(\pi) \subseteq W_0$. Therefore, $\mathsf{limsup}(f(\rho^n))_{n \in \mathbb{N}} = 0 \leq f(\rho)$. Therefore, f is upper semi-continuous.

Assume now that f is upper semi-continuous. We let $A := \{\pi \in \mathsf{K}^* \mid \mathsf{Cyl}(\pi) \subseteq W_0\}$. We claim that $\cup_{\pi \in A} \mathsf{Cyl}(\pi) = W_0$. By definition, we have $\cup_{\pi \in A} \mathsf{Cyl}(\pi) \subseteq W_0$. Consider now some $\rho \in W_0$. Assume towards a contradiction that, for all $n \in \mathbb{N}$, the finite path $\rho_{\leq n} \notin A$. Then, for all $n \in \mathbb{N}$, there is an infinite path $\theta^n \in \mathsf{Cyl}(\rho_{\leq n})$ such that $\theta^n \notin W_0$, that is such that $f(\theta^n) = 1$. Then, the infinite path ρ is the limit of the sequence $(\theta^n)_{n \in \mathbb{N}} \in (\mathsf{K}^\omega)^\mathbb{N}$ and $\mathsf{limsup}(f(\theta^n))_{n \in \mathbb{N}} = 1$. Hence, since f is upper semi-continuous, $f(\rho) = 1$ and $\rho \notin W_0$. Hence the contradiction. In fact, there is some $n \in \mathbb{N}$ such that $\rho_{\leq n} \in A$. That is, $\rho \in \cup_{\pi \in A} \mathsf{Cyl}(\pi)$. In fact, $W_0 = \cup_{\pi \in A} \mathsf{Cyl}(\pi)$, and it is therefore an open set. □

We therefore deduce as a corollary of what is done in the previous subsection that in all safety games (without stopping states) where each local interactions are maximizable w.r.t. Player A, Player A has a subgame optimal strategy. Moreover, this strategy can be chosen positional. This is stated in the corollary below.

Corollary 4.3. *Consider an arbitrary concurrent safety game* $\mathcal{G}$ *— whose set of states need not be finite — without stopping states. Let* $T := \mathsf{col}^{-1}[\{1\}] \subseteq Q$ *be the target that Player* A *wants to avoid. Assume either that all local interactions outside of* T *are maximizable w.r.t. Player* A*, or that there is a*

Player A *has an optimal strategy. Then, Player* A *has a positional strategy that is subgame optimal in* $\mathcal{G}$.

Proof. Since the reachability winning set is an open set — as it can be written as the union of the cylinders of paths that reach the target — the safety winning set is a closed set. Therefore, by Lemma 4.2 and Proposition 4.1, all Player-A locally optimal strategies are subgame optimal. Consider the valuation $\chi_{\mathcal{G}}[\mathsf{A}] : Q^+ \to [0,1]$. Let us argue that there is a positional Player-A strategy $\mathsf{s_A}$ that is locally optimal, i.e. such that, for all $\rho \in Q^+$, we have $\chi_{\mathcal{G}}[\mathsf{A}](\rho_{\mathsf{lt}}) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \chi_{\mathcal{G}}[\mathsf{A}]^\rho \rangle](\mathsf{s_A}(\rho_{\mathsf{lt}}))$. First, note that for all $\rho \in Q^* \cdot T \cdot Q^*$, we have $\chi_{\mathcal{G}}[\mathsf{A}]^\rho : Q \to \{0\}$. In addition, for all $\rho \in (Q \setminus T)^+$, we have $\chi_{\mathcal{G}}[\mathsf{A}]^\rho = \chi_{\mathcal{G}}[\mathsf{A}]$. Second, by Proposition 3.9, we have, for all $q \in Q$, $\chi_{\mathcal{G}}[\mathsf{A}](q) = \mathsf{val}[\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}]^q \rangle][\mathsf{A}]$. If we assume that all local interactions outside T are maximizable w.r.t. Player A, then such a Player-A strategy $\mathsf{s_A}$ does exist (note that it can play arbitrarily at states in T, of value 0). Assume now that Player A has an optimal strategy s'_A in $\mathcal{G}$. Then, by Lemma 3.10, we have, for all $q \in Q$, $\chi_{\mathcal{G}}[\mathsf{A}](q) = \chi_{\mathcal{G}}[\mathsf{s}'_\mathsf{A}](q) \leq \mathsf{val}[\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}]^q \rangle](\mathsf{s}'_\mathsf{A}(q))$. In fact, we have $\chi_{\mathcal{G}}[\mathsf{A}](q) \leq \mathsf{val}[\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle](\mathsf{s}'_\mathsf{A}(q))$. Therefore, a Player-A positional strategy $\mathsf{s_A}$ such that, for all $q \in Q$, we have $\mathsf{s_A}(q) := \mathsf{s}'_\mathsf{A}(q)$ ensures that, for all $q \in Q$, we have $\chi_{\mathcal{G}}[\mathsf{A}](q) \leq \mathsf{val}[\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle](\mathsf{s_A}(q))$. Hence, in any case, there is a Player-A positional strategy that is locally optimal, and therefore also subgame optimal. □

Note that, as soon as we drop the assumption that all local interactions are maximizable w.r.t. Player A, then the above corollary fails. Indeed, there is an MDP where Player A plays alone and where playing ε-optimally requires infinite choice. An example is provided in Figure 4.2, formally defined in Definition 4.3 below, and argued in Proposition 4.4. Note that a similar example is given in [48, Prop. 28].

Definition 4.3. *The game of Figure 4.2 is in fact an MDP* Γ *where Player* A *plays alone with two states:* $Q := \{q_0, \bot\}$. *The state* $\bot$ *is a self-looping sink and, at state* q_0, *Player* A *may play any integer* $n \in \mathbb{N}$ *which leads to a distribution* $d_n := \{q_0 \mapsto 1 - \frac{1}{2^n}; \bot \mapsto \frac{1}{2^n}\} \in \mathcal{D}(Q)$. *Player* A *has a safety objective* Safe *with* $\mathsf{K} = \{0,1\}$ *and* $\mathsf{col}(q_0) := 0$ *and* $\mathsf{col}(\bot) := 1$, *i.e. Player* A *wants to avoid the state* $\bot$.

Proposition 4.4. *In the safety game* $\mathcal{G}$ *of Definition 4.2, the state* q_0 *has value 1 but Player* A *has no optimal strategy from* q_0 *and any finite-choice strategy has value 0 from* q_0.

Proof. First, consider any positional Player-A strategy $\mathsf{s_A}$. Consider some $n \in \mathbb{N}$ such that $\mathsf{s_A}(q_0)(n) > 0$. Then, at each step, there is probability at least $\frac{\mathsf{s_A}(q_0)(n)}{2^n}$ to reach the target $\bot$, otherwise the game loops back on q_0. Hence, almost-surely, the state $\bot$ is reached. In fact, all Player-A positional strategies

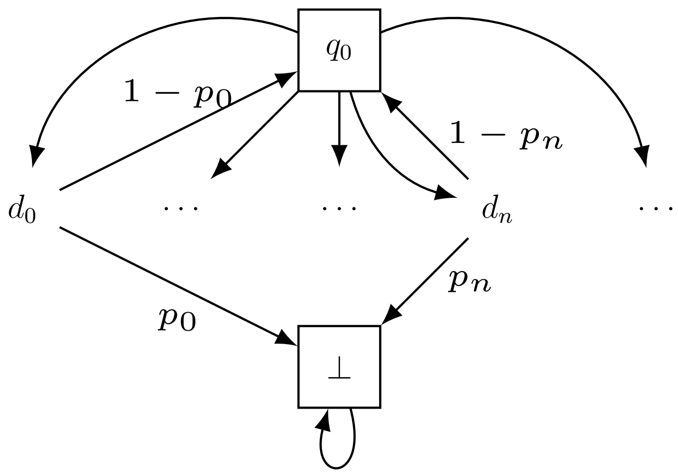

Figure 4.2: An MDP where Player A plays alone and wants to avoid the state $\bot$, with $p_i := \frac{1}{2^i}$: she does not have an optimal stategy, and playing almost-optimally requires infinite choice.

have value 0. Therefore, since this MDP is B-finite (recall, finitely many states, and Player B has finitely many actions), by Corollary 3.38, all finite-choice strategies have value 0 from q_0.

Consider some positive $\varepsilon > 0$. Let us build a Player-A strategy $\mathsf{s}_\mathsf{A}^\varepsilon$ of value at least $1 - \varepsilon$. Let $N \in \mathbb{N}$ be such that $\frac{1}{2^N} \leq \varepsilon$. For all $n \in \mathbb{N}$, we have $\mathsf{s}_\mathsf{A}(q_0^n) := \{N + n \mapsto 1\}$. That way, denoting s_B the only Player-B strategy in Γ, we have:

$$\mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\Gamma,q_0}(q_0^* \cdot \bot) = \sum_{n \in \mathbb{N}} \mathbb{P}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\Gamma,q_0}(q_0^n \cdot \bot) \leq \sum_{n \in \mathbb{N}} \frac{1}{2^{N+n+1}} = \frac{1}{2^N} \leq \varepsilon$$

Therefore, the value of this Player-A strategy s_A is at least $1 - \varepsilon$. □

Theorem 4.5. *In arbitrary finite-state concurrent safety games without stopping states:*

- *if the game is maximizable w.r.t. Player* A*, there is always a subgame optimal strategy that can be found among positional strategies;*
- *if not, there may not be optimal strategies and playing almost-optimally may require infinite choice;*
- *in any case if there is an optimal strategy, there is a subgame optimal positional one.*

These results are summarized in Table 4.1.

	GF	$\exists$ Opt. ?	ε-Opt.	Optimal	SubG. Opt.
Safety	Max. Arb.	Yes No	Positional ∞-choice	Positional	Positional

Table 4.1: Summary of how concurrent safety games behave.

Proof.

- This result is already known in the context of standard finite game local interactions, see [32, Theorem 1]. Corollary 4.3 above generalizes this result to arbitrary local interactions that are still maximizable w.r.t. Player A.

- It is already known that, in this context, infinite memory may be required for Player A, see [67, Theorem 3]. Note that, by using Corollary 3.38 on the examples provided to prove [67, Theorem 3], we would obtain that infinite-choice strategies is required to be almost-optimal. We also provide an example in Definition 4.3, argued in Proposition 4.4, where infinite-choice is required.

- This is given by Corollary 4.3.

□

4.2 Reachability games

In this section, we focus on reachability games — without stopping states, as for safety games — where the local interactions considered are arbitrary. Recall that, in such games, the goal of Player A is to reach a target, whereas Player B wants to avoid it. In all this section, given a reachability game $\mathcal{G}$, we will denote by $T := \mathsf{col}^{-1}[\{1\}] \subseteq Q$ the set of states that Player A wants to reach and we let $W_T := (\mathsf{col}^\omega)^{-1}(\mathsf{Reach}) \subseteq Q^\omega$ be the corresponding winning set for Player A. Without loss of generality, we assume that all states in the target T are self-looping sinks. It does not change the game since, once the target is reached, Player A has won regardless of what happens afterwards — since there are no stopping states. Therefore, the game can be seen as PI since the reachability objective W_T can be seen as a Büchi objective (where Player A wants to see infinitely often the target T) without changing the game since reaching once the target is equivalent to reaching it infinitely often. Hence, we may use Corollary 3.14 and Corollary 3.16 from the previous chapter, which only apply to PI games.

This section is an adaptation of the first part of [39] where, instead of considering only standard game forms with finitely many actions, we consider games with arbitrary interactions.

4.2.1 . Computing the Player-A value of reachability games

It is known for a long time that the values in reachability games can be computed with a least fixed point operator [10, 68], including with non-standard game forms in [10]. However, even in [10], the game forms considered were assumed valuable. Therefore, the reachability games considered had a value. We do not make such an assumption on local interactions in this subsection, and we show that it still holds that the Player-A value of the game can be computed with a least fixed point regardless of the local interactions involved.

First, we define the operator on which we will consider a least fixed point.

Definition 4.4. *Consider an arbitrary finite-state concurrent reachability game $\mathcal{G}$ without stopping states. We let $\mathsf{Val}_T := \{v : Q \to [0,1] \mid v[T] = \{1\}\}$ be the set of valuations mapping each state in the target to 1. We let $\Delta_\mathcal{G} : \mathsf{Val}_T \to \mathsf{Val}_T$ be such that, for all $v \in \mathsf{Val}_T$ and $q \in Q$, we have:*

$$\Delta_\mathcal{G}(v)(q) := \begin{cases} 1 & \text{if } q \in T \\ \mathsf{val}[\langle \mathsf{F}(q), v\rangle][\mathsf{A}] & \text{otherwise} \end{cases}$$

Hence, we do have $\Delta_\mathcal{G}(v) \in \mathsf{Val}_T$.

This operator ensures several useful properties that we describe below.

Lemma 4.6. *For all arbitrary finite-state concurrent reachability games $\mathcal{G}$ without stopping states, the operator $\Delta_\mathcal{G}$ ensures the following:*

- *it is non-decreasing, i.e. for all $v, v' \in \mathsf{Val}_T$ such that $v \leq v'$, we have $\Delta_\mathcal{G}(v) \leq \Delta_\mathcal{G}(v')$;*
- *it is 1-Lipschitz, i.e. for all $v, v' \in \mathsf{Val}_T$, we have $\|\Delta_\mathcal{G}(v) - \Delta_\mathcal{G}(v')\|_\infty \leq \|v - v'\|_\infty$;*
- *for all $n \in \mathbb{N}$ and $v \in \mathsf{Val}_T$, we denote by $\Delta_\mathcal{G}^{(n)}(v) \in \mathsf{Val}_T$ the vector obtained from v after n applications of the operator $\Delta_\mathcal{G}$. Denoting by $v_0 \in \mathsf{Val}_T$ the valuation such that $v_0[Q \setminus T] := 0$, we have that the sequence $(\Delta_\mathcal{G}^{(n)}(v_0))_{n\in\mathbb{N}}$ has a limit in Val_T that is equal to the least fixed point of the operator $\Delta_\mathcal{G}$.*

Proof. The two first properties come from Lemma 1.19. The third point comes from Kleene least fixed point theorem. □

Definition 4.5 (Notation least fixed point $\Delta_\mathcal{G}$). *For all arbitrary finite-state reachability games $\mathcal{G}$ without stopping states, we denote by $\mathsf{m}_\mathcal{G} : Q \to [0,1]$ the least fixed point of the operator $\Delta_\mathcal{G}$ (whose existence is ensured by Lemma 4.6).*

In fact, in all reachability games, this least fixed point is equal to the Player-A value of the game.

Proposition 4.7. *For all arbitrary finite-state reachability games $\mathcal{G}$ without stopping states, we have $\mathsf{m}_{\mathcal{G}} = \chi_{\mathcal{G}}[\mathsf{A}] : Q \to [0,1]$.*

Proof. Let $n \in \mathbb{N}$. Taking the notation of Lemma 4.6, let us show that $\chi_{\mathcal{G}}[\mathsf{A}] \geq \Delta_{\mathcal{G}}^{(n)}(v_0)$. Let $\varepsilon > 0$. We let $v : Q^+ \to [0,1]$ be such that, for all $\rho \in Q^+$, letting $k := |\rho| - 1$:

$$v(\rho) := \begin{cases} 1 & \text{if } \exists i \leq k,\ \rho_i \in T \\ \max(\Delta_{\mathcal{G}}^{(n-k)}(v_0)(\rho) - \frac{\varepsilon}{2^k}, 0) & \text{otherwise, if } k < n \\ 0 = v_0(\rho) & \text{otherwise} \end{cases}$$

Let us define a Player-A strategy $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}}$ dominating this valuation. Let $\rho \in Q^+$ and $k := |\rho|-1$. If $v(\rho) = 0$, then $\mathsf{s}_\mathsf{A}(\rho) \in \Sigma_\mathsf{A}^{\rho_{\mathsf{lt}}}$ is defined arbitrarily. Assume now that $v(\rho) > 0$. If $k \geq n$, $\mathsf{s}_\mathsf{A}(\rho)$ is also defined arbitrarily. Indeed, since $v(\rho) > 0$, this implies $v(\rho) = 1$ and therefore $v(\rho \cdot q) = 1$ for all $q \in Q$. Hence, $1 = \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle][\mathsf{s}_\mathsf{A}(\rho)] \geq v(\rho) = 1$. Assume now that it is not the case, that is $k < n$ and $v(\rho) > 0$. This means that $v(\rho) = \Delta_{\mathcal{G}}^{(n-k)}(v_0)(\rho) - \frac{\varepsilon}{2^k}$. Furthermore, $\Delta_{\mathcal{G}}^{(n-k-1)}(v_0) - \frac{\varepsilon}{2^{k+1}} \leq v^\rho$ (with $\Delta_{\mathcal{G}}^{(0)}(v_0) = v_0$). Hence, by Lemma 1.19, we have $\Delta_{\mathcal{G}}^{(n-k)}(v_0)(\rho) = \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \Delta_{\mathcal{G}}^{(n-k+1)}(v_0) \rangle][\mathsf{A}] \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle][\mathsf{A}] + \frac{\varepsilon}{2^{k+1}}$. That is, $v(\rho) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle][\mathsf{A}] - \frac{\varepsilon}{2^{k+1}}$. We let $\mathsf{s}_\mathsf{A}(\rho) \in \Sigma_\mathsf{A}^{\rho_{\mathsf{lt}}}$ be such that $\mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle][\mathsf{s}_\mathsf{A}(\rho)] \geq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle][\mathsf{A}] - \frac{\varepsilon}{2^{k+1}}$, which therefore ensures that $v(\rho) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle][\mathsf{s}_\mathsf{A}(\rho)]$. This concludes the definition of s_A which indeed dominates the valuation v.

Let us show that this strategy guarantees the valuation v by applying Theorem 3.12. The first condition of this theorem is satisfied. Furthermore, for all $\rho \in Q^\omega$, we have $\mathsf{limsup}_v(\rho) \in \{0,1\}$ with $\mathsf{limsup}_v(\rho) = 1$ if and only if $\rho \in W_T$. Therefore, the second condition of this theorem is also satisfied. Hence, the strategy s_A guarantees the valuation v with, for all $q \in Q$, $v(q) \geq \Delta_{\mathcal{G}}^{(n)}(v_0)(q) - \varepsilon$. As this holds for all $\varepsilon > 0$, it follows that $\chi_{\mathcal{G}}[\mathsf{A}] \geq \Delta_{\mathcal{G}}^{(n)}(v_0)$. As this holds for all $n \in \mathbb{N}$ and $\mathsf{m}_{\mathcal{G}} = \lim_{n \to \infty} \Delta_{\mathcal{G}}^{(n)}(v_0)$, it follows that $\chi_{\mathcal{G}}[\mathsf{A}] \geq \mathsf{m}_{\mathcal{G}}$.

Let us now show that $\chi_{\mathcal{G}}[\mathsf{A}] \leq \mathsf{m}_{\mathcal{G}}$. Fix a Player-$\mathsf{A}$ strategy $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}}$. Consider some $\varepsilon > 0$. For all $i \in \mathbb{N}$, we let $w_i : Q \to [0,1]$ be such that, we have: $w_i := \min(\mathsf{m}_{\mathcal{G}} + \frac{\varepsilon}{2^i}, 1)$ and we let $v : Q^+ \to [0,1]$ be such that, for all $\rho \in Q^+$, we have $v(\rho) := 1$ is ρ has visited T and $v(\rho) := w_{|\rho|-1}(\rho_{\mathsf{lt}}) \in [0,1]$ otherwise. Let us define a Player-B strategy $\mathsf{s}_\mathsf{B} \in \mathsf{S}_\mathsf{B}^{\mathcal{C}}$. For all $\rho \in Q^+$: if $v(\rho) = 1$, then $\mathsf{s}_\mathsf{B}(\rho)$ is defined arbitrarily and therefore $\mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v \rangle](\mathsf{s}_\mathsf{A}(\rho), \mathsf{s}_\mathsf{B}(\rho)) \leq v(\rho)$. Otherwise, we have $v(\rho) = \mathsf{m}_{\mathcal{G}}(\rho_{\mathsf{lt}}) + \frac{\varepsilon}{2^{|\rho|-1}}$. In addition, we have $v^\rho \leq \mathsf{m}_{\mathcal{G}} + \frac{\varepsilon}{2^{|\rho|}}$. Furthermore, we have $\mathsf{m}_{\mathcal{G}}(\rho_{\mathsf{lt}}) = \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \mathsf{m}_{\mathcal{G}} \rangle][\mathsf{A}]$. Hence, $\mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{s}_\mathsf{A}(\rho)) \leq \mathsf{m}_{\mathcal{G}}(\rho_{\mathsf{lt}})$. We deduce that $\mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle](\mathsf{s}_\mathsf{A}(\rho)) \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{s}_\mathsf{A}(\rho)) + \frac{\varepsilon}{2^{|\rho|}} \leq \mathsf{m}_{\mathcal{G}}(\rho_{\mathsf{lt}}) + \frac{\varepsilon}{2^{|\rho|}} < v(\rho)$. We set $\mathsf{s}_\mathsf{B}(\rho) \in \Sigma_\mathsf{B}^{\rho_{\mathsf{lt}}}$ such that $\mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), v^\rho \rangle](\mathsf{s}_\mathsf{A}(\rho), \mathsf{s}_\mathsf{B}(\rho)) \leq v(\rho)$. This concludes the definition of the strategy s_B. From its definition, we can deduce that in the stochastic tree

$\mathcal{T}_{\mathcal{C}}^{\mathsf{s_A},\mathsf{s_B}}$ induced by $\mathsf{s_A}$ and $\mathsf{s_B}$, the valuation v is non-increasing (recall Definition 2.3). Therefore, by Proposition 2.9 with a non-increasing valuation, for all $q \in Q$, we have $\mathsf{m}_{\mathcal{G}}(q) + \varepsilon \geq v(q) \geq \mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[\mathsf{limsup}_v]$. Furthermore, for all infinite paths $\rho \in Q^\omega$, we have if $\rho \in W_T$, then $\mathsf{limsup}_v(\rho) = 1$. Hence, $\mathsf{limsup}_v \geq \mathbb{1}_{W_T}$. Therefore, $\mathsf{m}_{\mathcal{G}}(q) + \varepsilon \geq \mathbb{P}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[W_T] \geq \chi_{\mathcal{G}}(\mathsf{s_A})[q] \geq \chi_{\mathcal{G}}(\mathsf{A})[q]$. Since this holds for all positive $\varepsilon > 0$, it follows that $\mathsf{m}_{\mathcal{G}} \geq \chi_{\mathcal{G}}(\mathsf{A})$. Hence, $\mathsf{m}_{\mathcal{G}} = \chi_{\mathcal{G}}(\mathsf{A})$. □

We conclude this subsection by stating a very useful proposition that we will apply in the next subsection to the valuation $\mathsf{m}_{\mathcal{G}} = \chi_{\mathcal{G}}[\mathsf{A}]$. The proof of this proposition is not long, but is very technical, and hence is postponed to the appendix.

Proposition 4.8 (Proof 4.4.1)**.** *Let $n \geq 1$. Consider a function $g : [0,1]^n \to [0,1]^n$ that is non-decreasing and 1-Lipschitz. Assume that its least fixed point $\mathsf{m} \in [0,1]$ is such that, for all $i \in [\![1,n]\!]$, we have $\mathsf{m}(i) > 0$. Then, for all $\varepsilon > 0$, there exists a valuation $v \in [0,1]^n$ such that $v \leq \mathsf{m}$, $\|\mathsf{m} - v\|_\infty \leq \varepsilon$ and for all $i \in [\![1,n]\!]$: $g(v)(i) > v(i)$.*

4.2.2 . Computing the set of maximizable states

Recall that in the snow-ball reachability (standard) game of Definition 3.6, Player A does not have an optimal strategy, even if the game has finitely many states. The aim of this subsection and the next is, given a finite-state concurrent reachability game to determine exactly from which states Player A has an optimal strategy. This, in turn, will give that whenever she has an optimal strategy, she has one that is positional. This extends Everett [10] (the existence of positional ε-optimal strategies). Note that in [10], arbitrary game forms are considered (not only standard ones), though they are assumed valuable.

In this subsection, we present the definitions and arguments in standard finite concurrent games — in particular, all local interactions are standard and finite. We also illustrate these definitions on examples. This subsection directly comes from [39]. In the Appendix 4.4.2, we give the formal definitions well suited for arbitrary game forms along with the formal proofs of correctness. Note that, although the underlying ideas are not too complicated, the formal proofs are quite technical. This is mainly due to the fact that we need to deal with infinite-memory strategies.

For the remainder of this subsection and the next, we consider an arbitrary finite-state concurrent reachability game $\mathcal{G}$, without stopping states. We still denote by $T := \mathsf{col}^{-1}[\{1\}] \subseteq Q$ the set of states that Player A wants to reach, and by $W_T := (\mathsf{col}^\omega)^{-1}(\mathsf{Reach}) \subseteq Q^\omega$ the set of infinite sequences of states reaching the set T. Let us first introduce some terminology that is relevant regardless of the game forms considered.

Definition 4.6 (Maximizable and sub-maximizable states)**.** *A state $q \in Q$*

from which Player A *has (resp. does not have) an optimal strategy is called* maximizable *(resp.* sub-maximizable*). The set of such states is denoted* $\mathsf{OptQ_A}$ *(resp.* $\mathsf{SubOptQ_A}$*).*

Remark 4.2. *One has to careful here: we have already used the "maximizable" terminology in this dissertation. Recall, this refers to the game forms where a player has optimal* GF*-strategies in all the games in normal form induced from that game form. The terminology we have defined above refers to states in a reachability game. In particular, the local interactions of maximizable states may not be maximizable for any player. The two notions are completely unrelated.*

For the remainder of this subsection, we assume that the game $\mathcal{G}$ is standard (and that all standard local interactions in $\mathcal{G}$ are finite). We want to build an optimal (and positional) strategy for Player A when possible. Recall Corollary 3.16: to be optimal, a Player-A positional strategy $\mathsf{s_A}$ has to play optimally at each local interaction $\mathsf{F}(q)$ (for $q \in Q$) with respect to the valuation $\chi_{\mathcal{G}}[\mathsf{A}] : Q \to [0,1]$. However, it is not sufficient: in the snow-ball game of Figure 3.1, when Player A plays optimally in $\mathsf{F}(q_0)$ w.r.t. the valuation $\chi_{\mathcal{G}}[\mathsf{A}]$ (that is, plays the top row with probability 1), Player B can enforce the game never to leave the state $q_0 \notin T$. Hence, locally, we want to have strategies that not only play locally optimally but also, regardless of the actions of Player B, have a non-zero probability to get closer to the target T. Such strategies will be called progressive strategies. To properly define this notion on standard game forms, we first introduce the notion of optimal Player-B actions.

Definition 4.7 (Optimal Player-B actions)**.** *Let $q \in Q$ be a state of the game. Consider the game in normal form $\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle$. For all* GF*-strategies $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathsf{F}(q))$, we define the set $\mathsf{Resp}^{\mathsf{B}}_{\sigma_\mathsf{A}}(q) \subseteq \mathsf{Act}^q_\mathsf{B}$ of* optimal actions *of Player* B *w.r.t. the* GF*-strategy* σ_A *by*

$$\mathsf{Resp}^{\mathsf{B}}_{\sigma_\mathsf{A}}(q) := \{ b \in \mathsf{Act}^q_\mathsf{B} \mid \mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_\mathsf{A}, b) = \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_\mathsf{A}) \}$$

In Figure 4.3, the set $\mathsf{Resp}^{\mathsf{B}}_{\sigma_\mathsf{A}}(q)$ of optimal Player-B actions w.r.t. the strategy σ_A are represented in bold purple: the expected values of these actions is the value of the GF-strategy: 1/2.

We can now define the set of progressive strategies on standard finite game forms, see Page 207 for a definition on arbitrary game forms.

Definition 4.8 (Progressive strategies in standard finite game forms)**.** *Consider a state $q \in Q$ and a set of good states $\mathsf{Gd} \subseteq Q$ that Player* A *wants to reach. We let $\mathsf{Gd}_{\mathcal{D}} \subseteq \mathcal{D}(Q)$ be the set of distributions over states with a non-zero probability to reach the set* Gd*: $\mathsf{Gd}_{\mathcal{D}} := \{ d \in \mathcal{D}(Q) \mid \mathsf{Sp}(d) \cap \mathsf{Gd} \neq \emptyset \}$. The set of* progressive strategies $\mathsf{Prog}_q(\mathsf{Gd})$ *at state q w.r.t.* Gd *is defined by*

$$\mathsf{Prog}_q(\mathsf{Gd}) := \{ \sigma_\mathsf{A} \in \mathsf{Opt}_\mathsf{A}(\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle) \mid \forall b \in \mathsf{Resp}^{\mathsf{B}}_{\sigma_\mathsf{A}}(q), \exists a \in \mathsf{Sp}(\sigma_\mathsf{A}),\ \delta_q(a,b) \in \mathsf{Gd}_{\mathcal{D}} \}$$

In Figure 4.3, the distributions over states in $\mathsf{Gd}_{\mathcal{D}}$ are arbitrarily chosen for the example and circled in green. The depicted Player-A GF-strategy is progressive as, for all bold purple actions, there is a green-circled outcome in the support of the strategy (the circled 3/4).

Progressive strategies are not enough. In reachability games in general, some states may be sub-maximizable. In that case, playing optimally implies avoiding these states. Given a set $\mathsf{Bd} \subseteq Q$ of states to avoid, an optimal GF-strategy that has a non-zero probability to reach that set of states Bd with an optimal Player-B action is called risky. We give and illustrate below the definition of risky strategies in standard finite game forms, see Page 208 for a definition on arbitrary game forms.

Definition 4.9 (Risky strategy in standard finite game forms). *Let $q \in Q$ be a state of the game and $\mathsf{Bd} \subseteq Q$ be a set of states that Player A wants to avoid. The set of distributions over states $\mathsf{Bd}_{\mathcal{D}} \subseteq \mathcal{D}(Q)$ is defined similarly to Gd_{D} in Definition 4.8: $\mathsf{Bd}_{\mathcal{D}} := \{d \in \mathcal{D}(Q) \mid \mathsf{Sp}(d) \cap \mathsf{Bd} \neq \emptyset\}$. Then, the set of* risky strategies $\mathsf{Risk}_q(\mathsf{Bd})$ *at state q w.r.t.* Bd *is defined by*

$$\mathsf{Risk}_q(\mathsf{Bd}) := \{\sigma_{\mathsf{A}} \in \mathsf{Opt}_{\mathsf{A}}(\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle) \mid \exists b \in \mathsf{Act}^q_{\mathsf{B}},\ \exists a \in \mathsf{Sp}(\sigma_{\mathsf{A}}),\ \delta_q(a,b) \in \mathsf{Bd}_{\mathsf{D}}\}$$

In Figure 4.3, the set of distributions over states $\mathsf{Bd}_{\mathcal{D}}$ are also arbitrarily chosen for the example and circled in red. The GF-strategy σ_{A} is not risky since no red-squared outcome appears in the intersection of the support of σ_{A} and the purple actions in $\mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)$.

Overall, we want for local strategies to be efficient, that is both progressive and not risky.

Definition 4.10 (Efficient strategies in arbitrary game forms). *Let $q \in Q$ be a state of the game and $\mathsf{Gd}, \mathsf{Bd} \subseteq Q$ be sets of states. The set of* efficient strategies $\mathsf{Eff}_q(\mathsf{Gd}, \mathsf{Bd})$ *at state q w.r.t.* Gd *and* Bd *is defined by* $\mathsf{Eff}_q(\mathsf{Gd}, \mathsf{Bd}) := \mathsf{Prog}_q(\mathsf{Gd}) \setminus \mathsf{Risk}_q(\mathsf{Bd})$.

In Figure 4.3, the GF-strategy σ_{A} is efficient as it is both progressive and not risky.

We can now compute inductively the set of maximizable and sub-maximizable states. First, given a set of sub-maximizable states Bd, we define iteratively below a set of secure states w.r.t. Bd, they are the states with a non-zero probability to get closer to the target $\top$ while avoiding the set Bd. The construction is illustrated in Figure 4.4.

Definition 4.11 (Secure states). *Consider a set of states $\mathsf{Bd} \subseteq Q$. We set $\mathsf{Sec}_0(\mathsf{Bd}) := T$ and, for all $i \geq 0$, $\mathsf{Sec}_{i+1}(\mathsf{Bd}) := \mathsf{Sec}_i(\mathsf{Bd}) \cup \{q \in Q \setminus \mathsf{Bd} \mid \mathsf{Eff}_q(\mathsf{Sec}_i(\mathsf{Bd}), \mathsf{Bd}) \neq \emptyset\}$. The set $\mathsf{Sec}(\mathsf{Bd})$ of states* secure *w.r.t.* Bd *is:* $\mathsf{Sec}(\mathsf{Bd}) := \cup_{n \in \mathbb{N}} \mathsf{Sec}_n(\mathsf{Bd}) \cup (\mathsf{m}_{\mathcal{G}})^{-1}[0]$.

Note that, as there are finitely many states, this procedure terminates in at most $n = |Q|$ steps. Furthermore, the states of value 0 are added since

$$\sigma_{\mathsf{A}} : \begin{matrix} 0 \\ 0.5 \\ 0.5 \\ 0 \end{matrix} \begin{pmatrix} 1/2 & \mathbf{1/2} & \boxed{\mathbf{0}} & 0 \\ 1 & \mathbf{3/4} & \mathbf{1/4} & 1 \\ \boxed{3/4} & \mathbf{1/4} & \mathbf{3/4} & 1/2 \\ 0 & \mathbf{0} & \boxed{\mathbf{1/2}} & 1 \end{pmatrix}$$

Figure 4.3: A game in normal form with an optimal GF-strategy depicted in brown on the left. Its value is $1/2 = 1/2 \cdot 3/4 + 1/2 \cdot 1/4$.

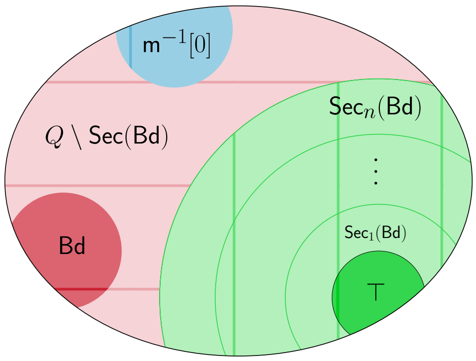

Figure 4.4: The construction of Definition 4.11 of the set of states $\mathsf{Sec}(\mathsf{Bd})$: it is the reunion of the blue and green vertical stripe areas.

any state of value 0 is maximizable. The benefit of this construction lies in the lemma below: if all states in Bd are sub-maximizable, then all states in $Q \setminus \mathsf{Sec}(\mathsf{Bd})$ also are.

Lemma 4.9 (Proof Page 209). *Assume that a set of states* $\mathsf{Bd} \subseteq Q$ *is such that* $\mathsf{Bd} \subseteq \mathsf{SubOptQ_A}$. *Then, the set of states* $Q \setminus \mathsf{Sec}(\mathsf{Bd})$ *is such that* $Q \setminus \mathsf{Sec}(\mathsf{Bd}) \subseteq \mathsf{SubOptQ_A}$ *(these correspond to the red horizontal stripe areas in Figure 4.4).*

Proof sketch. For an arbitrary Player A strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ to be optimal, it roughly needs, on all relevant paths, to be optimal. More precisely, on any finite path $\pi \in Q^+$ with a non-zero probability to occur if Player B plays optimal actions (recall Definition 4.7) against the strategy $\mathsf{s_A}$ — the path $\pi \in Q^+$ is called a relevant path — the strategy $\mathsf{s_A}$ needs to play an optimal GF-strategy in the local interaction $\mathsf{F}(\pi_{\mathsf{lt}})$ and the residual strategy $\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\pi)}$ has to be optimal from π_{lt} in the reachability game $\mathcal{G}$. Therefore, on all relevant paths, the strategy $\mathsf{s_A}$ has to play optimal GF-strategies that are not risky. However, in any local interaction of a state $q \in Q \setminus \mathsf{Sec}(\mathsf{Bd})$, there is no efficient strategies available to Player A. Therefore, if the game starts from a state $q \in Q \setminus \mathsf{Sec}(\mathsf{Bd})$ an optimal strategy $\mathsf{s_A}$ for Player A (which therefore is locally optimal but not progressive) would allow Player B to ensure staying in the set $Q \setminus \mathsf{Sec}(\mathsf{Bd})$ while playing optimal actions. In that case, the game never leaves the set $Q \setminus \mathsf{Sec}(\mathsf{Bd})$, which induces a value of 0, whereas $\chi_\mathcal{G}[\mathsf{A}](q) > 0$ since $q \notin \mathsf{Sec}(\mathsf{Bd})$. Thus, there is no optimal strategy for Player A from a state in $Q \setminus \mathsf{Sec}(\mathsf{Bd})$. □

We can now define inductively the set of bad states (which, in turn, will correspond to the set of sub-maximizable states).

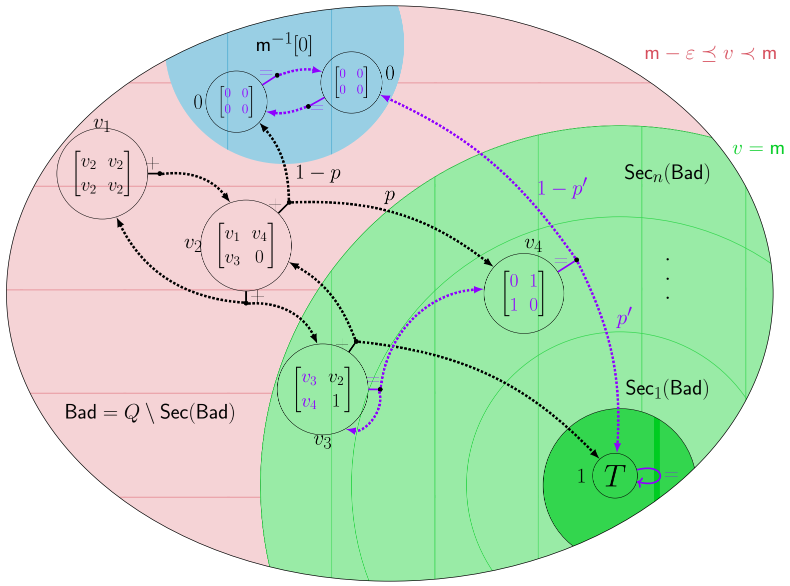

Figure 4.5: An illustration of the proof of Lemma 4.10 on the MDP induced by the strategy s_A. Labels $v_1, \ldots, v_4$ is the value of the corresponding states given by the valuation v.

Definition 4.12 (Set of sub-maximizable states). *Let* $\mathsf{Bad}_0 := \emptyset$ *and, for all* $i \geq 0$, $\mathsf{Bad}_{i+1} := Q \setminus \mathsf{Sec}(\mathsf{Bad}_i)$. *Then, the set* Bad *of bad states is equal to* $\mathsf{Bad} := \cup_{n \in \mathbb{N}} \mathsf{Bad}_n$ *for* $n = |Q|$.

Note that, as in the case of the set of secure states, since the game $\mathcal{G}$ is finite, this procedure ends in at most $n = |Q|$ steps.

Lemma 4.9 above ensures that the set of states Bad is included in $\mathsf{SubOptQ_A}$. In addition, we have that there exists a Player A positional strategy optimal from all states q in its complement $\mathsf{Sec}(\mathsf{Bad}) = Q \setminus \mathsf{Bad}$, as stated in the lemma below.

Lemma 4.10 (Proof Page 215). *For all* $\varepsilon > 0$, *there is a positional strategy* $s_A \in S_A^{\mathcal{C}}$ *such that:*

- *for all* $q \in \mathsf{Sec}(\mathsf{Bad})$, *we have* $\chi_{\mathcal{G}}[s_A](q) = \chi_{\mathcal{G}}[A](q)$;
- *for all* $q \in \mathsf{Bad}$, *we have* $\chi_{\mathcal{G}}[s_A](q) \geq \chi_{\mathcal{G}}[A](q) - \varepsilon$.

In particular, it follows that $\mathsf{Sec}(\mathsf{Bad}) \subseteq \mathsf{OptQ_A}$.

Proof sketch. To prove this lemma, we define a Player-A positional strategy $s_A \in S_A^{\mathcal{C}}$, a valuation $v \in [0,1]^Q$ of the states, we prove that the strategy s_A dominates that valuation and we prove that the only ECs compatible with s_A that are not the target have value 0. (Recall that all states in the target T are assumed self-looping sinks.) This will show that the strategy $\mathsf{s_A}$ guarantees the valuation v by applying Corollary 3.16. Recall that $\chi_{\mathcal{G}}[\mathsf{A}] = \mathsf{m}_{\mathcal{G}}$ by Proposition 4.7. As we want the strategy $\mathsf{s_A}$ to be optimal from all secure states, we consider a partial valuation v such that $v|_{\mathsf{Sec(Bad)}} := \mathsf{m}|_{\mathsf{Sec(Bad)}}$ (we will define it later on Bad). Then, on all secure states $q \in \mathsf{Sec}_i(\mathsf{Bad}) \setminus \mathsf{Sec}_{i-1}(\mathsf{Bad})$, we set $\mathsf{s_A}(q)$ to be an efficient strategy w.r.t. $\mathsf{Sec}_{i-1}(\mathsf{Bad})$ and Bad, i.e. $\mathsf{s_A}(q) \in \mathsf{Eff}_q(\mathsf{Sec}_{i-1}(\mathsf{Bad}), \mathsf{Bad})$. In particular, the GF-strategy $\mathsf{s_A}(q)$ is optimal in the game form $\mathsf{F}(q)$ w.r.t. the valuation $\mathsf{m}_{\mathcal{G}}$. However, we know that no strategy can be optimal from states in Bad. Hence, we consider a valuation v that is ε-close to the valuation $\mathsf{m}_{\mathcal{G}}$ on states in Bad for a well-chosen $\varepsilon > 0$. This ε is chosen such that, for all $q \in \mathsf{Sec(Bad)}$, the value of the GF-strategy $\mathsf{s_A}(q) \in \Sigma_{\mathsf{A}}^q$ in the game in normal form $\langle \mathsf{F}(q), v \rangle$ is at least $v(q)$[3]. We can now define the valuation v and the strategy $\mathsf{s_A}$ on Bad such that $\|v - \mathsf{m}_{\mathcal{G}}\|_\infty \leq \varepsilon$ and, for all $q \in \mathsf{Bad}$, the value of $\mathsf{s_A}(q)$ in $\mathsf{F}(q)$ w.r.t. v is greater than $v(q)$: $\mathsf{val}[\langle \mathsf{F}(q), v \rangle](\mathsf{s_A}(q)) > v(q)$. Note that this is where Proposition 4.8 comes into play. The valuation v and the strategy $\mathsf{s_A}$ are now completely defined on Q. By definition, the strategy $\mathsf{s_A}$ dominates the valuation v.

The MDP induced by the strategy $\mathsf{s_A}$ is schematically depicted in Figure 4.5. The different split arrows appearing in the figure correspond to the actions (or columns in the local interactions) available to Player B. Black +-labeled-split arrows correspond to the actions of Player B that increase the value of v, i.e. in a state q, such that the expected value w.r.t. to the probabilities chosen by the strategy $\mathsf{s_A}$ – of the values of the successor states of q given by v is greater than $v(q)$. For instance, we have $v_2 < p \cdot v_4 + (1-p) \cdot 0$, where the probability $p \in [0,1]$ is set by the strategy $\mathsf{s_A}$. On the other hand, purple =-labeled-split arrows correspond to the actions whose values are equal to the value of the state. For instance $v_4 = (1-p') \cdot 0 + p' \cdot 1$. We can see that the only split arrows exiting states in Bad (the red horizontal stripe area) are black (since $\mathsf{val}[\langle \mathsf{F}(q), v \rangle](\mathsf{s_A}(q)) > v(q)$ for all $q \in \mathsf{Bad}$). However, from a secure state $q \in \mathsf{Sec(Bad)}$ (the green and blue vertical stripe areas) there are also purple split arrows. Note that, in these secure states $q \in \mathsf{Sec(Bad)}$, purple split arrows correspond to the optimal actions $\mathsf{Resp}^{\mathsf{B}}_{\mathsf{s_A}(q)}(q)$ at the local interaction $\mathsf{F}(q)$. Furthermore, these split arrows cannot exit the set of secure states $\mathsf{Sec(Bad)}$ since the local strategy $\mathsf{s_A}(q)$ is not risky.

We can then prove that the strategy $\mathsf{s_A}$ guarantees the valuation v by applying Corollary 3.16: since $\mathsf{s_A}$ locally dominates the valuation v, it remains to show that all the ECs different that are not in the target T have only states of value 0. In the figure, this corresponds to having ECs only in the blue upper circle and dark green bottom right inner circle areas. In fact, Corollary 3.15 gives that any state q in an EC ensures $\mathsf{val}[\langle \mathsf{F}(q), v \rangle](\mathsf{s_A}(q)) = v(q)$, which

[3]Specifically, ε has to be chosen smaller than the smallest difference between the values of optimal actions in $\mathsf{Resp}^{\mathsf{B}}_{\mathsf{s_A}(q)}(q)$ and non-optimal action.

	∃ Opt. ?	ε-Opt.	Optimal	SubG. Opt.
Reach	No	Positional	Positional	Positional

Table 4.2: The summary of the situation in finite-state arbitrary concurrent reachability games.

implies that no state in Bad can be in an EC. This can be seen in the figure between the states of value v_1 and v_2: because of the black arrow from v_1 to v_2, we necessarily have $v_1 < v_2$. Then, v_2 cannot loop (with probability one) to v_1 since this would imply $v_2 < v_1$. As all the split arrows are black for states in Bad, no EC can occur in this region. Furthermore, the optimal actions in the secure states always have a non-zero probability to get closer to the target T. In the figure, this corresponds to the fact that there is always one tip of a purple split arrow that goes down in the $(\mathsf{Sec}_i(\mathsf{Bad}))_{i \in \mathbb{N}}$ hierarchy (since the strategy $\mathsf{s}_\mathsf{A}(q)$ is progressive): in the example, from v_3 to v_4 and from v_4 to the target T. Therefore, the only loop (with probability one) that can occur in the set $(\mathsf{Sec}_i(\mathsf{Bad}))_{i \in \mathbb{N}}$ is at the target T (recall that all states in the target T are assumed self-looping). We conclude by applying Corollary 3.16. □

Overall, we obtain the theorem below summarizing the results proved in this section. Note that we state it with arbitrary local interactions since it is what will be proved in Appendix 4.4.2, however we only argued the standard case in this subsection.

Theorem 4.11. *In an arbitrary finite-state concurrent reachability game $\mathcal{G}$ without stopping states, we have $\mathsf{Bad} = \mathsf{SubOptQ_A}$ and $\mathsf{Sec}(\mathsf{Bad}) = \mathsf{OptQ_A}$. Furthermore, for all $\varepsilon > 0$, there is a Player-A positional strategy s_A optimal from all states in $\mathsf{OptQ_A}$ and ε-optimal from all states in $\mathsf{SubOptQ_A}$.*

Proof. Initially, $\mathsf{Bad}_0 = \emptyset \subseteq \mathsf{SubOptQ_A}$. Then, by Lemma 4.9, for all $i \geq 0$, we have $\mathsf{Bad}_{i+1} = Q \setminus \mathsf{Sec}(\mathsf{Bad}_i) \subseteq \mathsf{SubOptQ_A}$. In particular, $\mathsf{Bad} = \mathsf{Bad}_n \subseteq \mathsf{SubOptQ_A}$. Furthermore, by Lemma 4.10, there exists a Player-A optimal strategy from all states in $\mathsf{Sec}(\mathsf{Bad}) = Q \setminus \mathsf{Bad}$. Hence, $\mathsf{Sec}(\mathsf{Bad}) \subseteq \mathsf{OptQ_A}$. As we have $Q = \mathsf{Bad} \uplus \mathsf{Sec}(\mathsf{Bad}) = \mathsf{OptQ_A} \uplus \mathsf{SubOptQ_A}$, it follows that: $\mathsf{Bad} = \mathsf{SubOptQ_A}$ and $\mathsf{Sec}(\mathsf{Bad}) = \mathsf{OptQ_A}$. Then the result is straightforwardly deduced from Lemma 4.10. □

We summarize the results in reachability games in the theorem below.

Theorem 4.12. *In arbitrary finite-state concurrent reachability games without stopping states:*

- *there does not always exist optimal strategies, which can be witnessed by a standard finite game;*
- *for all positive $\varepsilon > 0$, there is a positional strategy that is ε-optimal;*

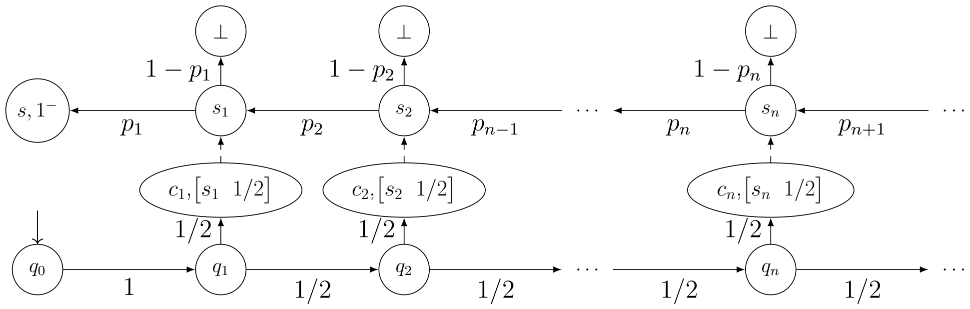

Figure 4.6: An infinite concurrent reachability game $\mathcal{C}$ (the Nature states are omitted). The probabilities p_k are such that, for all $i \geq 1$, the value of the state s_i is $\chi^{\mathcal{C}}(s_i) = \Pi_{k=1}^{i} p_k = (1/2 + 1/2^i)$.

- *whenever there exists an optimal strategy, there is one that is positional. This also holds for subgame optimal strategies.*

These results can be seen in Table 4.2.

Proof.

- This was originally shown in [10]. We provide an example in the snow-ball game of Definition 3.6, with its properties detailed in Proposition 3.8.
- This was first shown in [10] with valuable local interactions, non-necessarily standard ones. Theorem 4.11 above generalizes this results to arbitrary local interactions.
- The existence of a (subgame) optimal strategy implies that all states are maximizable. Hence, the same Theorem 4.11 gives that there exists a positional optimal strategy. Note that it is also subgame optimal since it is positional.

□

Infinite arenas. We conclude this section and chapter by showing that Theorem 4.11 fails in concurrent reachability with infinitely many states. This is already known, see for instance [48, Proposition 21]. Hence, we give here only informal explanations.

In Figure 4.6, we have depicted an infinite concurrent reachability game where the state q_0 is maximizable but, from q_0, Player A does not have any positional optimal strategy. Indeed, in state s is plugged the snow-Ball game of Definition 3.6 — the target is therefore denoted ⊤ — whose value is 1 but Player A does not have an optimal strategy. Then, for all $i \geq 0$, the probability to reach s from s_i is equal to $v_i = (1/2 + 1/2^i) > 1/2$. Hence, if Player A plays an $0 < \varepsilon_i$-optimal strategy in s such that $(1 - \varepsilon_i) \cdot q_i > 1/2$, then the value of the state s_i is greater than 1/2. In that case, in the states c_i, Player B plays the second columns obtaining the value 1/2. This induces that the value in all states q_i is 1/2. However, this is only possible if Player A has (infinite) memory, since the greater the index i considered, the smaller the value of ε_i needs to be to ensure $(1-\varepsilon_i) \cdot q_i \geq 1/2$ while still ensuring $\varepsilon_i > 0$ (since Player A does not have an optimal strategy from s). In particular, for any Player A positional strategy $\mathsf{s_A}$ from q_0 that is $0 < \varepsilon$-optimal in s, the value – w.r.t. the strategy $\mathsf{s_A}$ – of all states s_i for indexes i such that $(1-\varepsilon) \cdot q_i < 1/2$ is smaller than 1/2. In which case, Player B plays the first column in c_i, thus obtaining a value smaller than 1/2. It follows that the value of all states $(q_n)_{n \geq 0}$ – w.r.t. the strategy $\mathsf{s_A}$ – is smaller than 1/2. Hence, any Player A positional strategy is not optimal from q_0.

4.3 Discussion

The main result of this chapter is Theorem 4.11: in finite-state reachability games, for all positive $\varepsilon > 0$, Player A has a positional strategy that is optimal from every state where it is possible to be, and ε-optimal from all other states. Such a theorem does not hold in safety games, since playing almost-optimally may require infinite choice. However, it may be possible to to prove an adaptation of this theorem in safety games. We discuss it further in the discussion of the next chapter, i.e. in Section 5.4.

In fact, with Theorem 4.11, we can actually show that it is decidable if a given state is maximizable w.r.t. Player A. The reason why is because in a standard finite reachability game, given a pair of positional strategies, one for each player, it can be encoded in a decidable theory what is the outcome of the game (i.e. what is the probability to reach the target) with these strategies from any given state. This decidable theory is the first order theory of reals[4]. We have formally proved this result in [69, Theorem 30], which is the arXiv version of [39], on which Section 4.2 of this chapter is based.

4.4 Appendix

4.4.1 . Proof of Proposition 4.8

Proof. First, let us show by induction on k the following property $\mathcal{P}(k)$: assume that there exists a vector $w \in [0,1]^Q$ such that $w \leq \mathsf{m}$, $w \leq g(w)$ and for all $i \in [\![1,n]\!]$, $w(i) < g^{(k)}(w)(i)$. Then, there exists $w' \in [0,1]^n$ such that $w \leq w' \leq \mathsf{m}$ and for all $i \in [\![1,n]\!]$, $w'(i) < g(w')(i)$.

The property $\mathcal{P}(1)$ straightforwardly holds. Consider now some $k \geq 1$

[4]This corresponds to the set of well-formed formulas in first order logic using existential and universal quantifiers along with logical connectors between polynomial (in)equalities. We will use the first order theory of the reals in Section 9.2, hence we will give more details in that section.

and assume that $\mathcal{P}(k)$ holds and assume that there is a $w \in [0,1]^n$ such that $w \leq \mathsf{m}$, $w \leq g(w)$ and for all $i \in [\![1,n]\!]$, $w(i) < g^{(k+1)}(w)(i)$. Note that for all $j \in \mathbb{N}$, we have $g^j(w) \leq \mathsf{m}$. Now, let $n_= = \{i \in [\![1,n]\!] \mid w(i) = g^{(k)}(i)\}$ and $n_\uparrow = [\![1,n]\!] \setminus n_= = \{i \in [\![1,n]\!] \mid w(i) < g^{(k)}(w)(i)\}$. We define:

$$m_= := \min_{i \in n_=} g^{(k+1)}(w)(i) - g^{(k)}(w)(i) = \min_{i \in n_=} g^{(k+1)}(w)(i) - w(i) > 0$$

and:

$$m_\uparrow := \min_{i \in n_\uparrow} g^{(k)}(w)(i) - w(i) > 0$$

Let $m := \min(m_=, m_\uparrow)$ and $w' \in [0,1]^n$ be such that:

- $w'|_{n_=} = w|_{n_=} = g^{(k)}(w)\big|_{n_=}$;
- $w'|_{n_\uparrow} = g^{(k)}(w)\big|_{n_\uparrow} - m/2 \geq w|_{n_\uparrow}$.

With this choice, we have $w' \leq g^{(k)}(w) \leq \mathsf{m}$. Furthermore, we have:

- $w \leq w'$;
- $g^{(k)}(w) - m/2 \leq w'$.

Furthermore, note that:

$$\left\| g^{(k+1)}(w) - g(g^{(k)}(w) - m/2) \right\|_\infty \leq \left\| g^{(k)}(w) - (g^{(k)}(w) - m/2) \right\|_\infty = m/2$$

Hence, for all $i \in [\![1,n]\!]$, we have: $g^{(k+1)}(w)(i) - m/2 \leq g(g^{(k)}(w) - m/2)(i)$. Now, let us show that $w' \leq g(w')$. Let $i \in [\![1,n]\!]$:

- if $i \in n_=$: $w'(i) = w(i) \leq g(w)(i) \leq g(w')(i)$;
- if $i \in n_\uparrow$: $w'(i) = g^{(k)}(w)(i) - m/2 \leq g^{(k+1)}(w)(i) - m/2 \leq g(g^{(k)}(w) - m/2)(i) \leq g(w')(i)$.

We used the fact that $g^{(k)}(w)(i) \leq g^{(k+1)}(w)(i)$, which comes from the fact that $w \leq g(w)$, and the fat that g is non-decreasing. Finally, let us show that, for all $i \in [\![1,n]\!]$, we have $w'(i) < g^{(k)}(w')(i)$. Let $i \in [\![1,n]\!]$.

- if $i \in n_=$: $w'(i) = w(i) \leq g^{(k+1)}(w)(i) - m < g^{(k+1)}(w)(i) - m/2 \leq g(g^{(k)}(w) - m/2)(i) \leq g(w')(i) \leq g^{(k)}(w')(i)$;
- if $i \in n_\uparrow$: $w'(i) = g^{(k)}(w)(i) - m/2 < g^{(k)}(w)(i) \leq g^{(k)}(w')(i)$.

Again, we used the fact that $g(w') \leq g^{(k)}(w')$, which comes from the fact that $w' \leq g(w')$ and g is non-decreasing. We can then apply $\mathcal{P}(k)$ on w' to exhibit a vector $w'' \in [0,1]^Q$ such that $w \leq w' \leq w'' \leq \mathsf{m}$, $w'' \leq g(w'')$ and for all $i \in [\![1,n]\!]$, $w''(i) < g(w'')(i)$. Overall, $\mathcal{P}(k+1)$ holds and therefore $\mathcal{P}(j)$ holds for all $j \in \mathbb{N}$.

Consider some positive $\varepsilon > 0$. Let $\eta := \min_{i \in [\![1,n]\!]} \mathsf{m}(i) > 0$, $\iota := \min(\eta, \varepsilon) > 0$ and $w \in [0,1]^n$ be the valuation such that for all $i \in [\![1,n]\!]$, we have $w(i) := \mathsf{m}(i) - \iota < \mathsf{m}(i)$. First, let us argue that $w \leq g(w)$. Assume towards a contradiction that there is some $i \in [\![1,n]\!]$ such that $g(w)(i) < w(i)$. Then, $g(w)(i) \leq g(\mathsf{m})(i)$ since $w \leq \mathsf{m}$. Furthermore:

$$\mathsf{m}(i) = g(\mathsf{m})(i) \leq g(w)(i) + \|\mathsf{m} - w\| < w(i) + \iota = \mathsf{m}(i)$$

Hence the contradiction. In fact, $w(i) \leq g(w)(i)$ for all $i \in [\![1,n]\!]$. Thus, $w \leq g(w)$. Now, consider the sequence $(w_n)_{n \in \mathbb{N}}$ defined by $w_0 := w$ and for all $k \in \mathbb{N}$, $w_{k+1} := g(w_k) = g^{(k+1)}(w_0)$. We have, for all $k \in \mathbb{N}$, $w_k \leq w_{k+1}$. Hence, this sequence converges. In fact, its limit is equal to m (this directly derives from Kleene fixed-point theorem).

We can conclude that there exists a $k \in \mathbb{N}$ such that, for all $i \in [\![1,n]\!]$, we have $w(i) < w_k(i) = g^{(k)}(w)(i)$ since $w(i) < \mathsf{m}(i)$. We can then apply $\mathcal{P}(k)$ to obtain a valuation $v \in [0,1]^n$ such that $w \leq v \leq \mathsf{m}$ and for all $i \in [\![1,n]\!]$, $g(v)(i) > v(i)$. Furthermore, since $\|\mathsf{m} - v\| \leq \varepsilon$, we have $\|\mathsf{m} - v\| \leq \varepsilon$. □

4.4.2 . Computing the set of maximizable states: formal proofs with arbitrary game forms

In this subsection, we give a detailed proof of Theorem 4.11. To prove this theorem, we will adapt the definitions of the previous subsection to the case of arbitrary game forms and prove the same intermediate lemmas, that is Lemmas 4.9 and 4.10.

Progressive strategies

First, we define the notion of progressive Player-A GF-strategy on arbitrary game forms, as we cannot use Definition 4.8 as is. Indeed, there is no underlying action set in arbitrary game forms, hence, we cannot consider optimal Player-B actions. However, to grasp the idea behind the generalization of Definition 4.8, let us consider a standard game form with infinitely many Player-B actions. Even in that case, where the notion of optimal Player-B actions is defined, there are still two issues with Definition 4.8. First, we should not distinguish between optimal and non-optimal Player-B actions. The reason why is that Player B could have non-optimal actions, such that the gap between its values and the value of the Player-A GF-strategy is arbitrarily close to 0 (which cannot happen if she has only finitely many actions). Second, only requiring that, regardless Player-B GF-strategy, there is a positive probability to reach a good state in Gd is not enough as Player B could have strategies to ensure that this probability is arbitrarily close to 0. This is solved by requiring that the infimum, considered over all Player-B GF-strategies, of the maximum of both of these quantities is positive. With such a generalization, we do obtain a definition that carries over to arbitrary game forms. Before formally defining it, let us first introduce two notations we will use throughout this subsection.

Definition 4.13 (Two Notations)**.** *Consider a state $q \in Q$. For all Player-A GF-strategies $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}^q$, Player-$\mathsf{B}$ GF-strategies $\sigma_\mathsf{B} \in \Sigma_\mathsf{B}^q$ and subsets $S \subseteq Q$ of states, we let:*

- $v^q(\sigma_\mathsf{A}, \sigma_\mathsf{B}) := \mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_\mathcal{G} \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) - \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_\mathcal{G} \rangle](\sigma_\mathsf{A}) \geq 0$*;*
- $p_S^q(\sigma_\mathsf{A}, \sigma_\mathsf{B}) := \varrho_q(\sigma_\mathsf{A}, \sigma_\mathsf{B})[S] = \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_S \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B})$*.*

Definition 4.14 (Progressive strategies in arbitrary game forms)**.** *Consider a state $q \in Q$. Given a set of good states $\mathsf{Gd} \subseteq Q$ that Player A wants to reach, the set of* progressive strategies $\mathsf{Prog}_q(\mathsf{Gd})$ *at state q w.r.t. Gd is defined by*

$$\mathsf{Prog}_q(\mathsf{Gd}) := \{\sigma_\mathsf{A} \in \mathsf{Opt}_\mathsf{A}(\langle \mathsf{F}(q), \mathsf{m}_\mathcal{G} \rangle) \mid \inf_{\sigma_\mathsf{B} \in \Sigma_\mathsf{B}^q} \max(v^q(\sigma_\mathsf{A}, \sigma_\mathsf{B}), p_\mathsf{Gd}^q(\sigma_\mathsf{A}, \sigma_\mathsf{B})) > 0\}$$

In Page 218, we show that both definitions of progressive GF-strategies (i.e. Definition 4.8 and Definition 4.14) coincide on standard finite game forms.

Risky strategies

As for progressive strategies, Definition 4.9 of risky strategies in standard game forms does not carry over to arbitrary game forms. Consider a Player-A GF-strategies $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}^q$ at a state $q \in Q$. The idea is that, given a Player-B GF-strategy, for the GF-strategy σ_A not to be risky, sub-maximizable states may be seen with positive probability only if the outcome of the game in normal with both GF-strategies $\sigma_\mathsf{A}, \sigma_\mathsf{B}$ is greater than the value of the GF-strategy σ_A. However, it is not sufficient to allow that, for any Player-B GF-strategy σ_B, a sub-maximizable is reachable with positive probability as soon as there is an increase in value with σ_B, since that increase may be arbitrarily small. In fact, we need to consider the exact ratio between the gap between the outcome with both GF-strategies $\sigma_\mathsf{A}, \sigma_\mathsf{B}$ and the value of the GF-strategy σ_A and the probability to reach the set of states Bd. This is formally defined below in Definition 4.15.

Definition 4.15 (Risky strategies in arbitrary game forms)**.** *Consider a state $q \in Q$ and a set of bad states $\mathsf{Bd} \subseteq Q$ that Player A wants to avoid. For all Player-A GF-strategies $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}^q$, we let $\mathsf{PosPrb}_\mathsf{Bd}(q, \sigma_\mathsf{A}) := \{\sigma_\mathsf{B} \in \Sigma_\mathsf{B}^q \mid p_\mathsf{Bd}^q(\sigma_\mathsf{A}, \sigma_\mathsf{B}) > 0\}$ be the set of Player-B GF-strategies that induce, with the GF-strategy σ_A, a positive probability to reach the set Bd. The set of* risky strategies $\mathsf{Risk}_q(\mathsf{Gd})$ *at state q w.r.t. Bd is defined by*

$$\mathsf{Risk}_q(\mathsf{Bd}) := \{\sigma_\mathsf{A} \in \mathsf{Opt}_\mathsf{A}(\langle \mathsf{F}(q), \mathsf{m}_\mathcal{G} \rangle) \mid \inf_{\sigma_\mathsf{B} \in \mathsf{PosPrb}_\mathsf{Bd}(q, \sigma_\mathsf{A})} \frac{v^q(\sigma_\mathsf{A}, \sigma_\mathsf{B})}{p_\mathsf{Bd}^q(\sigma_\mathsf{A}, \sigma_\mathsf{B})} = 0\}$$

In Page 219, we show that the two definitions of risky GF-strategies (i.e. Definition 4.9 and Definition 4.15) coincide on standard finite game forms.

Proof of Lemma 4.9

The formal proof of Lemma 4.9 is quite technical and we will need two intermediary lemmas. Contrary to the standard finite game forms case (recall the proof of sketch of Lemma 4.9), we cannot define relevant paths as paths that can occur with positive probability with a Player-B strategy that would play only optimal actions, since we do not consider this notion with arbitrary game forms. Instead, we define below relevant successors (and consequently relevant paths) as successors that Player-B can enforce with positive probability while ensuring an increase in value arbitrarily small.

Definition 4.16 (Relevant successors). *Consider a state $q \in Q$ and a Player-A GF-strategy $\sigma_{\mathsf{A}} \in \Sigma^q_{\mathsf{A}}$. A state $q' \in Q$ is a* relevant successor *of q w.r.t. σ_{A} if there is some positive $\delta > 0$ such that for all $\varepsilon > 0$, there is some $\sigma_{\mathsf{B}} \in \Sigma^q_{\mathsf{B}}$ such that:*

$$v^q(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \leq \varepsilon \text{ and } p^q_{\{q'\}}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \geq \delta$$

We denote by $\mathsf{RelSucc}^q(\sigma_{\mathsf{A}}) \subseteq Q$ the set of all relevant successors of q w.r.t. $\sigma_{\mathsf{A}} \in \Sigma^q_{\mathsf{A}}$.

Then, given a Player-A strategy $\mathsf{s_A} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{A}}$, we let $\mathsf{RelPath}(\mathsf{s_A}) \subseteq Q^+$ denote the set of relevant paths *w.r.t. $\mathsf{s_A}$, i.e. the set of finite sequences of states with only relevant successors:*

$$\mathsf{RelPath}(\mathsf{s_A}) := \{\pi \in Q^+ \mid \forall 1 \leq i < |\pi| - 1,\ \pi_i \in \mathsf{RelSucc}^{\pi_{i-1}}(\mathsf{s_A}(\pi_{\leq i-1}))\}$$

We also define the set of finite paths after which a Player-A strategy does not play either an optimal GF-strategy nor a non-risky GF-strategy (called a problematic path).

Definition 4.17 (Problematic paths). *Given a set of states $\mathsf{Bd} \subseteq Q$ that Player A wants to avoid, for all Player-A strategies $\mathsf{s_A} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{A}}$, we denote by $\mathsf{Prbl}(\mathsf{s_A}) \subseteq Q^+$ the set of* problematic paths, *$\mathsf{Prbl}(\mathsf{s_A}, \mathsf{Bd}) := \{\pi \in Q^+ \mid \mathsf{s_A}(\pi) \notin \mathsf{Opt}_{\mathsf{A}}(\langle \mathsf{F}(\pi_{\mathsf{lt}}), \mathsf{m}_{\mathcal{G}} \rangle) \setminus \mathsf{Risk}_{\pi_{\mathsf{lt}}}(\mathsf{Bd})\}$.*

In fact, a Player-A strategy cannot be optimal after any problematic path, as stated in Lemma 4.13 below.

Lemma 4.13. *Assume that a set of states $\mathsf{Bd} \subseteq Q$ is such that $\mathsf{Bd} \subseteq \mathsf{SubOptQ}_{\mathsf{A}}$. Consider a Player-$\mathsf{A}$ strategy $\mathsf{s_A}$. For all $\pi \in \mathsf{Prbl}(\mathsf{s_A}, \mathsf{Bd})$, the residual strategy $\mathsf{s}^{\mathsf{tl}(\pi)}_{\mathsf{A}}$ is not optimal from π_{lt}.*

Proof. By Lemma 3.10, this holds as soon as $\mathsf{s_A}(\pi) \notin \mathsf{Opt}_{\mathsf{A}}(\langle \mathsf{F}(\pi_{\mathsf{lt}}), \mathsf{m}_{\mathcal{G}} \rangle)$. Assume now that $\mathsf{s_A}(\pi) \in \mathsf{Risk}_{\pi_{\mathsf{lt}}}(\mathsf{Bad})$. For all states $q \in \mathsf{Bad}$, we let $\varepsilon_q > 0$ be such that:

$$\varepsilon_q := \chi_{\mathcal{G}}[\mathsf{A}](q) - \chi_{\mathcal{G}}[\mathsf{s}^{\pi}_{\mathsf{A}}](q)$$

Note that $\varepsilon_q > 0$ since $\mathsf{Bd} \subseteq \mathsf{SubOptQ}_{\mathsf{A}}$. Let $\varepsilon = \min_{q \in \mathsf{Bd}} \varepsilon_q > 0$ since Q is finite. We now define a Player-B strategy $\mathsf{s_B} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$ as follows:

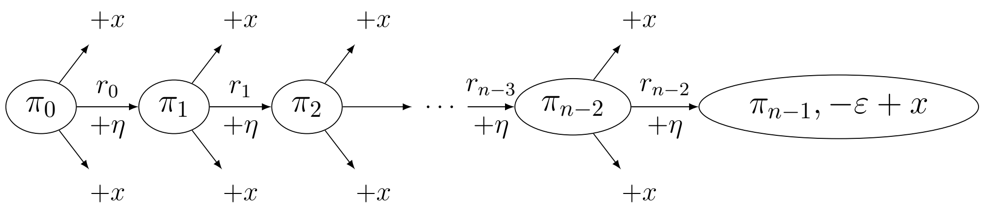

Figure 4.7: An illustration of the finite path π from the proof of Lemma 4.14.

- Since $\mathsf{s_A}(\pi) \in \mathsf{Risk}_{\pi_{\mathsf{lt}}}(\mathsf{Bad})$ there is some $\sigma_\mathsf{B} \in \mathsf{PosPrb}_\mathsf{Bd}(q, \mathsf{s_A}(\pi))$ such that $\frac{v^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi),\sigma_\mathsf{B})}{p^{\pi_{\mathsf{lt}}}_{\mathsf{Bd}}(\mathsf{s_A}(\pi),\sigma_\mathsf{B})} \leq \frac{\varepsilon}{3}$. We set $\mathsf{s_B}(\pi_{\mathsf{lt}}) := \sigma_\mathsf{B}$.
- We let $x := \frac{\varepsilon}{3} \cdot p^{\pi_{\mathsf{lt}}}_{\mathsf{Bd}}(\sigma_\mathsf{A}, \sigma_\mathsf{B})$. For all $q \in Q$, the residual strategy $\mathsf{s}_\mathsf{B}^{\pi_{\mathsf{lt}}}$ is chosen so that it is x-optimal against the strategy $\mathsf{s_A}$, that is for all $q \in Q$, we have $\mathbb{P}^{\mathsf{s}_\mathsf{A}^\pi, \mathsf{s}_\mathsf{B}^{\pi_{\mathsf{lt}}}}_{\mathcal{C},q}[\mathsf{Reach}] \leq \chi_\mathcal{G}[\mathsf{s}_\mathsf{A}^\pi](q) + x$

We obtain, denoting $p := \mathbb{P}^{\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\pi)}, \mathsf{s_B}}_{\mathcal{C},\pi_{\mathsf{lt}}}[\mathsf{Reach}]$:

$$
\begin{aligned}
p &= \sum_{q \in Q} \mathbb{P}^{\mathsf{s_A}(\pi), \mathsf{s_B}(\pi_{\mathsf{lt}})}_{\mathcal{C}}(\pi_{\mathsf{lt}}, q) \cdot \mathbb{P}^{\mathsf{s}_\mathsf{A}^\pi, \mathsf{s}_\mathsf{B}^{\pi_{\mathsf{lt}}}}_{\mathcal{C},q}[\mathsf{Reach}] \\
&\leq \sum_{q \in Q} \mathsf{out}[\langle \mathsf{F}(\pi_{\mathsf{lt}}), q \rangle](\mathsf{s_A}(\pi), \mathsf{s_B}(\pi_{\mathsf{lt}})) \cdot \chi_\mathcal{G}[\mathsf{s}_\mathsf{A}^\pi](q) + x && \text{by Def. 1.28} \\
&\leq \sum_{q \in \mathsf{Bd}} \mathsf{out}[\langle \mathsf{F}(\pi_{\mathsf{lt}}), q \rangle](\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \cdot (\chi_\mathcal{G}[\mathsf{A}](q) - \varepsilon) \\
&+ \sum_{q \in Q \setminus \mathsf{Bd}} \mathsf{out}[\langle \mathsf{F}(\pi_{\mathsf{lt}}), q \rangle](\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \cdot \chi_\mathcal{G}[\mathsf{A}](q) + x && \text{by Def. of } \mathsf{s_B} \\
&= \mathsf{out}[\langle \mathsf{F}(\pi_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}] \rangle](\mathsf{s_A}(\pi), \sigma_\mathsf{B}) - p^{\pi_{\mathsf{lt}}}_{\mathsf{Bd}}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \cdot \varepsilon + x && \text{by Def. of } p^{\pi_{\mathsf{lt}}}_{\mathsf{Bd}} \\
&= v^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) + \mathsf{val}[\langle \mathsf{F}(\pi_{\mathsf{lt}}), \chi_\mathcal{G}[\mathsf{A}] \rangle](\mathsf{s_A}(\pi)) - p^{\pi_{\mathsf{lt}}}_{\mathsf{Bd}}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \cdot \frac{2\varepsilon}{3} && \text{by Def. of } v^{\pi_{\mathsf{lt}}} \\
&\leq v^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) + \chi_\mathcal{G}[\mathsf{A}](\pi_{\mathsf{lt}}) - p^{\pi_{\mathsf{lt}}}_{\mathsf{Bd}}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \cdot \frac{2\varepsilon}{3} && \text{by Prop. 3.9} \\
&\leq \chi_\mathcal{G}[\mathsf{A}](\pi_{\mathsf{lt}}) - p^{\pi_{\mathsf{lt}}}_{\mathsf{Bd}}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \cdot \frac{\varepsilon}{3} && \text{by Def. of } \sigma_\mathsf{B}
\end{aligned}
$$

That is, $\mathbb{P}^{\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\pi)}, \mathsf{s_B}}_{\mathcal{C},\pi_{\mathsf{lt}}}[\mathsf{Reach}] < \chi_\mathcal{G}[\mathsf{A}](\pi_{\mathsf{lt}})$ since $p^{\pi_{\mathsf{lt}}}_{\mathsf{Bd}}(\sigma_\mathsf{A}, \sigma_\mathsf{B}) > 0$ as $\sigma_\mathsf{B} \in \mathsf{PosPrb}_\mathsf{Bd}(q, \sigma_\mathsf{A})$. Therefore, the residual strategy $\mathsf{s}_\mathsf{A}^{\mathsf{tl}(\pi)}$ is not optimal from π_{lt}. □

Furthermore, as soon as for a Player-A strategy, there is a relevant path w.r.t. that strategy that is problematic, then this strategy is not optimal (given that the set of states that Player A wants to avoid is a subset of sub-maximizable states).

Lemma 4.14. *Assume that a set of states* $\mathsf{Bd} \subseteq Q$ *is such that* $\mathsf{Bd} \subseteq \mathsf{SubOptQ_A}$. *Consider a Player-*$\mathsf{A}$ *strategy* $\mathsf{s_A}$ *and a state* $q_0 \in Q \setminus T$. *Assume that* $\mathsf{Prbl}(\mathsf{s_A}, \mathsf{Bd}) \cap \mathsf{RelPath}(\mathsf{s_A}) \cap q_0 \cdot (Q \setminus T)^* \neq \emptyset$. *Then, the strategy* $\mathsf{s_A}$ *is not optimal from* q_0.

Proof. Consider a path $\pi \in \mathsf{Prbl}(\mathsf{s_A}, \mathsf{Bd}) \cap \mathsf{RelPath}(\mathsf{s_A}) \cap q_0 \cdot (Q \setminus T)^*$. Let $n := |\pi|$. Recall Definition 4.16: for all $0 \leq i \leq n-2$, we let $r_i > 0$ be such that, for all $\theta > 0$, there is some $\sigma_\mathsf{B} \in \Sigma_\mathsf{B}^{\pi_i}$ such that:

$$v^{\pi_i}(\mathsf{s_A}(\pi_{\leq i}), \sigma_\mathsf{B}) \leq \theta \text{ and } p^{\pi_i}_{\{\pi_{i+1}\}}(\mathsf{s_A}(\pi_{\leq i}), \sigma_\mathsf{B}) \geq r_i \tag{4.1}$$

Let $r := \Pi_{i=0}^{n-2} r_i > 0$. Since, by Lemma 4.13, the residual strategy $\mathsf{s_A}^{\mathsf{tl}(\pi)} = \mathsf{s_A}^{\pi_{\leq n-2}}$ is not optimal from $\pi_{\mathsf{lt}} = \pi_{n-1}$, we let $\varepsilon := \chi_\mathcal{G}[\mathsf{A}](\pi_{n-1}) - \chi_\mathcal{G}[\mathsf{s_A}^{\pi_{\leq n-2}}](\pi_{n-1}) > 0$. Let also $\eta := \frac{\varepsilon \cdot r}{(n-1)\cdot 3} > 0$ and $x := \frac{\varepsilon \cdot r}{3} > 0$. For all $0 \leq i \leq n-2$, we let $\sigma_\mathsf{B}^i \in \Sigma_\mathsf{B}^{\pi_{i-1}}$ as in Equation 4.1 for $\theta := \eta$. We can now define a Player-B strategy $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$ in the following way:

- for all $0 \leq i \leq n-2$, we set $\mathsf{s_B}(\pi_{\leq i}) := \sigma_\mathsf{B}^i$;
- for all $0 \leq i \leq n-3$, the residual strategy $\mathsf{s_B}^{\pi_{\leq i}}$ is chosen x-optimal against $\mathsf{s_A}$ from all states but π_{i+1}: that is, for all $q \in Q \setminus \{\pi_{i+1}\}$, we have: $\mathbb{P}_{\mathcal{C},q}^{\mathsf{s_A}^{\pi_{\leq i}}, \mathsf{s_B}^{\pi_{\leq i}}}[\mathsf{Reach}] \leq \chi_\mathcal{G}[\mathsf{s_A}^{\pi_{\leq i}}](q) + x \leq \chi_\mathcal{G}[\mathsf{A}](q) + x$.
- the residual strategy $\mathsf{s_B}^{\pi_{\leq n-2}}$ is chosen x-optimal against $\mathsf{s_A}$: that is, for all $q \in Q$: $\mathbb{P}_{\mathcal{C},q}^{\mathsf{s_A}^{\pi_{\leq n-2}}, \mathsf{s_B}^{\pi_{\leq n-2}}}[\mathsf{Reach}] \leq \chi_\mathcal{G}[\mathsf{s_A}^{\pi_{\leq n-2}}](q) + x \leq \chi_\mathcal{G}[\mathsf{A}](q) + x$.

An illustration of the paths and some quantities involved in this proof is given in Figure 4.7. Roughly, this can read as follows. From π_0, there is probability at least r_0 to go to π_1. If another state is reached, then the Player-B strategy is chosen so that the value from there increases of at most x compared to the Player-A value of the state. Furthermore, by the choice of the strategy $\mathsf{s_B}(\pi_0)$, the expected Player-A value of the successors of π_0 has increased by at most η w.r.t. the Player-A value of π_0. This is repeated all along the path until $\pi_{\mathsf{lt}} = \pi_{n-1}$ is reached, from which we know that the Player-A strategy is not optimal. (The value of the Player-A strategy at π_{lt} is equal to the value of the Player-A value of π_{lt} minus ε.) These quantities are chosen so that $r \cdot \varepsilon > x + (n-1) \cdot \eta$ where $r = \Pi_0^{n-2} r_i$, i.e. the expected loss in the value — due to the Player-A strategy not being optimal at π_{lt} is greater then the increase in the value due to how the Player-B strategy $\mathsf{s_B}$ is defined.

Now let us show the equation below, for all $0 \leq i \leq n-2$:

$$\sum_{q \in Q} \mathbb{P}_{\mathcal{C},\pi_0}^{\mathsf{s_A},\mathsf{s_B}}(\pi_{1\ldots i} \cdot q) \cdot \chi_\mathcal{G}[\mathsf{A}](q) \leq \eta + \mathbb{P}_{\mathcal{C},\pi_0}^{\mathsf{s_A},\mathsf{s_B}}(\pi_{1\ldots i}) \cdot \chi_\mathcal{G}[\mathsf{A}](\pi_i) \tag{4.2}$$

where, $\pi_{1\ldots i}$ refers to the finite path $\pi_1 \cdots \pi_i$ — it is equal to ϵ when $i = 0$. Informally, this equation states that the expected value of the successors of π_i

is at most the Player-A value of the state π_i plus η, which is the margin in the increase of value that we chose for σ_{B}^i (since it was chosen from Equation 4.1 with $\theta = \eta$).

Let $0 \leq i \leq n-2$. Letting $p := \sum_{q \in Q} \mathbb{P}_{\mathcal{C},\pi_0}^{\mathsf{s_A},\mathsf{s_B}}(\pi_{1 \ldots i} \cdot q) \cdot \chi_{\mathcal{G}}[\mathsf{A}](q)$, we have:

$$
\begin{aligned}
p &= \mathbb{P}_{\mathcal{C},\pi_0}^{\mathsf{s_A},\mathsf{s_B}}(\pi_{1 \ldots i}) \cdot \left(\sum_{q \in Q} \mathbb{P}_{\mathcal{C}}^{\mathsf{s_A}(\pi_{\leq i}),\mathsf{s_B}(\pi_{\leq i})}(\pi_i, q) \cdot \chi_{\mathcal{G}}[\mathsf{A}](q) \right) \\
&= \mathbb{P}_{\mathcal{C},\pi_0}^{\mathsf{s_A},\mathsf{s_B}}(\pi_{1 \ldots i}) \cdot \left(\sum_{q \in Q} \mathsf{out}[\langle \mathsf{F}(\pi_i), q \rangle](\mathsf{s_A}(\pi_{\leq i}), \mathsf{s_B}(\pi_{\leq i})) \cdot \chi_{\mathcal{G}}[\mathsf{A}](q) \right) \\
&= \mathbb{P}_{\mathcal{C},\pi_0}^{\mathsf{s_A},\mathsf{s_B}}(\pi_{1 \ldots i}) \cdot \mathsf{out}[\langle \mathsf{F}(\pi_i), \chi_{\mathcal{G}}[\mathsf{A}] \rangle](\mathsf{s_A}(\pi_{\leq i}), \mathsf{s_B}(\pi_{\leq i})) \\
&= \mathbb{P}_{\mathcal{C},\pi_0}^{\mathsf{s_A},\mathsf{s_B}}(\pi_{1 \ldots i}) \cdot \mathsf{out}[\langle \mathsf{F}(\pi_i), \chi_{\mathcal{G}}[\mathsf{A}] \rangle](\mathsf{s_A}(\pi_{\leq i}), \sigma_{\mathsf{B}}^i) \\
&= \mathbb{P}_{\mathcal{C},\pi_0}^{\mathsf{s_A},\mathsf{s_B}}(\pi_{1 \ldots i}) \cdot (\mathsf{val}[\langle \mathsf{F}(\pi_i), \chi_{\mathcal{G}}[\mathsf{A}] \rangle](\mathsf{s_A}(\pi_{\leq i})) + v^{\pi_i}(\mathsf{s_A}(\pi_{\leq i}), \sigma_{\mathsf{B}}^i)) \\
&\leq \mathbb{P}_{\mathcal{C},\pi_0}^{\mathsf{s_A},\mathsf{s_B}}(\pi_{1 \ldots i}) \cdot \mathsf{val}[\langle \mathsf{F}(\pi_i), \chi_{\mathcal{G}}[\mathsf{A}] \rangle][\mathsf{A}] + \eta \\
&= \mathbb{P}_{\mathcal{C},\pi_0}^{\mathsf{s_A},\mathsf{s_B}}(\pi_{1 \ldots i}) \cdot \chi_{\mathcal{G}}[\mathsf{A}](\pi_i) + \eta
\end{aligned}
$$

where the last equality comes from Proposition 3.9. We do obtain Equation 4.2. Now, for the readability of the series of (in)equalities below, we let $p := \mathbb{P}_{\mathcal{C},\pi_0}^{\mathsf{s_A},\mathsf{s_B}}[\mathsf{Reach}]$ and, for all $0 \leq i \leq n-1$ and $q \in Q$, $p_i(q) := \mathbb{P}_{\mathcal{C},\pi_0}^{\mathsf{s_A},\mathsf{s_B}}(\pi_{1 \ldots i} \cdot q)$

and $p_i := \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi_0}(\pi_{1\ldots i})$. We have:

$$
\begin{aligned}
p &= \sum_{i=0}^{n-2} \sum_{q \in Q \setminus \{\pi_{i+1}\}} p_i(q) \cdot \mathbb{P}^{\mathsf{s_A}^{\pi_{\leq i}},\mathsf{s_B}^{\pi_{\leq i}}}_{\mathcal{C},q}[\mathsf{Reach}] + p_{n-1} \cdot \mathbb{P}^{\mathsf{s_A}^{\pi_{\leq n-2}},\mathsf{s_B}^{\pi_{\leq n-2}}}_{\mathcal{C},\pi_{n-1}}[\mathsf{Reach}] \\
&\leq \sum_{i=0}^{n-2} \sum_{q \in Q \setminus \{\pi_{i+1}\}} p_i(q) \cdot (\chi_{\mathcal{G}}[\mathsf{A}](q) + x) + p_{n-1} \cdot (\chi_{\mathcal{G}}[\mathsf{s_A}^{\pi_{\leq n-2}}](\pi_{n-1}) + x) \\
&= \sum_{i=0}^{n-2} \sum_{q \in Q \setminus \{\pi_{i+1}\}} p_i(q) \cdot \chi_{\mathcal{G}}[\mathsf{A}](q) + p_{n-1} \cdot (\chi_{\mathcal{G}}[\mathsf{A}](\pi_{n-1}) - \varepsilon) + x \qquad (0) \\
&= \sum_{i=0}^{n-3} \sum_{q \in Q \setminus \{\pi_{i+1}\}} p_i(q) \cdot \chi_{\mathcal{G}}[\mathsf{A}](q) + \sum_{q \in Q} p_{n-2}(q) \cdot \chi_{\mathcal{G}}[\mathsf{A}](q) + x - \varepsilon \cdot p_{n-1} \qquad (1) \\
&\leq \sum_{i=0}^{n-3} \sum_{q \in Q \setminus \{\pi_{i+1}\}} p_i(q) \cdot \chi_{\mathcal{G}}[\mathsf{A}](q) + p_{n-2} \cdot \chi_{\mathcal{G}}[\mathsf{A}](\pi_{n-2}) + \eta + x - \varepsilon \cdot p_{n-1} \qquad (1') \\
&= \sum_{i=0}^{n-4} \sum_{q \in Q \setminus \{\pi_{i+1}\}} p_i(q) \cdot \chi_{\mathcal{G}}[\mathsf{A}](q) + \sum_{q \in Q} p_{n-3}(q) \cdot \chi_{\mathcal{G}}[\mathsf{A}](q) + \eta + x - \varepsilon \cdot p_{n-1} \quad (2) \\
&\leq \sum_{i=0}^{n-4} \sum_{q \in Q \setminus \{\pi_{i+1}\}} p_i(q) \cdot \chi_{\mathcal{G}}[\mathsf{A}](q) + p_{n-3} \cdot \chi_{\mathcal{G}}[\mathsf{A}](\pi_{n-3}) + 2 \cdot \eta + x - \varepsilon \cdot p_{n-1} \quad (2') \\
&\ldots \\
&\leq p_0 \cdot \chi_{\mathcal{G}}[\mathsf{A}](\pi_0) + (n-1) \cdot \eta + x - \varepsilon \cdot p_{n-1} \\
&= \chi_{\mathcal{G}}[\mathsf{A}](\pi_0) + \frac{\varepsilon \cdot r}{3} + \frac{\varepsilon \cdot r}{3} - \varepsilon \cdot p_{n-1}
\end{aligned}
$$

The equalities from (0) to (1) and from (1′) to (2) are obtained by realizing that $p_{n-1} = p_{n-2}(\pi_{n-1})$ and $p_{n-2} = p_{n-3}(\pi_{n-2})$. Furthermore, the inequalities from (1) to (1′) and from for (2) to (2′) are obtained by applying Equation 4.2. This is by iterating the application of this equation that we obtain the last inequality. In addition, we have:

$$
\begin{aligned}
p_{n-1} &= \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}(\pi) = \Pi_{i=0}^{n-2} \mathbb{P}^{\mathsf{s_A}(\pi_{\leq i}),\mathsf{s_B}(\pi_{\leq i})}_{\mathcal{C}}(\pi_i, \pi_{i+1}) \\
&= \Pi_{i=0}^{n-2} p^{\pi_i}_{\{\pi_{i+1}\}}(\mathsf{s_A}(\pi_{\leq i}), \sigma^i_{\mathsf{B}}) \geq \Pi_{i=0}^{n-1} r_i = r
\end{aligned}
$$

Overall, we obtain:

$$
\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi_0}[\mathsf{Reach}] \leq \chi_{\mathcal{G}}[\mathsf{A}](\pi_0) - \frac{\varepsilon \cdot r}{3} < \chi_{\mathcal{G}}[\mathsf{A}](\pi_0)
$$

Hence, the Player-A strategy $\mathsf{s_A}$ is not optimal from $\pi_0 = q_0$. □

We can now proceed to the proof of Lemma 4.9.

Proof. Consider a state $q_0 \in Q \setminus \mathsf{Sec}(\mathsf{Bd})$ and a Player-A strategy $\mathsf{s_A}$. Let $u := \chi_{\mathcal{G}}[\mathsf{A}](q_0) > 0$. If we have $\mathsf{Prbl}(\mathsf{s_A}, \mathsf{Bd}) \cap \mathsf{RelPath}(\mathsf{s_A}) \cap q_0 \cdot (Q \setminus T)^* \neq \emptyset$, then Lemma 4.14 gives that the strategy $\mathsf{s_A}$ is not optimal from q. Hence, let us assume that $\mathsf{Prbl}(\mathsf{s_A}, \mathsf{Bd}) \cap \mathsf{RelPath}(\mathsf{s_A}) \cap q_0 \cdot (Q \setminus T)^* = \emptyset$.

We let $\mathsf{Path}(\mathsf{s_A}) := \{\rho \in (Q \setminus \mathsf{Sec}(\mathsf{Bd}))^* \mid q_0 \cdot \rho \in \mathsf{RelPath}(\mathsf{s_A})\}$. We also consider a sequence $(\varepsilon_n)_{n\geq 1}$ of positive reals such that $\sum_{n\geq 1} \varepsilon_n < u$. Let us define a Player-B strategy $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$, by induction on $i \in \mathbb{N}$, on paths in $Q^i \cap \mathsf{Path}(\mathsf{s_A})$ starting at state q_0 ensuring:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[Q^i \setminus \mathsf{Path}(\mathsf{s_A})] \leq \sum_{n=1}^{i} \varepsilon_n$$

Informally, this means that Player B is able to ensure that, with very large probability, only relevant non-secure states are seen. Hence, Player B will be able to ensure that the probability to reach the target from q_0 is less than u.

This straightforwardly holds for $i = 0$ since $\epsilon \in \mathsf{Path}(\mathsf{s_A})$. Assume now that this holds for some $i \in \mathbb{N}$. Let $\rho \in Q^i \cap \mathsf{Path}(\mathsf{s_A})$ and $\pi := q_0 \cdot \rho$. The goal is to define a Player-B GF-strategy $\sigma_\mathsf{B} \in \Sigma_\mathsf{B}^{\pi_{\mathsf{lt}}}$ such that:

$$p^{\pi_{\mathsf{lt}}}_{\mathsf{Sec}(\mathsf{Bd}) \cup Q \setminus \mathsf{RelSucc}^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi))}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \leq \varepsilon_{i+1} \tag{4.3}$$

Let $k := |Q \setminus \mathsf{RelSucc}^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi))| + 1 \in \mathbb{N}$. For all $q \in Q \setminus \mathsf{RelSucc}^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi))$, there is some $\varepsilon_q > 0$ such that, for all $\sigma_\mathsf{B} \in \Sigma_\mathsf{B}^{\pi_{\mathsf{lt}}}$, if $v^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \leq \varepsilon_q$ then $p^{\pi_{\mathsf{lt}}}_{\{q\}}(\sigma_\mathsf{A}, \sigma_\mathsf{B}) < \frac{\varepsilon_{i+1}}{k}$. Let $\varepsilon := \min_{q \in Q \setminus \mathsf{RelSucc}^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi))} \varepsilon_q$.

Furthermore, since $\pi \in \mathsf{RelPath}(\mathsf{s_A}) \cap q_0 \cdot (Q \setminus T)^*$ (since $T \subseteq \mathsf{Sec}(\mathsf{Bd})$), we have by assumption $\pi \notin \mathsf{Prbl}(\mathsf{s_A}, \mathsf{Bd})$. That is, $\mathsf{s_A}(\pi) \in \mathsf{Opt_A}(\langle \mathsf{F}(\pi_{\mathsf{lt}}), \chi_{\mathcal{G}}[\mathsf{A}]\rangle) \setminus \mathsf{Risk}_{\pi_{\mathsf{lt}}}(\mathsf{Bd})$. However, since $\pi_{\mathsf{lt}} \in Q \setminus \mathsf{Sec}(\mathsf{Bd})$, it must be that $\mathsf{s_A}(\pi) \notin \mathsf{Prog}_{\pi_{\mathsf{lt}}}(\mathsf{Sec}(\mathsf{Bd}))$. Hence,

$$\inf_{\sigma_\mathsf{B} \in \Sigma_\mathsf{B}(\mathsf{F}(\pi_{\mathsf{lt}}))} \max(v^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}), p^{\pi_{\mathsf{lt}}}_{\mathsf{Sec}(\mathsf{Bd})}(\mathsf{s_A}(\pi), \sigma_\mathsf{B})) = 0$$

We can therefore consider some $\sigma_\mathsf{B} \in \Sigma_\mathsf{B}^{\pi_{\mathsf{lt}}}$ such that:

$$\max(v^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}), p^{\pi_{\mathsf{lt}}}_{\mathsf{Sec}(\mathsf{Bd})}(\mathsf{s_A}(\pi), \sigma_\mathsf{B})) \leq \min(\varepsilon, \frac{\varepsilon_{i+1}}{k})$$

In particular, this GF-strategy ensures that, we have $p^{\pi_{\mathsf{lt}}}_{\mathsf{Sec}(\mathsf{Bad})}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \leq \frac{\varepsilon_{i+1}}{k}$. Furthermore, for all $q \in Q \setminus \mathsf{RelSucc}^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi))$, by definition of ε_q and ε, we also have $p^{\pi_{\mathsf{lt}}}_{\{q\}}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \leq \frac{\varepsilon_{i+1}}{k}$. For such a GF-strategy σ_B we have:

$$\begin{aligned}
p^{\pi_{\mathsf{lt}}}_{\mathsf{Sec}(\mathsf{Bd}) \cup Q \setminus \mathsf{RelSucc}^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi))}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) &\leq p^{\pi_{\mathsf{lt}}}_{\mathsf{Sec}(\mathsf{Bd})}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \\
&\quad + p^{\pi_{\mathsf{lt}}}_{Q \setminus \mathsf{RelSucc}^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi))}(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \\
&\leq \frac{\varepsilon_{i+1}}{k} + \sum_{q \in Q \setminus \mathsf{RelSucc}^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi))} p^{\pi_{\mathsf{lt}}}_q(\mathsf{s_A}(\pi), \sigma_\mathsf{B}) \\
&\leq \frac{\varepsilon_{i+1}}{k} + \sum_{q \in Q \setminus \mathsf{RelSucc}^{\pi_{\mathsf{lt}}}(\mathsf{s_A}(\pi))} \frac{\varepsilon_{i+1}}{k} = \varepsilon_{i+1}
\end{aligned}$$

We obtain Equation 4.3. We set $\mathsf{s_B}(\pi) := \sigma_\mathsf{B}$. This is done for all paths $\pi := q_0 \cdot \rho$ for $\rho \in Q^i \cap \mathsf{Path}(\mathsf{s_A})$. We obtain:

$$\begin{aligned}\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[Q^{i+1} \setminus \mathsf{Path}(\mathsf{s_A})] = \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[Q^{\leq i} \setminus \mathsf{Path}(\mathsf{s_A})] \\ + \sum_{\rho \in Q^i \cap \mathsf{Path}(\mathsf{s_A})} \sum_{\substack{q \in \mathsf{Sec}(\mathsf{Bd}) \\ \cup Q \setminus \mathsf{RelSucc}^{\pi_{\mathsf{lt}}}_{\mathsf{s_A}(\pi)}}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}(\rho) \cdot p^{\rho_{\mathsf{lt}}}_q(\mathsf{s_A}(q_0 \cdot \rho), \mathsf{s_B}(q_0 \cdot \rho)) \\ \leq \sum_{n=1}^{i} \varepsilon_n + \sum_{\rho \in Q^i \cap \mathsf{Path}(\mathsf{s_A})} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}(\rho) \cdot \varepsilon_{i+1} = \sum_{n=1}^{i+1} \varepsilon_n\end{aligned}$$

Hence, the property holds at index $i+1$. In fact, it holds for all $i \in \mathbb{N}$. With such a strategy $\mathsf{s_B}$, we have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[Q^* \setminus \mathsf{Path}(\mathsf{s_A})] \leq \sum_{n=1}^{\infty} \varepsilon_n < u$$

Furthermore, since $\mathsf{Path}(\mathsf{s_A}) \subseteq (Q \setminus \mathsf{Sec}(\mathsf{Bd}))^*$ and $T \subseteq \mathsf{Sec}(\mathsf{Bd})$, we have

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[\mathsf{Reach} \cap \mathsf{Path}(\mathsf{s_A})^\omega] = 0$$

That is:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[\mathsf{Reach}] \leq \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[Q^* \setminus \mathsf{Path}(\mathsf{s_A})] < u = \chi_\mathcal{G}[\mathsf{A}](q)$$

That is, the Player-A strategy $\mathsf{s_A}$ is not optimal from $q \in Q \setminus \mathsf{Sec}(\mathsf{Bd})$. □

Proof of Lemma 4.10

The idea of the proof of this lemma is close to the informal ideas given as proof sketch. However, the exact details are quite technical, although we do not need intermediate lemmas to establish this result.

Proof. Let $\varepsilon > 0$. We define a positional Player-A strategy $\mathsf{s_A} \in \mathsf{S}^\mathcal{C}_\mathsf{A}$ along with a valuation of states $v : Q \to [0,1]$. First, we let $v|_{\mathsf{Sec}(\mathsf{Bad})} := \mathsf{m}_\mathcal{G}|_{\mathsf{Sec}(\mathsf{Bad})}$. For all $q \in \mathsf{Sec}(\mathsf{Bad})$, if $\mathsf{m}_\mathcal{G}(q) = 0$, we define $\mathsf{s_A}(q)$ arbitrarily. It is also the case is $q \in T$, since all states in T are assume self-looping sinks. Otherwise, we let $i_q \in \mathbb{N}$ be the least integer such that $q \in \mathsf{Sec}_{i_q}(\mathsf{Bad})$. We let $\mathsf{s_A}(q) \in \mathsf{Eff}_q(\mathsf{Sec}_{i_q-1}(\mathsf{Bad}), \mathsf{Bad})$. Since $\mathsf{s_A}(q) \notin \mathsf{Risk}_q(\mathsf{Bad})$, it follows that:

$$e_q := \inf_{\sigma_\mathsf{B} \in \mathsf{PosPrb}_\mathsf{Bad}(q, \mathsf{s_A}(q))} \frac{v^q(\mathsf{s_A}(q), \sigma_\mathsf{B})}{p^q_\mathsf{Bad}(\mathsf{s_A}(q), \sigma_\mathsf{B})} > 0$$

We then let $e := \min_{q \in \mathsf{Sec}(\mathsf{Bad})} e_q > 0$. We use Proposition 4.8 to define the valuation v on the states in Bad. Indeed, since the operator $\Delta_\mathcal{G}$ is non-decreasing and 1-Lipschitz (by Lemma 4.6), it follows that we can define $v|_\mathsf{Bad}$ such that:

- $\mathsf{m}_{\mathcal{G}}|_{\mathsf{Bad}} - \min(e, \varepsilon) \leq v|_{\mathsf{Bad}} \leq v|_{\mathsf{m}_{\mathcal{G}}}$;
- for all $q \in \mathsf{Bad}$, we have $v(q) < \Delta_{\mathcal{G}}(v)(q)$.

For all $q \in \mathsf{Bad}$, since $\Delta_{\mathcal{G}}(v)(q) = \mathsf{val}[\langle \mathsf{F}(q), v \rangle](\mathsf{A})$, we can then define $\mathsf{s}_\mathsf{A}(q) \in \Sigma_\mathsf{A}^q$ such that $v(q) < \mathsf{val}[\langle \mathsf{F}(q), v \rangle](\mathsf{s}_\mathsf{A}(q))$.

This concludes the definitions of v and s_A. Let us show that the strategy s_A guarantees the valuation v by applying Corollary 3.14. First, let us show that the strategy s_A dominates the valuation v. This holds for all states in Bad. Consider now some state $q \in \mathsf{Sec}(\mathsf{Bad})$ and Player-B GF-strategy $\sigma_\mathsf{B} \in \Sigma_\mathsf{B}^q$. There are two possibilities:

- Assume that $\sigma_\mathsf{B} \notin \mathsf{PosPrb}_{\mathsf{Bad}}(q, \mathsf{s}_\mathsf{A}(q))$, that is we have $p_{\mathsf{Bad}}^q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B}) = \varrho_q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B})[\mathsf{Bad}] = 0$. In that case, we have:

$$\begin{aligned}
\mathsf{out}[\langle \mathsf{F}(q), v \rangle](\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B}) &= \sum_{q' \in Q} \varrho_q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B})(q') \cdot v(q') \\
&= \sum_{q' \in \mathsf{Sec}(\mathsf{Bad})} \varrho_q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B})(q') \cdot v(q') \\
&= \sum_{q' \in \mathsf{Sec}(\mathsf{Bad})} \varrho_q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B})(q') \cdot \mathsf{m}_{\mathcal{G}}(q') \\
&= \mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B}) \geq \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{s}_\mathsf{A}(q)) \\
&= \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{A}) = \Delta(\mathsf{m}_{\mathcal{G}})(q) = \mathsf{m}_{\mathcal{G}}(q) = v(q)
\end{aligned}$$

  Note that we have $\mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{s}_\mathsf{A}(q)) = \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{A})$ because $\mathsf{s}_\mathsf{A}(q) \in \mathsf{Opt}_\mathsf{A}(\mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle])$.

- Assume now that $\sigma_\mathsf{B} \in \mathsf{PosPrb}_{\mathsf{Bad}}(q, \mathsf{s}_\mathsf{A}(q))$. This implies

$$v^q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B}) \geq e \cdot p_{\mathsf{Bad}}^q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B})$$

  with

$$v^q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B}) = \mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B}) - \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{s}_\mathsf{A}(q)) \geq 0$$

  and $\mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{s}_\mathsf{A}(q)) = \mathsf{m}_{\mathcal{G}}(q) = v(q)$. Therefore, we have:

$$\begin{aligned}
\mathsf{out}[\langle \mathsf{F}(q), v \rangle](\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B}) &= \sum_{q' \in Q} \varrho_q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B})(q') \cdot v(q') \\
&\geq \sum_{q' \in \mathsf{Sec}(\mathsf{Bad})} \varrho_q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B})(q') \cdot \mathsf{m}_{\mathcal{G}}(q') \\
&+ \sum_{q' \in \mathsf{Bad}} \varrho_q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B})(q') \cdot (\mathsf{m}_{\mathcal{G}}(q') - \min(e, \varepsilon)) \\
&= \mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B}) - p_{\mathsf{Bad}}^q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B}) \cdot \min(e, \varepsilon) \\
&\geq \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{s}_\mathsf{A}(q)) + p_{\mathsf{Bad}}^q(\mathsf{s}_\mathsf{A}(q), \sigma_\mathsf{B}) \cdot (e - \min(e, \varepsilon)) \\
&\geq \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\mathsf{s}_\mathsf{A}(q)) = v(q)
\end{aligned}$$

In both cases, we have $\mathsf{out}[\langle \mathsf{F}(q), v\rangle](\mathsf{s_A}(q), \sigma_\mathsf{B}) \geq v(q)$. As this holds for all $q \in \mathsf{Sec}(\mathsf{Bad})$, it follows that the strategy $\mathsf{s_A}$ dominates the valuation v.

Let us now consider the second condition of Corollary 3.14

By definition of the strategy $\mathsf{s_A}$, for all $q \in \mathsf{Bad}$, we have $d_q := \mathsf{val}[\langle \mathsf{F}(q), v\rangle](\mathsf{s_A}(q)) - v(q) > 0$. We let $d := \min_{q \in \mathsf{Bad}} d_q > 0$. Furthermore, for all $u \in v[Q]$, we let $Q^v_{>u} := \{q \in Q \mid v(q) > u\}$ and $Q^v_{\leq u} := Q \setminus Q^v_{>u}$. Consider now some state $q \in \mathsf{Bad}$. Let $u := v(q)$. For any Player-B GF-strategy $\sigma_\mathsf{B} \in \Sigma^q_\mathsf{B}$, we have:

$$\begin{aligned}
u + d \leq \mathsf{out}[\langle \mathsf{F}(q), v\rangle](\mathsf{s_A}(q), \sigma_\mathsf{B}) &= \sum_{q' \in Q} \varrho_q(\mathsf{s_A}(q), \sigma_\mathsf{B})(q') \cdot v(q') \\
&\leq \sum_{q' \in Q^v_{>u}} \varrho_q(\mathsf{s_A}(q), \sigma_\mathsf{B})(q') + \sum_{q' \in Q^v_{\leq u}} \varrho_q(\mathsf{s_A}(q), \sigma_\mathsf{B})(q') \cdot u \\
&= p^q_{Q^v_{>u}}(\mathsf{s_A}(q), \sigma_\mathsf{B}) + (1 - p^q_{Q^v_{>u}}(\mathsf{s_A}(q), \sigma_\mathsf{B})) \cdot u \\
&= p^q_{Q^v_{>u}}(\mathsf{s_A}(q), \sigma_\mathsf{B}) \cdot (1 - u) + u
\end{aligned}$$

Hence, $p^q_{Q^v_{>u}}(\mathsf{s_A}(q), \sigma_\mathsf{B}) \geq \frac{d}{1-u} \geq d > 0$. This holds for all Player-B GF-strategies $\sigma_\mathsf{B} \in \Sigma^q_\mathsf{B}$ and state $q \in \mathsf{Bad}$. Furthermore, recall Definition 1.28: for all $q \in \mathsf{Bad}$, letting $u := v(q)$, Player-B GF-strategies $\sigma_\mathsf{B} \in \Sigma^q_\mathsf{B}$ and Player-B strategies $\mathsf{s_B} \in \mathsf{S}^\mathcal{C}_\mathsf{B}$ such that $\mathsf{s_B}(q) = \sigma_\mathsf{B}$, we have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[Q^v_{>u}] = p^q_{Q^v_{>u}}(\mathsf{s_A}(q), \sigma_\mathsf{B}) \geq d$$

Hence, for any Player-B strategy with the Player-A strategy $\mathsf{s_A}$, if the state $q \in \mathsf{Bad}$ is seen infinitely often, almost surely, the set $Q^v_{>u}$ is seen infinitely often almost-surely. Furthermore, by Corollary 3.13, almost-surely all states visited infinitely often have the same values w.r.t. v. That is, for any Player-B strategy, almost-surely the state $q \in \mathsf{Bad}$ is seen only finitely often. Note that, alternatively, we could have invoked Corollary 3.15.

Let us now deal with the states in $\mathsf{Sec}(\mathsf{Bad})$. Let $q \in \mathsf{Sec}(\mathsf{Bad}) \setminus (T \cup v^{-1}[0])$. We have $q \in \mathsf{Sec}_{i_q}(\mathsf{Bad})$ (with $i_q \geq 1$) and $\mathsf{s_A}(q) \in \mathsf{Eff}_q(\mathsf{Sec}_{i_q-1}(\mathsf{Bad}), \mathsf{Bad})$. Therefore, $\mathsf{s_A}(q) \in \mathsf{Prog}_q(\mathsf{Sec}_{i_q-1}(\mathsf{Bad}))$. Hence,

$$y := \inf_{\sigma_\mathsf{B} \in \Sigma^q_\mathsf{B}} \max(v^q(\mathsf{s_A}(q), \sigma_\mathsf{B}), p^q_{\mathsf{Sec}_{i_q-1}(\mathsf{Bad})}(\mathsf{s_A}(q), \sigma_\mathsf{B})) > 0$$

As above, for all $\sigma_\mathsf{B} \in \Sigma^q_\mathsf{B}$, if $v^q(\mathsf{s_A}(q), \sigma_\mathsf{B}) \geq y$, then $p^q_{Q^v_{>u}}(\mathsf{s_A}(q), \sigma_\mathsf{B}) \geq y > 0$. Fix a Player-B strategy $\mathsf{s_B} \in \mathsf{S}^\mathcal{C}_q$. We let $\mathsf{IncVal}_q(y) := \{\rho \in Q^+ \mid v^q(\mathsf{s_A}(q), \mathsf{s_B}(\rho \cdot q)) \geq y\}$ and $\mathsf{Prog}_q(y) := \{\rho \in Q^+ \mid p^q_{\mathsf{Sec}_{i_q-1}(\mathsf{Bad})}(\mathsf{s_A}(q), \sigma_\mathsf{B}) \geq y\}$. As for the states in Bad, we have that if the set $\mathsf{IncVal}_q(y)$ occurs infinitely often then almost-surely the state q is seen infinitely often and almost-surely the set $Q^v_{>u}$ is seen infinitely often. Furthermore, as mentioned above for states in Bad, by Corollary 3.13, almost-surely all states visited infinitely often have the same values w.r.t. v. Hence the set $\mathsf{IncVal}_q(y)$ almost-surely is seen finitely often.

Therefore, the definition of y implies that, if q is seen infinitely often, then so is the set $\mathsf{Prog}_q(y)$ almost-surely. Furthermore, recall Definition 1.28: for all Player-B GF-strategies $\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}^q$ and Player-B strategies $\mathsf{s}_{\mathsf{B}} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$ such that $\mathsf{s}_{\mathsf{B}}(q) = \sigma_{\mathsf{B}}$, we have:

$$\mathbb{P}_{\mathcal{C},q}^{\mathsf{s}_{\mathsf{A}},\mathsf{s}_{\mathsf{B}}}[\mathsf{Sec}_{i_q-1}(\mathsf{Bad})] = p^q_{\mathsf{Sec}_{i_q-1}(\mathsf{Bad})}(\mathsf{s}_{\mathsf{A}}(q), \sigma_{\mathsf{B}})$$

Therefore, if the set $\mathsf{Prog}_q(y)$ occurs infinitely often, then almost-surely the set $\mathsf{Sec}_{i_q-1}(\mathsf{Bad})$ also occurs infinitely often. Overall, we obtain that if the state q is seen infinitely often, then the set $\mathsf{Sec}_{i_q-1}(\mathsf{Bad})$ is also seen infinitely often almost-surely. This holds for all $q \in \mathsf{Sec}(\mathsf{Bad}) \setminus (T \cup v^{-1}[0])$. Hence, it follows that if the game does not settle in Q_0^v, then almost-surely the set $\mathsf{Sec}_0(\mathsf{Bad}) = T$ is seen infinitely often. That is, Player A wins the reachability game. Hence, the second condition of Corollary 3.14 is satisfied by the strategy s_{A} w.r.t. the valuation v. We can therefore apply Corollary 3.14 to obtain that the strategy s_{A} guarantees the valuation v. □

The two definitions of progressive GF-strategies coincide

Lemma 4.15. *Consider a standard finite concurrent reachability game $\mathcal{G}$. Then, for all sets of good states $\mathsf{Gd} \subseteq Q$ that Player A wants to reach, for all $q \in Q$, a Player-A GF-strategy $\sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}}^q$ is progressive w.r.t. Gd in the sense of Definition 4.8 if and only if it is in the sense of Definition 4.14.*

Proof. Recall that, since the standard game $\mathcal{G}$ is finite, then the set of Player-B actions $\mathsf{Act}_{\mathsf{B}}^q$ at state q is finite. Furthermore, in both definitions of progressive, the GF-strategy considered is optimal in game in normal form $\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle$.

Now, consider a Player-A GF-strategy $\sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}}^q$ and assume that it is progressive w.r.t. Gd in the sense of Definition 4.8. For all Player-B actions $b \in \mathsf{Resp}_{\sigma_{\mathsf{A}}}^{\mathsf{B}}(q)$, we let $p_b := \sum_{a \in \mathsf{Sp}(\sigma_{\mathsf{A}})} \sigma_{\mathsf{A}}(a) \cdot \varrho_q(a,b)[\mathsf{Gd}] > 0$ and $p := \min_{b \in \mathsf{Resp}_{\sigma_{\mathsf{A}}}^{\mathsf{B}}(q)} p_b > 0$. Furthermore, for all Player-B actions $b \in \mathsf{Act}_{\mathsf{B}}^q \setminus \mathsf{Resp}_{\sigma_{\mathsf{A}}}^{\mathsf{B}}(q)$, we let $v_b := v^q(\sigma_{\mathsf{A}}, b) > 0$ (in the sense of Definition 4.13) and $v := \min_{b \in \mathsf{Act}_{\mathsf{B}}^q \setminus \mathsf{Resp}_{\sigma_{\mathsf{A}}}^{\mathsf{B}}(q)} v_b > 0$. We let $\delta := \min(v,p)/2$. Now, consider any Player-B GF-strategy $\sigma_{\mathsf{B}} \in \mathcal{D}(\mathsf{Act}_{\mathsf{B}}^q)$. There are two possibilities:

- Assume that $\sigma_{\mathsf{B}}[\mathsf{Resp}_{\sigma_{\mathsf{A}}}^{\mathsf{B}}(q)] \geq \frac{1}{2}$. Then, we have:

$$\begin{aligned}
p^q_{\mathsf{Gd}}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) &= \varrho_q(\mathsf{s}_{\mathsf{A}}(q), \sigma_{\mathsf{B}})[\mathsf{Gd}] \\
&\geq \sum_{b \in \mathsf{Resp}_{\sigma_{\mathsf{A}}}^{\mathsf{B}}(q)} \sigma_{\mathsf{B}}(b) \cdot \sum_{a \in \mathsf{Sp}(\sigma_{\mathsf{A}})} \sigma_{\mathsf{A}}(a) \cdot \varrho_q(a,b)[\mathsf{Gd}] \\
&\geq \sum_{b \in \mathsf{Resp}_{\sigma_{\mathsf{A}}}^{\mathsf{B}}(q)} \sigma_{\mathsf{B}}(b) \cdot p \\
&= \sigma_{\mathsf{B}}[\mathsf{Resp}_{\sigma_{\mathsf{A}}}^{\mathsf{B}}(q)] \cdot p \geq p/2 \geq \delta
\end{aligned}$$

- Assume now that $\sigma_{\mathsf{B}}[\mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)] < \frac{1}{2}$. In that case, $\sigma_{\mathsf{B}}[\mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)] \geq \frac{1}{2}$. Hence:

$$\begin{aligned}
v^q(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) &= \mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) - \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}}) \\
&= \sum_{b \in \mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} \sigma_{\mathsf{B}}(b) \cdot (\mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}}, b) - \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}})) \\
&+ \sum_{b \in \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} \sigma_{\mathsf{B}}(b) \cdot (\mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}}, b) - \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}})) \\
&\geq \sum_{b \in \mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} \sigma_{\mathsf{B}}(b) \cdot v^q(\sigma_{\mathsf{A}}, b) \geq \sum_{b \in \mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} \sigma_{\mathsf{B}}(b) \cdot v \\
&= \sigma_{\mathsf{B}}[\mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)] \cdot v \geq v/2 \geq \delta
\end{aligned}$$

That is, in any case, we have $\max(p^q_{\mathsf{Gd}}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}), v^q(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})) \geq \delta > 0$. That is, the Player-A GF-strategy is progressive w.r.t. Gd in the sense of Definition 4.8.

Assume now that the GF-strategy $\sigma_{\mathsf{A}} \in \Sigma^q_{\mathsf{A}}$ is not progressive in the sense of Definition 4.8. Consider some Player-B action $b \in \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)$ such that, for all $a \in \mathsf{Sp}(\sigma_{\mathsf{A}})$, we have $\varrho_q(a, b)[\mathsf{Gd}] = 0$. Then, we have $\mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}}, b) = \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}})$, hence $v^q(\sigma_{\mathsf{A}}, b) = 0$. Furthermore, we have $p^q_{\mathsf{Gd}}(\sigma_{\mathsf{A}}, b) = \sum_{a \in \mathsf{Sp}(\sigma_{\mathsf{A}})} \varrho_q(\sigma_{\mathsf{A}}, b)[\mathsf{Gd}] = 0$. That is, $\max(v^q(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}), p^q_{\mathsf{Gd}}(\sigma_{\mathsf{A}}, b)) = 0$. Therefore the player-A GF-strategy σ_{A} is not progressive in the sense of Definition 4.14. □

The two definitions of risky GF-strategies coincide

Lemma 4.16. *Consider a standard finite concurrent reachability game $\mathcal{G}$. Then, for all sets of bad states $\mathsf{Bd} \subseteq Q$ that Player A wants to avoid, for all $q \in Q$, a Player-A GF-strategy $\sigma_{\mathsf{A}} \in \Sigma^q_{\mathsf{A}}$ is risky w.r.t. Bd in the sense of Definition 4.9 if and only if it is in the sense of Definition 4.15.*

Proof. As for the case of progressive strategies, recall that, since the standard game $\mathcal{G}$ is finite, then the set of Player-B actions $\mathsf{Act}^q_{\mathsf{B}}$ at state q is finite. Furthermore, in both definitions of risky, the GF-strategy considered is optimal in game in normal form $\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle$.

Assume now that the GF-strategy $\sigma_{\mathsf{A}} \in \Sigma^q_{\mathsf{A}}$ is risky in the sense of Definition 4.8. Consider some Player-B action $b \in \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)$ such that there is some $a_b \in \mathsf{Sp}(\sigma_{\mathsf{A}})$ such that $\varrho_q(a_b, b)[\mathsf{Bd}] > 0$. Then, we have $\mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}}, b) = \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}})$, hence $v^q(\sigma_{\mathsf{A}}, b) = 0$. Furthermore, $p^q_{\mathsf{Bd}}(\sigma_{\mathsf{A}}, b) \geq \varrho_q(a_b, b)[\mathsf{Bd}] > 0$. That is, $\sigma_{\mathsf{B}} \in \mathsf{PosPrb}_{\mathsf{Bd}}(q, \sigma_{\mathsf{A}})$ and $\frac{v^q(\sigma_{\mathsf{A}}, b)}{p^q_{\mathsf{Bd}}(\sigma_{\mathsf{A}}, b))} = 0$. Therefore the Player-A GF-strategy σ_{A} is risky in the sense of Definition 4.14.

Assume now that the GF-strategy $\sigma_{\mathsf{A}} \in \Sigma^q_{\mathsf{A}}$ is not risky in the sense of Definition 4.8. Hence, for all Player-B actions $b \in \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)$ and for all $a \in \mathsf{Sp}(\sigma_{\mathsf{A}})$, we have $\varrho_q(a, b)[\mathsf{Bd}] = 0$. For all Player-B actions $b \in \mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)$, we let $p_b := p^q_{\mathsf{Bd}}(\sigma_{\mathsf{A}}, b) = \sum_{a \in \mathsf{Sp}(\sigma_{\mathsf{A}})} \varrho_q(a,b)[\mathsf{Bd}]$ and $v_b := v^q(\sigma_{\mathsf{A}}, b) > 0$. Let $p := \max_{b \in \mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} p_b$ and $v := \min_{b \in \mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} v_b > 0$. If $p = 0$, then for all Player-B GF-strategies $\sigma_{\mathsf{B}} \in \mathcal{D}(\mathsf{Act}^q_{\mathsf{B}})$, we have $p^q_{\mathsf{Bd}}(\sigma_{\mathsf{A}}, b) = 0$, hence $\mathsf{PosPrb}_{\mathsf{Bd}}(q, \sigma_{\mathsf{A}}) = \emptyset$ and therefore σ_{A} is not risky in the sense of Definition 4.15. Assume now that $p > 0$. We let $\delta := \frac{v}{p}$. Consider any Player-B GF-strategy $\sigma_{\mathsf{B}} \in \mathcal{D}(\mathsf{Act}^q_{\mathsf{B}})$. We have:

$$\begin{aligned}
p^q_{\mathsf{Bd}}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) = & \sum_{b \in \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} \sigma_{\mathsf{B}}(b) \cdot \sum_{a \in \mathsf{Sp}(\sigma_{\mathsf{A}})} \varrho_q(a,b)[\mathsf{Bd}] \\
& + \sum_{b \in \mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} \sigma_{\mathsf{B}}(b) \cdot \sum_{a \in \mathsf{Sp}(\sigma_{\mathsf{A}})} \varrho_q(a,b)[\mathsf{Bd}] \\
& = \sum_{b \in \mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} \sigma_{\mathsf{B}}(b) \cdot \sum_{a \in \mathsf{Sp}(\sigma_{\mathsf{A}})} \varrho_q(a,b)[\mathsf{Bd}] \\
& \leq \sum_{b \in \mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} \sigma_{\mathsf{B}}(b) \cdot p = \sigma_{\mathsf{B}}[\mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)] \cdot p
\end{aligned}$$

Furthermore:

$$\begin{aligned}
v^q(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) & = \mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) - \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}}) \\
& = \sum_{b \in \mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}} \sigma_{\mathsf{B}}(b) \cdot (\mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}}, b) - \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}})) \\
& + \sum_{b \in \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} \sigma_{\mathsf{B}}(b) \cdot (\mathsf{out}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}}, b) - \mathsf{val}[\langle \mathsf{F}(q), \mathsf{m}_{\mathcal{G}} \rangle](\sigma_{\mathsf{A}})) \\
& = \sum_{b \in \mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} \sigma_{\mathsf{B}}(b) \cdot v^q(\sigma_{\mathsf{A}}, b) \geq \sum_{b \in \mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)} \sigma_{\mathsf{B}}(b) \cdot v \\
& = \sigma_{\mathsf{B}}[\mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)] \cdot v
\end{aligned}$$

Hence, whenever $\sigma_{\mathsf{B}} \in \mathsf{PosPrb}_{\mathsf{Bd}}(q, \sigma_{\mathsf{A}})$, we have $\sigma_{\mathsf{B}}[\mathsf{Act}^q_{\mathsf{B}} \setminus \mathsf{Resp}^{\mathsf{B}}_{\sigma_{\mathsf{A}}}(q)] > 0$ and therefore $\frac{v^q(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})}{p^q_{\mathsf{Bd}}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})} \geq \frac{v}{p} > 0$. That is, the GF-strategy σ_{A} is not risky. $\square$

5 - Büchi, co-Büchi and parity objectives

In this chapter, we study finite-state concurrent games with arbitrary local interactions, with Büchi, co-Büchi and parity objectives. Contrary to the previous chapter, we now consider games with stopping states. Note that in this chapter, whenever we exhibit a game to witness a negative result, we will invoke Corollary 3.38. Indeed, thanks to this result, it suffices to show that positional strategies are not enough to achieve a value to show that finite choice strategies, and in particular finite-memory strategies, cannot achieve it either.

We first focus on Büchi games. The positive result we show is that, as for reachability games, whenever there is an optimal strategy, there is one that is positional (which is therefore also subgame optimal). This is a generalization of what we did in [40] with standard finite local interactions. We also exhibit a standard finite concurrent game where playing almost-optimally requires infinite choice. Note that this game was already used in the litterature to show that infinite memory is needed to play almost-optimally. We can then complete the picture of how arbitrary concurrent Büchi games behave, see Theorem 5.5.

We then consider co-Büchi games. The main positive result for this objective is that, with standard finite game forms, positional strategies are enough to be almost-optimal. This was proved in [50, Theorem 3.1]. It is an open question if this still holds in games with arbitrary game forms which are maximizable w.r.t. Player A. The other positive result is that playing subgame optimally in standard finite games can be done positionally. This is a direct consequence of Corollary 3.23. As for the negative results, we show that infinite choice is required for playing almost-optimally with arbitrary game forms — it is a direct corollary of the fact that this is already the case for safety objectives. We also show that playing optimally in standard finite games requires infinite choice. This is exemplified by a co-Büchi game we have already discussed twice in this dissertation. We also exhibit an arbitrary concurrent game, with local interactions maximizable w.r.t. Player A, where playing subgame optimally requires infinite choice. Overall, all the results summarizing how arbitrary concurrent co-Büchi games behave are gathered in Theorem 5.13.

Finally, we consider parity objectives with at least 3 colors. These objectives inherit all the negative results of the objectives studied before in this chapter. In fact, it only remains to exhibit a standard finite game where playing subgame optimally requires infinite choice. This example is already known, and we have already briefly discussed it in Subsection 3.4.2. The results are gathered in Theorem 5.15.

Seeing finitely and infinitely often a set of states. Before diving into how each objective behaves, we recall how to write with union and intersection

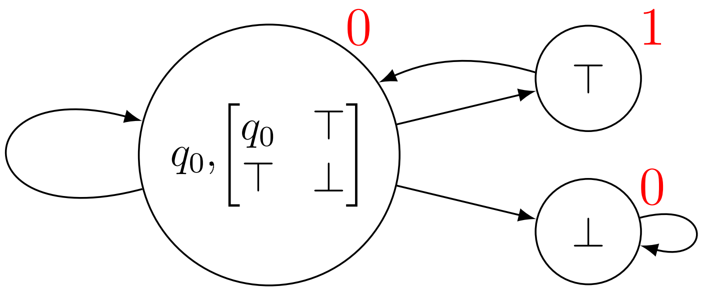

Figure 5.1: A deterministic standard concurrent Büchi game $\mathcal{G} = \langle \mathcal{C}, \mathsf{Buchi} \rangle$ where Player A wants to visit the state $\top$ infinitely often.

the event that a set of states is seen (in)finitely often. This will be particularly useful in this section when considering the probability of such events.

Given a set of states Q and subset of states $T \subseteq Q$, the event where the set of states T is seen infinitely often can be written as follows:

$$(Q^* \cdot T)^\omega = \bigcap_{n \in \mathbb{N}} (Q^* \cdot T)^n$$

Symmetrically, the event where the set of states T is seen only finitely often can be written as follows:

$$\bigcup_{n \in \mathbb{N}} Q^n \cdot (Q \setminus T)^\omega$$

5.1 Büchi objectives

Let us first deal with Büchi objectives. Recall that we only consider games with finitely many states. Since there does not always exist optimal strategies in reachability games — even when all local interactions are standard finite — it is also the case for Büchi games. Furthermore, we have shown in the previous chapter that almost-optimal strategies can always be found among positional strategies in reachability games, without any assumptions on the local interactions. This does not hold in Büchi games, since in general infinite-choice strategies may be required to be almost-optimal. We provide a Büchi game witnessing this fact below in Figure 5.1 and Definition 5.1. Note that this Büchi game is very close to the snow ball reachability game of Figure 3.1. The only difference is that the target $\top$ that Player A wants to visit infinitely often loops back on q_0, instead of self-looping. In addition, note that this example is already known and comes from [47, Figure 1].

Definition 5.1 (Game described in Figure 5.1). *Consider the game depicted in Figure 5.1. This game $\mathcal{G} = \langle \mathcal{C}, \mathsf{Buchi} \rangle$ is standard and deterministic. There is only one non-trivial state: q_0. The set of colors considered is $\mathsf{K} := \{0, 1\}$*

and the colors of the states $q_0, \top, \bot$ are given in red near them: $\mathsf{col}(q_0) := 0$, $\mathsf{col}(\bot) := 0$ *and* $\mathsf{col}(\top) := 1$. *This game is win/lose with a Büchi objective (recall Definition 1.25): Player* A *wins if and only if the state* $\top$ *is visited infinite often. The Player-*A *set of actions at state* q_0 *is* $\mathsf{Act}_{\mathsf{A}}^{q_0} := \{a_1, a_2\}$ *where* a_1 *refers to the top row and* a_2 *refers to the bottom row and similarly we have* $\mathsf{Act}_{\mathsf{B}}^{q_0} := \{b_1, b_2\}$ *where* b_1 *refers to the leftmost column and* b_2 *refers to the rightmost column.*

Proposition 5.1. *The standard finite concurrent Büchi game* $\mathcal{G}$ *from Definition 5.1 is such that:*

- *The value of the game from* q_0 *is 1;*
- *Player* A *has no optimal strategy;*
- *All finite-choice Player-*A *strategies have value 0 from* q_0.

Proof. Let us first show the third item. Consider any positional Player-A strategy $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$. Let $p := \mathsf{s_A}(q_0)(a_1) \in [0, 1]$ be the probability that the strategy $\mathsf{s_A}$ plays the top row in q_0. If $p = 1$, then a Player-B deterministic positional strategy that plays action b_1 in q_0 (the left column) ensures that the game never leaves the state q_0, and therefore never reaches the state $\top$. Hence, such a Player-A strategy $\mathsf{s_A}$ has value 0. Assume now that $p < 1$. Then, a Player-B deterministic positional strategy that plays action b_2 in q_0 (the right column) ensures that, at each step, there is probability $1 - p > 0$ to reach the sink state $\bot$ of value 0. Otherwise, the state $\top$ is reached, and the game loops back to q_0 and has once again probability $1 - p > 0$ to reach the sink state $\bot$. Hence, almost surely, with both of these strategies, the sink state $\bot$ is reached. Therefore, this Player-A strategy $\mathsf{s_A}$ has value 0 from q_0. In fact, all Player-A positional strategies have value 0 from q_0. Therefore, by Corollary 3.38, all Player-A finite-choice strategies have value 0 from q_0.

Consider now the first item. Let $\varepsilon > 0$. Let us define a Player-A strategy $\mathsf{s_A}$ of value at least $1 - \varepsilon$. Consider a sequence $(\varepsilon_k)_{k \in \mathbb{N}}$ such that, for all $k \in \mathbb{N}$, we have $\varepsilon_k > 0$ and $\sum_{k \in \mathbb{N}} \varepsilon_k \leq \varepsilon$. Furthermore, for all $\rho \in \{q_0, \top, \bot\}^+$, we let $|\rho|_\top \in \mathbb{N}$ denote the number of times the state $\top$ occurred in ρ. We now define $\mathsf{s_A}$, as follows. For all $\rho \in \{q_0, \top, \bot\}^+$ such that $\rho_{\mathsf{lt}} = q_0$, we let

$$\mathsf{s_A}(\rho) := \{a_1 \mapsto 1 - \varepsilon_{|\rho|_\top},\ a_2 \mapsto \varepsilon_{|\rho|_\top}\}$$

Clearly, this strategy has infinite choice. Consider any Player-B strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$. We have:

$$\begin{aligned}\mathbb{P}_{\mathcal{C},q_0}^{\mathsf{s_A},\mathsf{s_B}}[(q_0 \cup \top)^* \cdot \bot] &= \sum_{k \in \mathbb{N}} \mathbb{P}_{\mathcal{C},q_0}^{\mathsf{s_A},\mathsf{s_B}}[(q_0 \cup \top)^k \cdot \bot] \leq \sum_{k \in \mathbb{N}} \sum_{\rho \in (q_0 \cup \top)^k} \mathbb{P}_{\mathcal{C},q_0}^{\mathsf{s_A},\mathsf{s_B}}(\rho) \cdot \mathsf{s_A}(\rho)(a_2) \\ &\leq \sum_{k \in \mathbb{N}} \sum_{\rho \in (q_0 \cup \top)^k} \mathbb{P}_{\mathcal{C},q_0}^{\mathsf{s_A},\mathsf{s_B}}(\rho) \cdot \varepsilon_k \leq \sum_{k \in \mathbb{N}} \varepsilon_k \leq \varepsilon\end{aligned}$$

Furthermore, we have — by Proposition 1.3 (the continuity of probability function) for the third equality:

$$\begin{aligned}
\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[(q_0 \cup \top)^* \cdot q_0^\omega] &= \sum_{k \in \mathbb{N}} \sum_{\rho \in (q_0 \cup \top)^*,\ |\rho|_\top = k} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}(\rho) \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0 \cdot \rho}(q_0^\omega) \\
&= \sum_{k \in \mathbb{N}} \sum_{\rho \in (q_0 \cup \top)^*,\ |\rho|_\top = k} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}(\rho) \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0 \cdot \rho}(\bigcap_{n \in \mathbb{N}} q_0^n) \\
&= \sum_{k \in \mathbb{N}} \sum_{\rho \in (q_0 \cup \top)^*,\ |\rho|_\top = k} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}(\rho) \cdot (\lim_{n \to \infty} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0 \cdot \rho}(q_0^n)) \\
&= \sum_{k \in \mathbb{N}} \sum_{\rho \in (q_0 \cup \top)^*,\ |\rho|_\top = k} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}(\rho) \cdot (\lim_{n \to \infty} (1 - \varepsilon_k)^n) = 0
\end{aligned}$$

Overall, we obtain that:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[(q_0^* \cup \top)^\omega] = 1 - (\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[(q_0 \cup \top)^* \cdot \bot] + \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q_0}[(q_0 \cup \top)^* \cdot q_0^\omega]) \geq 1 - \varepsilon$$

As this holds for all Player-B strategies $\mathsf{s_B}$, it follows that the strategy $\mathsf{s_A}$ has value at least $1 - \varepsilon$ from q_0. In fact, the value of the game from q_0 is 1.

As for the second item, since in the snow-ball game of Figure 3.1 Player-A has no optimal strategy (i.e. strategy of value 1) from q_0, then it is also the case for this game. □

Hence, playing almost-optimally can be very costly in Büchi games, whereas it is not the case in reachability games. However, whenever it is possible to play optimally in Büchi games, just like in reachability games, it can be done with a positional strategy.

Proposition 5.2. *In all arbitrary finite-state concurrent Büchi games, when there is an optimal strategy, there is one that is positional. (This strategy is therefore subgame optimal.)*

To prove this result, we will transform a Büchi game into a reachability game, and use the result already existing on reachability games. We define below how to translate a Büchi game into a reachability game. Informally, this is done by replacing every state q in the target by a trivial state with only one possible outcome: a probability distribution that goes with probability $\chi_\mathcal{G}[\mathsf{A}](q)$ to the new target (in the reachability game) and with probability $1 - \chi_\mathcal{G}[\mathsf{A}](q)$ to a sink state of value 0.

Definition 5.2. *Consider an arbitrary finite-state Büchi game* $\mathcal{G} = \langle \mathcal{C}, \mathsf{Buchi} \rangle$. *We let* $T := \mathsf{col}^{-1}[\{1\}] \subseteq Q$ *and* $Q_{\mathsf{ch}} := T \cup Q_{\mathsf{ns}}$. *We define the reachability game* $\mathcal{G}^{\mathsf{Reach}} := \langle \mathcal{C}^{\mathsf{Reach}}, \mathsf{Reach} \rangle$ *with* $\mathcal{C}^{\mathsf{Reach}} := \langle Q \uplus \{\top\} \uplus \{\bot\}, \mathsf{F}^{\mathsf{Reach}}, \mathsf{K}, \mathsf{col}^{\mathsf{Reach}} \rangle$ *as follows:*

- $\{\top\}$ *and* $\{\bot\}$ *are two fresh states not in* Q*;*

- *all states* $q \in Q_{\mathsf{ch}}$ *are made into trivial states with only one outcome:* $\mathsf{F}^{\mathsf{Reach}}(q) := \langle *, *, \{\top \mapsto \chi_{\mathcal{G}}[\mathsf{A}](q),\ \bot \mapsto 1 - \chi_{\mathcal{G}}[\mathsf{A}](q)\}, *\rangle$*;*

- *all states* $q \in Q \setminus Q_{\mathsf{ch}}$ *are left unchanged:* $\mathsf{F}(q) := \mathsf{F}^{\mathsf{Reach}}(q)$*;*

- *both states* $\top$ *and* $\bot$ *are self-looping sinks;*

- $\mathsf{col}^{\mathsf{Reach}}(\top) := 1$ *and* $\mathsf{col}^{\mathsf{Reach}}[Q \cup \bot] := \{0\}$*;*

Remark 5.1. *From the definition of the arena* $\mathcal{C}^{\mathsf{Reach}}$*, one can realize that there are no stopping states in* $\mathcal{C}^{\mathsf{Reach}}$*. This is done so that we can use the results from Chapter 4 that only apply to games without stopping states.*

Furthermore, the game $\mathcal{G}^{\mathsf{Reach}}$ *would be identical if we changed it into Büchi game with the same target, since this target is a self-looping sink.*

This transformation ensures the lemma below.

Lemma 5.3. *Consider an arbitrary finite-state Büchi game* $\mathcal{G} = \langle \mathcal{C}, \mathsf{Buchi} \rangle$ *and the reachability game* $\mathcal{G}^{\mathsf{Reach}}$*. Then, for all* $q \in Q$*, we have* $\chi_{\mathcal{G}}[\mathsf{A}](q) = \chi_{\mathcal{G}^{\mathsf{Reach}}}[\mathsf{A}](q)$*. Furthermore, if Player-*A *has an optimal strategy in* $\mathcal{G}$*, then she also has one in* $\mathcal{G}^{\mathsf{Reach}}$*.*

Proof. We want to apply Lemma 1.16, Page 48. However, our transformation does not fit exactly the statement of Lemma 1.16. First, we have replaced the states in Q_{ch} with trivial states instead of stopping states. However, it is rather straightforward that if we had replaced every state $q \in Q_{\mathsf{ch}}$ by a stopping states of value $\chi_{\mathcal{G}}[\mathsf{A}](q)$, all the values of the states in the game $\mathcal{G}^{\mathsf{Reach}}$ would stay the same. Second, we changed the objective from Büchi to reachability. However, since the games $\mathcal{G}$ and $\mathcal{G}^{\mathsf{Reach}}$ can be seen as over once a state in Q_{ch} is reached, what matters is what happens if no state in Q_{ch} is reached. In both games $\mathcal{G}$ and $\mathcal{G}^{\mathsf{Reach}}$, what happens in that case is the same: Player A loses. Hence, we can apply Lemma 1.16 to obtain the desired statement. □

Furthermore, we have also the following lemma.

Lemma 5.4. *Consider an arbitrary finite-state Büchi game* $\mathcal{G} = \langle \mathcal{C}, \mathsf{Buchi} \rangle$ *and the reachability game* $\mathcal{G}^{\mathsf{Reach}}$*. Assume that there is an optimal strategy in* $\mathcal{G}^{\mathsf{Reach}}$ *and, for all* $q \in T \setminus Q_{\mathsf{ns}}$*, Player* A *has* GF*-strategies in* Σ_{A}^{q} *that are optimal in the game in normal form* $\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle$*. Then, Player-*$\mathsf{A}$ *has a positional optimal strategy in* $\mathcal{G}$*.*

Proof. By Theorem 4.11, since Player A has an optimal strategy in the reachability game $\mathcal{G}^{\mathsf{Reach}}$, then she has one $\mathsf{s}_{\mathsf{A}}^{\mathsf{Reach}}$ that is positional. Let us now define a Player-A positional strategy $\mathsf{s}_{\mathsf{A}} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ that is optimal. For all $q \in Q \setminus Q_{\mathsf{ch}}$, we let $\mathsf{s}_{\mathsf{A}}(q) := \mathsf{s}_{\mathsf{A}}^{\mathsf{Reach}}(q) \in \Sigma_{\mathsf{A}}^{q}$. Furthermore, for all $q \in T \setminus Q_{\mathsf{ns}}$, we let $\mathsf{s}_{\mathsf{A}}(q)$ be a Player-A GF-strategy that is optimal in the game in normal form

$\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle$. Let us show that the positional Player-A strategy $\mathsf{s_A}$ is (subgame) optimal in $\mathcal{G}$ by applying Corollary 3.14[1]. Consider the first condition of this corollary: that is, let us show that the strategy $\mathsf{s_A}$ is locally optimal. By Lemma 5.3 and since the Player-A strategy $\mathsf{s_A^{Reach}}$ is optimal in $\mathcal{G}^{\mathsf{Reach}}$, for all $q \in Q \setminus Q_{\mathsf{ch}}$, we have $\chi_{\mathcal{G}}[\mathsf{A}](q) \leq \mathsf{val}[\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle](\mathsf{s_A}(q))$. This also holds for all $q \in T \setminus Q_{\mathsf{ns}}$ by definition of the strategy $\mathsf{s_A}$ and by Lemma 3.9. Hence, the Player-A strategy $\mathsf{s_A}$ is locally optimal in $\mathcal{G}$. Consider now the second condition of Corollary 3.14. We let $V := \chi_{\mathcal{G}}[\mathsf{A}][Q] \subseteq [0, 1]$ and for all $u \in V$, we let $Q_u := \{q \in Q \mid \chi_{\mathcal{G}}[\mathsf{A}](q) = u\}$. When applied to a Büchi game, the second condition of Corollary 3.14 states that, against all Player-B strategies, if the game loops ever indefinitely on states of the same positive value, then the target is seen infinitely often almost surely. Said otherwise, the probability of, at some point, always seeing states of the same positive value while avoiding the target is 0. Since the Player-A strategy $\mathsf{s_A}$ is positional, this amounts to show that, from every state $q \in Q \setminus Q_{\mathsf{ch}}$ of value $u > 0$, the probability to always see Q_u and never Q_{ch} is 0. Formally, we show that, for all Player-B strategies $\mathsf{s_B} \in \mathsf{S_B^{\mathcal{C}}}$, for all $u \in V \setminus \{0\}$ and $q \in Q_u$, we have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[(Q_u \setminus Q_{\mathsf{ch}})^\omega] = 0 \tag{5.1}$$

This actually comes straightforwardly from Corollary 3.14 applied to the (subgame) optimal strategy $\mathsf{s_A^{Reach}}$ in $\mathcal{G}^{\mathsf{Reach}}$, because the strategies $\mathsf{s_A}$ and $\mathsf{s_A^{Reach}}$ coincide on $Q \setminus Q_{\mathsf{ch}}$. Note that Corollary 3.14 only applies to games with PI payoff functions, however, as mentioned in Remark 5.1, the game $\mathcal{G}^{\mathsf{Reach}}$ can be seen as a Büchi game since the target is a self-looping sink. Let $u \in V \setminus \{0\}$, $q \in Q_u$ and $\mathsf{s_B} \in \mathsf{S_B^{\mathcal{C}}}$. The Player-$\mathsf{B}$ strategy $\mathsf{s_B}$ can be seen as a strategy in the arena $\mathsf{S_B^{\mathcal{C}^{Reach}}}$. We have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[(Q_u \setminus Q_{\mathsf{ch}})^\omega] = \mathbb{P}^{\mathsf{s_A^{Reach}},\mathsf{s_B}}_{\mathcal{C}^{\mathsf{Reach}},q}[(Q_u \setminus Q_{\mathsf{ch}})^\omega]$$

Furthermore, if $\mathbb{P}^{\mathsf{s_A^{Reach}},\mathsf{s_B}}_{\mathcal{C}^{\mathsf{Reach}},q}[(Q_u \setminus Q_{\mathsf{ch}})^\omega] > 0$, then the second condition of Corollary 3.14 would not hold for the strategy $\mathsf{s_A^{Reach}}$ in the game $\mathcal{G}^{\mathsf{Reach}}$, since $\mathbb{P}^{\mathsf{s_A^{Reach}},\mathsf{s_B}}_{\mathcal{C}^{\mathsf{Reach}},q}[(Q_u \setminus Q_{\mathsf{ch}})^\omega \cap Q^* \cdot T] = 0$. In fact, Equation 5.1 does hold, and by Corollary 3.14 the Player-A positional $\mathsf{s_A}$ is (subgame) optimal in $\mathcal{G}$. $\square$

We can now proceed to the proof of Proposition 5.2.

Proof. By Lemma 5.3, Player A has an optimal strategy in the game $\mathcal{G}^{\mathsf{Reach}}$. Furthermore, consider the Player-A optimal strategy $\mathsf{t_A}$ in $\mathcal{G}$. For all states $q \in T \setminus Q_{\mathsf{ns}}$, we have $\mathsf{t_A}(q) \in \Sigma_{\mathsf{A}}^q$. By Lemma 3.10, for all $q \in T \setminus Q_{\mathsf{ns}}$, we have $\chi_{\mathcal{G}}[\mathsf{A}](q) = \chi_{\mathcal{G}}[\mathsf{t_A}](q) \leq \mathsf{val}[\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}]^q \rangle](\mathsf{t_A}(q)) = \mathsf{val}[\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle](\mathsf{t_A}(q)) \leq \mathsf{val}[\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle][\mathsf{A}] = \chi_{\mathcal{G}}[\mathsf{A}](q)$. This last equality comes from Lemma 3.9.

[1]Note that we cannot apply Corollary 3.16 since the game $\mathcal{G}$ is not standard.

	∃ Opt. ?	ε-Opt.	Optimal	SubG. Opt.
Buchi	No	∞-choice	Positional	Positional

Table 5.1: The summary of the situation in arbitrary finite-state concurrent Büchi games.

Hence, the Player-A GF-strategy $\mathsf{t}_\mathsf{A}(q)$ is optimal in the game in normal form $\langle \mathsf{F}(q), \chi_\mathcal{G}[\mathsf{A}] \rangle$. Hence by Lemma 5.4, Player A has an optimal positional strategy in $\mathcal{G}$. □

We summarize the results on Büchi games in the theorem below.

Theorem 5.5. *In arbitrary finite-state concurrent Büchi games:*

- *there does not always exist optimal strategies;*
- *almost-optimal strategies may require infinite choice, which can be witnessed by a standard finite game;*
- *whenever there exists an optimal strategy, there is one that is positional. This also holds for subgame optimal strategies.*

These results can be seen in Table 5.1.

Proof.

- This is consequence of the fact that this is already the case for reachability games, see Theorem 4.12.
- It was already known that infinite memory may be required to play almost-optimally in Büchi games (see [32, Theorem 2]). We have shown this result in Proposition 5.1, by reusing an example already known (that comes from [47, Figure 1]).
- This is given by Proposition 5.2.

□

5.2 co-Büchi objectives

Let us now consider the case of co-Büchi objectives. Recall that we only consider games with finitely many states. Since there does not always exist optimal strategies in standard finite reachability games, it is also the case for co-Büchi objectives. Let us now deal with how to play almost-optimally, optimally and subgame optimally in co-Büchi games.

5.2.1 . Almost-optimal strategies

Consider first how to play almost-optimally in co-Büchi games. As stated in Proposition 4.4, if no assumption is made on the local interactions (i.e. if we do not assume that they are maximizable w.r.t. Player A), in safety games, infinite-choice strategies may be required to play almost-optimally. Furthermore, note that safety games can be seen as special cases of co-Büchi games where the target is self-looping. Hence, it is also the case of co-Büchi games that playing almost-optimally may require infinite choice, if no assumption is made on the local interactions.

However, in co-Büchi games where all local interactions are standard finite, then playing almost-optimally can be done positionally.

Theorem 5.6. *In all standard finite concurrent co-Büchi games, for all $\varepsilon > 0$, Player A has a positional strategy that is ε-optimal.*

As mentioned at the end of Subsection 3.2.3, this result is already known, it was proved in [50, Theorem 3.1]. The proof is quite involved as it is byproduct a memory transfer result from limit-sure winning (i.e. almost-optimal for the value 1) to almost-optimal strategies.

However, what happens in arbitrary games with local interactions maximizable w.r.t. Player A is unknown.

Open Question 5.1. *Can playing almost-optimally be done positionally in arbitrary finite-state co-Büchi games where all local interactions are maximizable w.r.t. Player A?*

5.2.2 . Optimal strategies

Let us now consider how to play optimally in co-Büchi games. Contrary to the Büchi case where, whenever it can be done, it can be done positionally regardless of the local interactions, playing optimally may require infinite choice in co-Büchi games even in standard finite games. This is witnessed by the co-Büchi game described in Figure 5.2. In fact, we have already discussed this game twice in this dissertation. Once, in Figure 2.5 as a slightly modified game, to witness that it is possible that action strategies[2] may achieve a value that regular strategy can only approach. The second time was in Figure 3.2 to witness that there can be optimal strategies without subgame optimal strategies. For a formal description, one can consider Definition 2.19 that referred to the game of Figure 2.5. The difference with the game of Figure 5.2 is that, in Figure 2.5, both states q_1 and q_1' are merged into a single state q_1. Note

[2]Action strategy is a notion defined in Section 2.5, in the context of standard games, and used only in that section. These are strategies that take into account the states and actions played by the strategies. To distinguish the strategy we usually use from the action strategies of this section, we will sometimes refer to them as regular strategies.

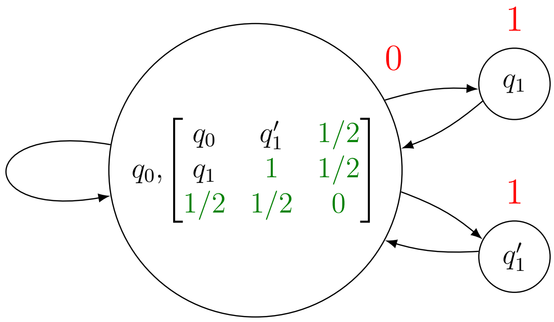

Figure 5.2: A co-Büchi game where playing optimally requires infinite choice.

that the most complicated properties ensured by this game have already been proven in Proposition 2.21 and the subsequent lemmas stated to prove it.

Proposition 5.7. *The game of Figure 5.2 ensures the following:*

- *the state q_0 has value $\frac{1}{2}$;*
- *no finite-choice Player-A strategy is optimal;*
- *there is an infinite-choice Player-A strategy that is optimal.*

Proof.

- This is given by Lemma 2.22.
- Consider any positional player-A strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$. If $\mathsf{s_A}(q_0)(a_3) > 0$ (i.e. the bottom row is played with positive probability), then Player B can play the rightmost column with probability 1 and ensure that the expected value of the stopping states seen is less than $\frac{1}{2}$. Therefore, such a strategy has value less than $\frac{1}{2}$. Otherwise, if $\mathsf{s_A}(q_0)(a_1) = 1$, then Player B can play the middle column with probability 1 and ensure that the state q_1' will be seen infinitely often. Therefore, such a strategy has value 0. Finally, if $\mathsf{s_A}(q_0)(a_3) = 0$ and $\mathsf{s_A}(q_0)(a_2) > 0$, then Player B can play the leftmost column with probability 1 and ensure that, almost-surely, the state q_1 will be seen infinitely often. Hence, such a strategy has also value 0. In fact, no positional strategy can achieve the value $\frac{1}{2}$ in this game. Thus, by Corollary 3.38, all Player-A finite-choice strategies are not optimal from q_0.
- An optimal action strategy is described in Lemma 2.23. With the difference between the games of Figure 2.5 and of Figure 5.2, this can be done with a regular infinite-choice strategy.

□

5.2.3 . Subgame optimal strategies

What happens for subgame optimally depends on the local interactions in the game. First, let us consider the case of standard finite local interactions. In this context, subgame optimal strategies can be found among positional strategies.

Proposition 5.8. *In all standard finite concurrent co-Büchi games, if there is a subgame optimal strategy, there is one that is positional.*

Proof. Let us apply Corollary 3.24. To do so, we need to show that in all standard finite co-Büchi games, if there is a subgame almost-surely winning strategy, there is one that is positional. Hence, consider such standard finite co-Büchi game $\mathcal{G}$ and assume that Player A has a subgame almost-surely winning strategy. In that case, all states in $\mathcal{G}$ have value 1. By Theorem 5.6, there is a positional strategy $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ that is $\frac{1}{2}$-optimal (from all states). Consider any Player-B strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$. Then, from all states, there is probability at least $\frac{1}{2}$ that the PI objective $\mathsf{coBuchi}$ holds. Hence, by Proposition 3.5[3], the probability of the PI objective $\mathsf{coBuchi}$ is 1 from all states. That is, the positional strategy $\mathsf{s_A}$ is subgame almost-surely winning. We can therefore apply Corollary 3.24 to obtain the result. □

However, if the local interactions are not maximizable w.r.t. Player A, playing subgame optimally may require infinite choice. This is witnessed in the game of Figure 5.3. Note that it is very close to the game of Figure 4.2 that witnessed that playing almost-optimally may require infinite choice in safety games. The only change is that the state $\perp$ that Player A wants to see only finitely often loops back to q_0.

We described formally this game below.

Definition 5.3 (Game depicted in Figure 5.3). *The game of Figure 5.3 is an MDP Γ where Player A plays alone with two states: $Q := \{q_0, \perp\}$. The state $\perp$ is trivial and loops back to q_0 and, at state q_0, Player A may play any integer $n \in \mathbb{N}$ which leads to a distribution $d_n := \{q_0 \mapsto 1 - \frac{1}{2^n}; \perp \mapsto \frac{1}{2^n}\} \in \mathcal{D}(Q)$. Player A has a co-Büchi objective $\mathsf{coBuchi}$ with $\mathsf{K} = \{0, 1\}$ and $\mathsf{col}(q_0) := 0$ and $\mathsf{col}(\perp) := 1$, i.e. Player A wants to see the state $\perp$ only finitely often.*

Proposition 5.9. *In the co-Büchi game $\mathcal{G}$ of Definition 5.3:*

- *the state q_0 has value 1;*
- *all finite-choice Player-A strategies have value 0;*

[3]This proposition applies to stochastic trees for a prefix-independent Borel objective. However, in the game $\mathcal{G}$, there could be some stopping states (of value 1), which make the objective in the stochastic tree not PI. However, it suffices to replace these stopping states by a self-looping sink of color 0, which will therefore be of value 1. In that way, the value of the game in unchanged and the objective is PI.

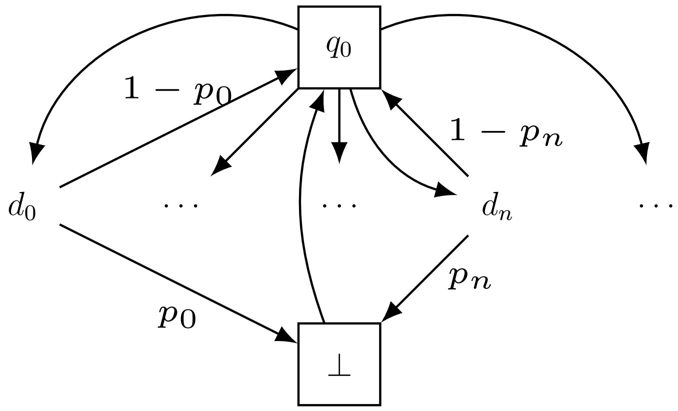

Figure 5.3: An MDP where Player A plays alone and wants to see the state $\bot$ only finitely often, with $p_i := \frac{1}{2^i}$.

- *there is an infinite-choice Player-*A *strategy that is subgame almost-surely winning.*

Proof. This is quite similar to the proof of Proposition 4.4. First, consider any positional Player-A strategy $\mathsf{s_A}$. Consider some $n \in \mathbb{N}$ such that $\mathsf{s_A}(q_0)(n) > 0$. Then, at each step, there is probability at least $\frac{\mathsf{s_A}(q_0)(n)}{2^n}$ to reach the target $\bot$, and in any case the game loops back on q_0. Hence, almost-surely, the state $\bot$ is reached infinitely often. In fact, all Player-A positional strategies have value 0. Therefore, since this MDP is standard and B-finite (recall, finitely many states, and Player B has finitely many actions), by Corollary 3.38, all finite-choice strategies have value 0 from q_0.

Let us now build a Player-A infinite-choice strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ that is subgame almost-surely winning. For all $\rho \in Q^*$, we let $\mathsf{s_A}(\rho) := \{|\rho| + 1 \mapsto 1\}$. Let $\pi \in Q^+$. Denoting $\mathsf{s_B}$ the only Player-B strategy in Γ, for all $n \in \mathbb{N}$, we have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\Gamma,\pi}(Q^n \cdot (Q^* \cdot \bot)) = \sum_{k \in \mathbb{N}} \sum_{\rho \in \pi \cdot Q^n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\Gamma,\rho}(q_0^k \cdot \bot) = \sum_{k \in \mathbb{N}} \frac{1}{2^{|\rho|+k}} \leq \sum_{k \in \mathbb{N}} \frac{1}{2^{n+k+1}} = \frac{1}{2^n}$$

Hence, by Proposition 1.3 (the continuity of probability measure), we have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\Gamma,\pi}(\bigcap_{n \in \mathbb{N}} Q^n \cdot (Q^* \cdot \bot)) = \lim_{n \to \infty} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\Gamma,\pi}(Q^n \cdot (Q^* \cdot \bot)) = 0$$

That is, from π, the state $\bot$ is seen infinitely often with probability 0. In other words, it is seen finitely often with probability 1. This holds for all $\pi \in Q^+$. Therefore, the infinite-choice strategy $\mathsf{s_A}$, is subgame almost-surely winning. □

The fact that playing subgame optimally may require infinite choice with co-Büchi objectives can be witnessed in games where all local interactions are

$$\mathcal{F} = \begin{bmatrix} [q_0^- - q_1] & \frac{\top + q_1}{2} \\ [\top^- - \bot] & \frac{\top + q_1}{2} \end{bmatrix}$$

Figure 5.4: A standard infinite local interaction maximizable w.r.t. Player A.

maximizable w.r.t. Player A. However, the local interactions considered will not be finite, otherwise we would be in the scope of Proposition 5.8. We conclude this subsection by providing an example of such a co-Büchi game. Let us first define the only non-trivial game form occurring in that game. To do so, let us first introduce a notation for game form.

Definition 5.4 (Notation for a standard game form). *Given any two outcome x and y, the notation $[x^- - y]$ refers to the standard game form $\langle \mathbb{N}, *, \{x, y\}, \varrho \rangle_{\mathsf{s}}$ where Player A plays alone and, for all $n \in \mathbb{N}$, we have $\varrho(n, *) := \{x \mapsto 1 - \frac{1}{2^n}, y \mapsto \frac{1}{2^n}\}$.*

The game form of interest for us is depicted in Figure 5.4 and formally defined in Definition 5.5 below.

Definition 5.5 (Game form depicted in Figure 5.4). *The game form $\mathcal{F} = \langle \mathsf{Act}_\mathsf{A}, \mathsf{Act}_\mathsf{B}, \{q_0, q_1, \top, \bot\}, \varrho \rangle_{\mathsf{s}}$ of Figure 5.4 is standard, Player A has infinitely many actions available, whereas Player B has two. We have $\mathsf{Act}_\mathsf{A} := \{a_1, a_2\} \times \mathbb{N}$ and $\mathsf{Act}_\mathsf{B} := \{b_1, b_2\}$. The two Player-$\mathsf{B}$ actions correspond to the two columns of the game form, with b_1 corresponding to the leftmost column and b_2 corresponding to the rightmost one. If Player B plays b_2, then the outcome is a uniform distribution on $\top$ and q_1. Otherwise, i.e. if Player B plays b_1, for all $n \in \mathbb{N}$, if Player A plays (a_1, n), then we obtain the outcome of the game form $[q_0^- - q_1]$ for the action n whereas if Player A plays (a_2, n), then we obtain the outcome of the game form $[\top^- - \bot]$ for the action n.*

Let us first show that this game form is maximizable.

Proposition 5.10. *The game form defined in Definition 5.5 is maximizable w.r.t. Player A.*

Proof. We let $\mathsf{O} := \{q_0, q_1, \top, \bot\}$. Consider a valuation $v : \mathsf{O} \to [0, 1]$. Let $u := \mathsf{val}[\langle \mathcal{F}, v \rangle]$. Clearly, since Player B can play the action b_2 we have:

$$u \leq \frac{v(q_1) + v(\top)}{2}$$

Furthermore, note that if $v(\top) > u$, then Player A has an optimal GF-strategy in the game in normal form $\langle \mathcal{F}, v \rangle$. Indeed, it suffices to play an action (a_1, n) with $n \in \mathbb{N}$ such that $v(\top) \cdot (1 - \frac{1}{2^n}) \geq u$. Assume now that $v(\top) \leq u \leq$

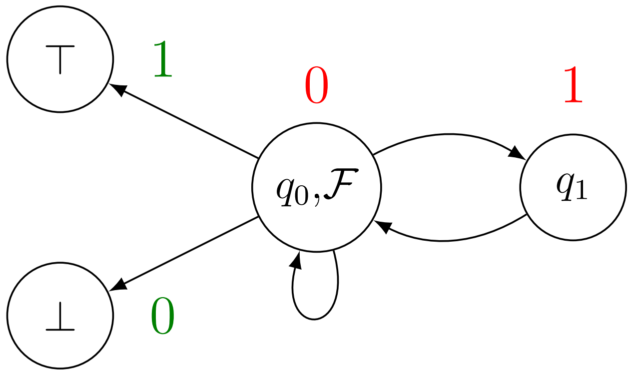

Figure 5.5: A co-Büchi game where Player A wants to see only finitely often the states of color 1: q_1 and $\bot$. The local interaction at state q_0 is depicted in Figure 5.4.

$\frac{v(q_1)+v(\top)}{2}$. It follows that we have $\frac{v(q_1)+v(\top)}{2} \leq v(q_1)$. Therefore, we have $u \leq v(q_1)$. Hence, Player A can play the action $(a_1, 0)$ — if Player B plays b_1, this leads to probability 1 to go to q_1 — to be optimal in the game in normal form $\langle \mathcal{F}, v \rangle$. In any case, Player A has an optimal GF-strategy in the game in normal form $\langle \mathcal{F}, v \rangle$. □

We define below in Definition 5.6 a concurrent co-Büchi game with local interactions maximizable w.r.t. Player A where playing subgame optimally requires infinite choice.

Definition 5.6 (Game depicted in Figure 5.5). *The game of Figure 5.5 has four states $q_0, q_1, \top, \bot$. The two states $\top$ and $\bot$ are stopping states, with $\top$ of value 1 and $\bot$ of value 0. The state q_1 is looping on q_0 and the local interaction of the state q_0 in the game form of Definition 5.5. In particular, $Q_{\mathsf{ns}} = \{q_0, q_1\}$. Player A has a co-Büchi objective with $\mathsf{K} = \{0, 1\}$ and $\mathsf{col}(q_0) := 0$ and $\mathsf{col}(q_1) := 1$.*

Proposition 5.11. *In the co-Büchi game $\mathcal{G}$ of Definition 5.6:*

- *the state q_0 has value 1;*
- *no finite-choice Player-A strategy is optimal;*
- *there is an infinite-choice Player-A strategy that is subgame optimal.*

Proof. • Let $\varepsilon > 0$. Consider a Player-A positional strategy s_A such that $\mathsf{s}_\mathsf{A}(q_0)((a_2, n)) := 1$ for some $n \in \mathbb{N}$ such that $\frac{1}{2^n} \leq \varepsilon$. Then, if Player-B plays action b_1 (i.e. the left column), there is probability $1 - \frac{1}{2^n} \geq 1 - \varepsilon$ to see the state $\top$ of value 1. Furthermore, if Player B plays action b_2 (i.e. the right column), then there is probability $\frac{1}{2}$ to go the state $\top$ and probability $\frac{1}{2}$ to see the state q_1 and loop back to q_0. Hence, if this action is played indefinitely, the state $\top$ is reached almost-surely. In fact, this Player-A positional strategy $\mathsf{s_A}$ has value at least $1 - \varepsilon$. Therefore, the state q_0 has value 1.

- Consider any Player-A positional strategy $\mathsf{s_A}$. If it plays an action of the shape $(a_2, \cdot)$ with positive probability, then if Player B plays action b_1 (i.e. the left column) there is a positive probability to reach the state $\perp$ of value 0. Therefore, such a strategy is not optimal. Otherwise, consider some $n \in \mathbb{N}$ such that $\mathsf{s_A}(q_0)((a_1, n)) > 0$. Then, if Player B always plays the action b_1 (i.e. the left column), at each step there is probability at least $\frac{\mathsf{s_A}(q_0)((a_1,n))}{2^n} > 0$ to see the state q_1. Hence, almost-surely, it seen infinitely often. Therefore, such a strategy has value 0. In fact, no positional Player-A strategy is optimal. Therefore, since this game is standard and B-finite (recall, finitely many states, and Player B has finitely many actions), by Corollary 3.38, no finite-choice Player-A strategy can be optimal from q_0.

- Let us describe a Player A subgame optimal strategy. It is very similar to the one described in the proof of Proposition 5.9. Indeed, for all $\rho \in Q_{\mathsf{ns}}^*$, we let $\mathsf{s_A}(\rho) := \{(a_1, |\rho| + 1) \mapsto 1\}$. Let $\pi \in Q_{\mathsf{ns}}^+$. Player B never has an interest of playing action b_2 since this leads to a stopping state of value 1 with probability $\frac{1}{2}$ and otherwise it loops back on q_0. Hence, let us denote by $\mathsf{s_B}$ the Player-B strategy that always plays action b_1 in q_0, for all $n \in \mathbb{N}$, we have:

$$\mathbb{P}_{\Gamma,\pi}^{\mathsf{s_A},\mathsf{s_B}}(Q_{\mathsf{ns}}^n \cdot (Q_{\mathsf{ns}}^* \cdot q_1)) = \sum_{k \in \mathbb{N}} \sum_{\rho \in \pi \cdot Q_{\mathsf{ns}}^n} \mathbb{P}_{\Gamma,\rho}^{\mathsf{s_A},\mathsf{s_B}}(q_0^k \cdot q_1) = \sum_{k \in \mathbb{N}} \frac{1}{2^{|\rho|+k}} \leq \sum_{k \in \mathbb{N}} \frac{1}{2^{n+k+1}} = \frac{1}{2^n}$$

Hence, by Proposition 1.3 (the continuity of probability measures), we have:

$$\mathbb{P}_{\Gamma,\pi}^{\mathsf{s_A},\mathsf{s_B}}(\bigcap_{n \in \mathbb{N}} Q_{\mathsf{ns}}^n \cdot (Q^* \cdot q_1)) = \lim_{n \to \infty} \mathbb{P}_{\Gamma,\pi}^{\mathsf{s_A},\mathsf{s_B}}(Q_{\mathsf{ns}}^n \cdot (Q_{\mathsf{ns}}^* \cdot q_1)) = 0$$

That is, from π, the state q_1 is seen infinitely often with probability 0. Therefore, the infinite-choice Player-A strategy $\mathsf{s_A}$ is subgame almost-surely optimal.

□

It may seem that this result proves that playing almost-optimally in finite-state co-Büchi games with maximizable[4] local interactions may not be done with positional strategies. Indeed, assume that it is the case, i.e. that positional strategies are sufficient to play almost-surely in such co-Büchi games.

[4]In this paragraph and the next, until Proposition 5.12, we use maximizable for maximizable w.r.t. Player A.

$$\mathcal{F} = \left[\begin{bmatrix} q_0^- - q_1 \end{bmatrix} \quad \frac{\top + q_1}{2} \right]$$

Figure 5.6: The game form $\mathsf{Opt}(\mathcal{F}, v)$ for $\mathcal{F}$ the game form of Definition 5.5 and $v : \mathsf{O} \to [0,1]$ such that $v(\top) = v(q_1) = v(q_0) := 1$ and $v(\bot) := 0$.

This would imply, as in the proof of Proposition 5.8, that with maximizable local interactions, playing subgame almost-surely can be done positionally. Furthermore, (the second part of) Theorem 3.17 (informally) states that the amount of memory to be subgame optimal (when possible) corresponds to the amounts of memory to be subgame almost-surely winning (when possible). Hence, we would obtain that playing subgame optimally with maximizable local interactions can be done positionally, which is in contradiction with Proposition 5.11.

The issue with this argument is that the informal statement of Theorem 3.17 that we gave above is true up to a modification of the local interactions. That modification consists in only considering the Player-A GF-strategies that are optimal w.r.t. some valuations (i.e. the set of local interactions considered in $\mathsf{Opt}(\{\mathsf{F}(q) \mid q \in Q\})$, recall Definition 3.11). In fact, in the game of Definition 5.6, up to this modification, the local interaction at state q_0 is not maximizable, we show this fact below Proposition 5.12. Therefore, even if positional strategies were enough to play almost-optimally in co-Büchi games with maximizable local interactions, it would not imply that playing subgame optimally can be done with positional strategies in co-Büchi games with maximizable local interactions.

Proposition 5.12. *The game form $\mathcal{F}$ of Definition 5.5 is maximizable w.r.t. Player* A*. However, letting $v : \mathsf{O} \to [0,1]$ such that $v(\top) = v(q_1) = v(q_0) := 1$ and $v(\bot) := 0$, the game form $\mathsf{Opt}(\mathcal{F}, v)$ is not maximizable w.r.t. Player* A*.*

Proof. The game form $\mathcal{F}$ is maximizable w.r.t. Player A, by Proposition 5.10. Then, consider the valuation v. The set of optimal Player-A GF-strategies is equal to $\mathcal{D}(\{a_1\} \times \mathbb{N})$. Recall that, letting $\mathcal{F} = \langle \mathcal{D}(\mathsf{Act_A}), \mathcal{D}(\mathsf{Act_B}), \mathsf{O}, \varrho \rangle$, we have $\mathsf{Opt}(\mathcal{F}, v) = \langle \mathsf{Opt_A}(\langle \mathcal{F}, v \rangle), \mathcal{D}(\mathsf{Act_B}), \mathsf{O}, \varrho \rangle$ where $\mathsf{Opt_A}(\langle \mathcal{F}, v \rangle) \subseteq \mathcal{D}(\mathsf{Act_A})$ denotes the set of Player-A GF-strategies optimal in the game in normal form $\langle \mathcal{F}, v \rangle$. The game form $\mathsf{Opt}(\mathcal{F}, v)$ is depicted in Figure 5.6 and is not maximizable w.r.t. Player A. This is witnessed, for instance, by a valuation w such that $w(q_0) := \frac{1}{2}$, $w(\top) := 1$ and $w(q_1) := 0$. □

We summarize the results on co-Büchi games in the theorem below.

Theorem 5.13. *In arbitrary finite-state concurrent co-Büchi games:*

	GF	∃ Opt. ?	ε-Opt.	Optimal	SubG. Opt.
coBuchi	Max. Arb.	No	Pos[1]/? ∞-choice	∞-choice	Pos[1]/∞-choice ∞-choice

Table 5.2: The summary of the situation in arbitrary finite-state concurrent co-Büchi games. The results [1] hold with standard finite local interactions, but a priori does not with arbitrary maximizable local interactions.

- *there does not always exist optimal strategies, even in standard finite games;*
- *with standard finite local interactions, playing almost-optimally can always be done positionally. With arbitrary local interactions, otherwise it may require infinite choice. With arbitrary local interactions maximizable w.r.t. Player* A*, the question is still open.*
- *playing optimally may require infinite choice, even with standard finite local interactions;*
- *when the local interactions are standard finite, playing subgame optimally can be done positionally, otherwise it may require infinite choice.*

These results can be seen in Table 5.2.3.

Proof.
- that is because this holds in reachability games, see Theorem 4.12;
- this is already known, see [50, Theorem 3.1]. The fact that this does not hold anymore for arbitrary local interactions is a consequence of the fact that this is the case for the safety objective, see Proposition 4.4. The case of arbitrary local interactions maximizable w.r.t. Player A is still open (see 5.1).
- we have first proved this result in [41, Section 6]. We state it in Proposition 5.7;
- We have also proved this result in [41, Corollary 1]. However, it does not hold with arbitrary local interactions, see Proposition 5.9, even if all local interactions are maximizable w.r.t. Player A, see Proposition 5.11.

□

5.3 Parity objectives

Recall that we only consider games with finitely many states. We already know that there does not always exist optimal strategies in standard finite parity games since this holds in reachability games (see 4.12). We also know

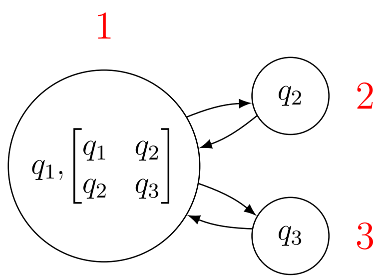

Figure 5.7: A parity game.

that playing almost-optimally may require infinite choice in standard finite parity games since this holds in Büchi games (see 5.5). Furthermore, playing optimally (when possible) may require infinite choice in standard finite parity games since this holds in co-Büchi games (see 5.13). In fact, playing subgame optimally (when possible) may also require infinite choice. That is witnessed by the game defined in Definition 5.7. Note that we already talked about this game in Subsection 3.4.2, as it was depicted in Figure 3.12.

Definition 5.7 (Game depicted in Figure 5.7). *The game $\mathcal{G}$ of Figure 5.7 has three states: $Q = \{q_1, q_2, q_3\}$. The two states q_2 and q_3 have a trivial local interaction and loop back to q_1. We denote by a_1 and a_2 the two actions available to Player A in q_1 where a_1 refers to the top row and a_2 refers to the bottom row. Similarly, We denote by b_1 and b_2 the two actions available to Player B in q_1 where b_1 refers to the leftmost row and b_2 refers to the rightmost column. Player A has a parity objective with $\mathsf{K} = \{1, 2, 3\}$, $\mathsf{col}(q_1) := 1$, $\mathsf{col}(q_2) := 2$ and $\mathsf{col}(q_3) := 3$. That is, Player A wants to see only finitely often the state q_3 while seeing infinitely often the state q_2.*

Proposition 5.14. *The game of Definition 5.7 is such that:*

- *all finite-choice Player-A strategies have value 0 from q_1;*
- *there is an infinite-choice Player-A strategy that is subgame almost-surely winning.*

Proof.
- Consider any positional Player-A strategy $\mathsf{s_A}$. If $\mathsf{s_A}(q_0)(a_2) > 0$ (i.e. if the bottom row is played with positive probability), then by playing the action b_2 (i.e. the rightmost column) Player B ensures that almost-surely the state q_3 is seen infinitely often. Otherwise, if $\mathsf{s_A}(q_0)(a_1) = 1$, then by playing the action b_1 (i.e. the leftmost column) Player B ensures that the state q_2 is never seen. In fact, the value of all Player-A positional strategies is 0. Therefore, since this game is standard finite, by Corollary 3.38, all finite-choice Player-A strategies have value 0 from q_1.

- Let us now describe a Player-A subgame almost-surely winning strategy. For all $\rho \in Q^+$, we let $|\rho|_{2,3} \in \mathbb{N}$ denote the number of times the finite paths ρ has visited the states q_2 and q_3. Consider also a sequence $(\varepsilon_k)_{k\in\mathbb{N}}$ such that, for all $k \in \mathbb{N}$, we have $\varepsilon_k \in (0,1]$ with $\sum_{k\in\mathbb{N}} \varepsilon_k < \infty$. We define a Player-$\mathsf{A}$ strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ as follows, for all $\rho \in Q^+$, we have:

$$\mathsf{s_A}(\rho) := \{a_1 \mapsto 1 - \varepsilon_{|\rho|_{2,3}},\ a_2 \mapsto \varepsilon_{|\rho|_{2,3}}\}$$

Consider now any Player-B deterministic strategy $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$ and some $\pi \in Q^+$. Let $n \in \mathbb{N}$. We have:

$$\begin{aligned}\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi}[Q^n \cdot q_1^\omega] &= \sum_{\rho\in\pi\cdot Q^n} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[q_1^\omega] = \sum_{\rho\in\pi\cdot Q^n} \lim_{k\to\infty} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[q_1^k] \\ &\leq \sum_{\rho\in\pi\cdot Q^n} \lim_{k\to\infty} (1-\varepsilon_{|\rho|_{2,3}})^k = 0\end{aligned}$$

As this holds for all $n \in \mathbb{N}$, it follows that, from π, almost-surely, the set of states $\{q_2, q_3\}$ is seen infinitely often.

Let us now show that almost-surely the state q_3 is seen only finitely often. Let $n \in \mathbb{N}$ and consider some $\rho \in (\{q_1,q_2\}^* \cdot q_3)^n \cdot \{q_1,q_2\}^*$. Since the strategy $\mathsf{s_B}$ is deterministic, for all $i \in \mathbb{N}$, we have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi\cdot\rho}[\{q_1\}^i \cdot q_3] = 0 \text{ or } \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi\cdot\rho}[\{q_1\}^{i+1}] = 0$$

Therefore, there is at most one $i \in \mathbb{N}$ such that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi\cdot\rho}[q_1^i \cdot q_3] > 0$ with $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi\cdot\rho}[q_1^i \cdot q_3 \mid q_1^i] \leq \varepsilon_{n+k}$ where k denotes the number of times the state q_2 occurs in ρ after the last q_3. Therefore, $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi\cdot\rho}[q_1^* \cdot q_3] \leq \varepsilon_{n+k}$. It follows that, for all $\theta \in (\{q_1,q_2\}^* \cdot q_3)^n$:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi\cdot\theta}[\{q_1,q_2\}^* \cdot q_3] = \sum_{k\in\mathbb{N}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi\cdot\theta}[(q_1^* \cdot q_2)^k \cdot q_1^* \cdot q_3] \leq \sum_{k\in\mathbb{N}} \varepsilon_{n+k}$$

Therefore, by Proposition 1.3 (the continuity of probability measures):

$$\begin{aligned}\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi}[\bigcap_{n\in\mathbb{N}} (\{q_1,q_2\}^* \cdot q_3)^n] &\leq \lim_{n\to\infty} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi}[(\{q_1,q_2\}^* \cdot q_3)^{n+1}] \\ &\leq \lim_{n\to\infty} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\pi}[\{q_1,q_2\}^* \cdot q_3 \mid (\{q_1,q_2\}^* \cdot q_3)^n] \\ &\leq \lim_{n\to\infty} \sum_{k\in\mathbb{N}} \varepsilon_{n+k} = 0\end{aligned}$$

This comes from the fact that $\sum_{k\in\mathbb{N}} \varepsilon_k < \infty$. Therefore, the state q_3 is seen infinitely often with probability 0. That is, against the Player-B deterministic strategy $\mathsf{s_B}$, the parity objective is ensured almost-surely from π. Since this holds for all Player-B deterministic strategies and by Corollary 2.17 (since we obtain an MDP once Player A has fixed her

	GF	∃ Opt. ?	ε-Opt.	Optimal	SubG. Opt.
Parity	M./A.	No	∞-choice	∞-choice	∞-choice

Table 5.3: The summary of the situation in arbitrary finite-state concurrent parity games with at least three colors.

strategy), the value of the Player-A strategy $\mathsf{s_A}$ is 1 from π. As this holds for all $\pi \in Q^+$, it follows that the infinite-choice Player-A strategy $\mathsf{s_A}$ is subgame almost-surely winning.

□

We summarize how concurrent parity games behave below:

Theorem 5.15. *In arbitrary finite-state concurrent parity games:*

- *there does not always exist optimal strategies;*
- *playing almost-optimally may require infinite choice;*
- *playing optimally, when possible, may require infinite choice;*
- *playing subgame optimally, when possible, may require infinite choice.*

All of these results can be witnessed by standard finite games, with at most three colors. These results can be seen in Table 5.3.

Proof.

- this was already the case for reachability games, see Theorem 4.12;
- this was already the case for Büchi games, see Theorem 5.5;
- this was already the case for co-Büchi games, see Theorem 5.13;
- It was already known that playing subgame optimally may require infinite memory, see [47, Theorem 7]. We have further proved that it may require infinite choice in Proposition 5.14 by using the same game used in [47].

□

5.4 Discussion and future work

In this chapter, we have studied Büchi, co-Büchi and parity objectives. As stated in Subsection 5.2.1, we leave Open Question 5.1 unanswered. A possible direction to try and answer this question could be a local-global transfer. We discuss this notion extensively in Part III.

In the previous chapter, among other things, we have designed a procedure to compute, in reachability games, the set of states $\mathsf{Opt_A}$ from which Player A has an optimal strategy. In turn, this allowed us to establish Theorem 4.11: for all positive $\varepsilon > 0$, Player A has a positional strategy that is optimal from all states in $\mathsf{Opt_A}$, and ε-optimal from all other states in $Q \setminus \mathsf{Opt_A}$. It seems natural to look for a similar result in the Büchi, co-Büchi and parity games we have studied in this chapter. We believe that it could be possible to obtain one for Büchi objectives: for all positive $\varepsilon > 0$, Player A has a strategy that is optimal and positional[5] from all states in $\mathsf{Opt_A}$, and ε-optimal from all other states in $Q \setminus \mathsf{Opt_A}$. The difference with the reachability case is that the strategy may have infinite choice at states in $Q \setminus \mathsf{Opt_A}$. Such a result could be obtained by using arguments akin to the ones used in Section 5.1 to prove Proposition 5.2. As a corollary, we would obtain that this also holds with safety objectives, since safety games can be seen as special cases of Büchi games. However, this cannot be extended to co-Büchi objectives and to parity objectives, since, with these objectives, playing optimally may require infinite choice.

Finally, we would like to discuss a possible future work that extends some of what we have done in this chapter. In [70], the authors study Muller objectives. Given a finite set K of colors, a Muller objective is defined by a set $S \subseteq 2^{\mathsf{K}}$ of subsets of colors such that an infinite sequence of colors is winning for Player A if and only if the set of colors seen infinitely often is in S. In particular, Muller objectives are PI and more general than parity objectives. In that paper [70], among other things, the authors show that in finite turn-based games, for those Muller objectives that are upward-closed, positional strategies are sufficient to be almost-surely winning for Player A. A Muller objective defined by $S \subseteq 2^{\mathsf{K}}$ is upward-closed if, for all $C, C' \in 2^{\mathsf{K}}$ if $C \in S$ and $C \subseteq C'$, then $C' \in S$. Note that, by [58, Theorem 4.5] — alternatively, Corollary 3.24 — we can deduce that with upward-closed Muller objectives, Player A always has positional optimal strategies in finite turn-based games.

We believe that this could be extended to standard concurrent games, up to adding the assumption that optimal strategies do exist. We conjecture the following:

Conjecture 5.16. *In all standard finite concurrent games with an upward-closed Muller objective, if Player* A *has an optimal strategy, then she has one that is positional.*

An upward-closed Muller objective can be seen as a disjunction of conjunctions of Büchi objectives. That is, Player A wants to see infinitely often as many states as possible. Therefore, we believe that, to play optimally, Player A only needs to play a locally optimal strategy whose support, at each state, is maximal (which is possible because the game is standard). We also believe that what is stated in Conjecture 5.16 gives a characterization of upward-closed Muller objectives. That is, given any Muller objective that is not upward-

[5] In other words, the strategy plays a single GF-strategy at each local interaction at states in $\mathsf{Opt_A}$.

closed, there is a standard finite game where Player A can play optimally, but it can only be done with infinite-choice strategies. Note that upward-closed Muller objectives are also characterized in [70].

Following a similar idea, we believe that it is possible to characterize those Muller objectives for which, in standard finite concurrent games, whenever there is a subgame optimal strategy, there is a finite-memory one. As can be seen in Table 3.1, in finite parity games with at least three colors, infinite-choice strategies may be required to play subgame optimally. Parity objectives with at least three colors can be written as a conjunction of a Büchi and a co-Büchi objective. We believe that this is the issue. We make the conjecture below.

Conjecture 5.17. *In all standard finite concurrent games with a Muller objective that can be written as a disjuntion of either a conjunction of Büchi or a conjunction of co-Büchi objectives, if Player* A *has an optimal strategy, then she has a finite-memory one.*

Part III

Restricting game forms in concurrent games

This final part is arguably the most important one of this dissertation: here, we adopt an entirely new approach towards concurrent games, which we believe is promising. The general idea is the following. Since concurrent (uncolored) arenas behave poorly in general, especially compared to turn-based (uncolored) arenas, see for instance the previous part, we restrict ourselves to subsets of concurrent (uncolored) arenas. These subsets will always strictly include turn-based arenas and contain only arenas enjoying some of the nice properties that turn-based arenas enjoy. To properly define these subsets, we consider a set of colors K and a win/lose objective $W \subseteq \mathsf{K}^\omega$ — say a reachability objective — and a type of strategies τ — say optimal positional strategies. Then, we want to define a subset of (W, τ)-well-behaved concurrent uncolored arenas. Informally, we define (W, τ)-well-behaved concurrent uncolored arenas as arenas for which in all colored games with objective W that can be obtained from them, there are τ-strategies, for one or both of the players.

The main novelty lies in the way we define these subsets of concurrent uncolored arenas. Indeed, they are defined via the crucial notion of *game form*, and *local interaction*. That is, given an objective W and a type of strategies τ, the goal is to identify a set of game forms $S_{(W,\tau)}$ such that all the concurrent uncolored arenas whose local interactions are included in $S_{(W,\tau)}$ are (W, τ)-well-behaved.

To define this set $S_{(W,\tau)}$, we proceed in two steps, described below.

- First, we characterize the game forms that are individually well-behaved w.r.t. (W, τ). More precisely, a game form $\mathcal{F}$ is individually well-behaved w.r.t. (W, τ) if all "simple arenas" that can be built from $\mathcal{F}$ are (W, τ)-well-behaved. Informally, a simple arena built from $\mathcal{F}$ is an arena with only one non-trivial local interaction, $\mathcal{F}$, which occurs in a central state q_{init}. Every other state is either stopping or trivial (i.e. with only one possible distributions over successor states regardless of what the players do) and looping back to q_{init}. That way, the only source of interaction, and therefore concurrency, in the arena comes from $\mathcal{F}$.

- Second, we prove that all concurrent arenas where all local interactions are individually well-behaved w.r.t. (W, τ) are (W, τ)-well-behaved. We then define the set $S_{(W,\tau)}$ to be equal to the set of game forms individually well-behaved w.r.t. (W, τ).

When this second step is conclusive, in some cases it will be the case even for arenas with infinitely many states, however in some other cases it will be only for finite-state arenas.

Assuming that we have achieved both of these steps, the set $S_{(W,\tau)}$ of game forms can be seen as maximal w.r.t. (W, τ). Indeed, we can reformulate the two steps above as follows:

- all game forms not in $S_{(W,\tau)}$ behave poorly w.r.t. (W, τ), even individually — since, given any game form that is individually poorly-behaved (i.e. non individually well-behaved) w.r.t. (W, τ), there is a simple arena built on $\mathcal{F}$ that is not (W, τ)-well-behaved;

- all game forms in $S_{(W,\tau)}$ behave well w.r.t. (W, τ), even collectively — since all concurrent arenas with local interactions individually well-behaved w.r.t (W, τ) are (W, τ)-well-behaved.

Alternatively, we could say that being individually well-behaved w.r.t. (W, τ) is a local necessary and sufficient condition for arenas to be (W, τ)-well-behaved. We formally define this notion in page 247.

As argued above, the set $S_{(W,\tau)}$ of game forms can be seen as maximal w.r.t. (W, τ). However note that for various (W, τ), there exists a concurrent arena with some individually poorly-behaved local interactions that is still (W, τ)-well-behaved. The only thing we claim about game forms individually poorly-behaved w.r.t. (W, τ) is that there are some simple arenas built on them that are not (W, τ)-well-behaved. However, it is not the case of all the arenas built on them.

The main purpose of this part is to apply this two-step procedure to define sets of game forms $S_{(W,\tau)}$ for various win/lose objectives W and types of strategies τ. In addition, we will perform this transfer not only on win/lose objectives, but also on more general payoff functions. In fact, we will perform this transfer with sets S of (non necessarily win/lose) payoff functions. In that case, a game form will be deemed individually safe w.r.t. (S, τ) if it is individually safe w.r.t. (f, τ), for all payoff functions $f \in S$.

Benefits of this characterization. Let us now give an idea of how the above characterization can be used in practice. We believe that the main benefit of characterizing the local interactions that are individually well-behaved w.r.t. payoff functions and types of strategies lies in the design of games. Indeed, when designing a game with a specification in mind — for instance, the existence of positional optimal strategies in finite-state reachability games — our characterization provides exactly the safe building blocks (i.e. local interactions) that can be used to ensure that the desired specification holds in the compound arenas. Furthermore, as long as they are individually well-behaved, all building blocks can be used, regardless of the other individually well-behaved blocks used in the game. Hence, the design of games can be done locally, without knowing the other local components, or even the number of components involved.

As mentioned above, the desired specification holds in any compound arena whose building blocks are individually well-behaved. In addition this arena can be dynamically modified while maintaining that this specification holds. This modification may consist in adding an individually well-behaved building block to this compound arena, removing one building block or replacing a building block by another individually well-behaved one.

Arguably, it has a second benefit, though it may not have a strong practical value. Indeed, to decide if a game $\mathcal{G}$ with objective W enjoys the existence of τ-strategies, one may consider every local interaction occurring in the underlying arena, decide if they are all individually well-behaved w.r.t. (W, τ) (assuming it is possible), and if so conclude that τ-strategies exist in $\mathcal{G}$. However, it has two main drawbacks: first, this requires to handle all local interactions, which may be costly. Second, we can only conclude if all local interactions are individually well-behaved. Perhaps this approach is promising in contexts where we already know, a priori, that a large amount of the local interactions involved are individually well-behaved w.r.t. (W, τ), and there are only a few of them to check.

Some formal definitions. Before we give an overview of what we do specifically in each chapter of this part, let us formally introduce the notion of simple games built from a game form as we will use this notion throughout this part. Recall Definition 1.21: a game form $\mathcal{F}'$ is obtained from a game form $\mathcal{F}$ if $\mathcal{F}'$ is equal to $\mathcal{F}$ up to a — not necessarily injective — renaming of the outcomes. Also recall Definition 1.11: a game form is trivial if both players have only one available GF-strategy.

Definition 5.1 (Simple games). *Consider a game form $\mathcal{F}$. A* simple game *built on $\mathcal{F}$ is a game $\mathcal{G}$ such that:*

- *there is a central state $q_{\mathsf{init}} \in Q_{\mathsf{ns}}$ such that $\mathsf{F}(q_{\mathsf{init}})$ is obtained from $\mathcal{F}$;*
- *all non-stopping non-central states $q \in Q_{\mathsf{ns}} \setminus \{q_{\mathsf{init}}\}$ are trivial and looping on the state q_{init}.*

Note that the game $\mathcal{G}$ is indeed built from $\mathcal{F}$ according to Definition 1.22.

In this part, we prove instances of what we call local-global transfers. These correspond to the (somewhat) necessary and sufficient condition discussed earlier in page 245. We define below how we will formulate these transfers in this part.

Definition 5.2 (Necessary and Sufficient Condition Transfer, NSC-Transfer). *Let X denote a set of payoff functions. A game with payoff function in X is called an X game. Let also Y be a subset of game forms — for instance, it may be the set of standard game forms. Consider a predicate φ on X games and a predicate ψ_{GF} on Y game forms. When we say that: "among Y game forms, satisfying ψ_{GF} is an NSC-transfer for (possibly finite, possibly without stopping states) X games to satisfy the property φ", it means that:*

1. *from any game form $\mathcal{F}$ in Y that does not satisfy the predicate ψ_{GF}, one can build on $\mathcal{F}$ a simple (possibly finite, possibly without stopping states) X game that does not satisfy the predicate φ;*

2. *all (possibly finite, possibly without stopping states) X games, whose local interactions are game forms in Y satisfying the predicate ψ_{GF}, satisfy the predicate φ.*

We state several results in this part with the notion of NSC-transfer: Proposition 6.1, Theorem 6.6, Proposition 7.1, Theorem 7.5, Proposition 7.7 and Corollary 8.10.

Let us argue why what we have defined in Definition 5.2 is called a necessary and sufficient condition. Consider an NSC-transfer statement as in Definition 5.2: among Y game forms, satisfying ψ_{GF} is an NSC-transfer for X games to satisfy the property φ. This statement hides in fact an equivalence that could be stated as follows. Consider a non-empty set S_{GF} of Y game forms. Then, the two following assertions are equivalent:

a. The set S_{GF} only contains game forms satisfying the predicate ψ_{GF};

b. All X games built on S_{GF} satisfy the property φ.

Indeed, since it is an NSC-transfer, item a. implies item b. by item 2. of Definition 5.2 and item b. implies item a. by item 1. of Definition 5.2. Thus, all NSC-transfers hide an equivalence stated with sets of game forms (hence the NSC terminology).

This part contains four chapters. Chapter 6 deals with the restrictions on game forms to be used in concurrent games so that they ensure nice properties that can be directly deduced from Theorem 2.3. In particular, we provide an NSC-transfer for the existence of winning strategies in concurrent games. In the next two chapters, we state NSC-transfers for the existence of positional optimal strategies in parity games. Specifically, in Chapter 7, we consider arbitrary local interactions and show NSC-transfers for safety, reachability and Büchi objectives. In Chapter 8, we consider only standard finite game forms, but we prove NSC-transfers for arbitrary parity objectives. Finally, in Chapter 9, we study the different classes of game forms we have defined in this part. Note that these classes are studied outside of any concurrent game context.

6 - Game forms for general objectives

In this chapter, we make use of the new version of Blackwell determinacy stated in Theorem 2.3. Specifically, we use items $(1.a)$ but also (2) of this theorem to obtain two things rather straightforwardly. First, somewhat necessary and sufficient conditions on game forms that, when they are ensured by all the local interactions in a game behaves in a good way, i.e. NSC-transfer. Second, sufficient conditions on game forms that, when they are ensured by all the local interactions in a game, ensure that the game behaves in a good way. In the next two paragraphs, we present the two NSC-transfers of this chapter.

We first realize that item (2) alone of Theorem 2.3 gives straightforwardly an NSC-transfer on arbitrary game forms for infinite games (with arbitrary payoff functions) to have a value. This is stated in Proposition 6.1.

We then consider the existence of winning strategies. There are very simple concurrent games in which no player has a winning strategy. We consider, among standard deterministic game forms, the ones ensuring the existence of winning strategies in infinite concurrent games. These are called determined game forms. We obtain an NSC-transfer stated in Theorem 6.6. The proof is direct from Theorem 2.3 (items $(1.a), (2)$). Note that we have studied determined game forms in [38]. In that paper, we have already shown (see [38, Theorem 17]) that, when used in concurrent games, determined game forms ensure the existence of winning strategies from every state for either of the players. However, in this paper, the proof is more elaborate since could not use Theorem 2.3 as we had not proved it yet. Instead, we used the notions of parallelization and sequentialization of games and strategies, that we have defined in this dissertation in Section 3.4.

Nonetheless, in [38], in addition to proving the existence of winning strategies, we also show that the memory requirement to win in concurrent games with determined local interactions is the same as in (deterministic) turn-based games. This is stated, without a proof, in Theorem 6.7. As an additional remark on determined game forms, Theorem 6.6 is a generalization of Borel determinacy (Theorem 2.1), that uses Borel determinacy as a black box in its proof.

Finally, we quickly discuss in Subsection 6.2.2 an application of determined game forms: discrete-bidding games.

The applications of Theorem 2.3 to determined and valuable game forms actually constitute the only NSC-transfer that we will give in this chapter. In the remainder of this chapter, we will slightly modify the definition of determined game forms to obtain sometimes weaker, sometimes incomparable restrictions on game forms.

When proving Theorem 6.6, we actually use a definition (see Proposition 6.5) of determined game forms that is equivalent to the one we gave in [38], but that is more suited for the application of Theorem 2.3. Another interest of this equivalent definition is that it is easier to generalize. The first generalization we propose makes the determinacy of game forms asymmetric in the players: we obtain game forms that are semi-determined w.r.t. Player A (or Player B). When such game forms are used in concurrent games, they ensure that the game has a value, that the value of every state is either 0 or 1, and that from every state where this value is 1, Player A has a winning strategy. This is stated in Proposition 6.9.

The second generalization we consider induces the definition of game forms finitely maximizable w.r.t. a player. These are standard finite game forms that are maximized by a finite set of GF-strategies for Player A (or Player B), see Definition 6.3. When these game forms are used in a finite concurrent arena $\mathcal{C}$, the arena $\mathcal{C}$ behaves, somehow, like a finite turn-based arena. We give two applications. In games obtained from such arenas with a payoff function that is PI upward well-founded, Player A has a subgame optimal strategy, see Theorem 6.11. This generalizes Corollary 3.25 (the same result in the context of turn-based games). Furthermore, if this payoff function corresponds to a parity objective, Player A has a positional strategy that is optimal, see Corollary 6.12. To prove both these applications, we use the aforementioned notions of parallelization and sequentialization of games and strategies from Section 3.4.

Finally, we present a strengthening of the notion of finitely maximizable game forms: uniquely maximizable game forms. These are arbitrary game forms that are maximized by a single Player-A GF-strategy. We first show that, though this requirement is very strong, there is a natural way to construct uniquely maximizable game forms, see Proposition 6.14. We then show that, when all local interactions of an arena are uniquely maximizable w.r.t. Player A, she has a positional strategy that is (subgame) optimal regardless of the payoff function considered, see Theorem 6.15.

As mentioned above, with these last three classes of game forms, we do not state an NSC-transfer. Furthermore, everything we prove in this chapter, except what does concern determined game forms, is unpublished.

6.1 Valuable game forms

We start with a very brief section dealing with valuable game forms. This first local-global transfer in concurrent games is a straightforward consequence of Theorem 2.3 (item 2). Indeed, if one only considers this item, this theorem can be read as follows: as soon as all local interactions are valuable, the game has a value. Furthermore, it is straightforward that from a game form that is not valuable, one can build a simple form that does not have a value. Hence,

$$\mathcal{F}_1 = \begin{bmatrix} x & x & z \\ x & y & y \\ z & y & z \end{bmatrix}$$

Figure 6.1: A game form (that is determined).

$$\mathcal{F}_2 = \begin{bmatrix} x & x & z \\ x & z & y \\ z & y & y \end{bmatrix}$$

Figure 6.2: A game form (that is not determined).

we obtain the proposition below:

Proposition 6.1. *Among arbitrary game forms, being valuable is an NSC-transfer for infinite games (with measurable payoff functions into* $[0, 1]$*) to have a value.*

Proof. Consider any game form $\mathcal{F}$ that is not valuable. Let $v : \mathsf{O} \to [0, 1]$ be a valuation of the outcomes of $\mathcal{F}$ such that the game in normal form $\langle \mathcal{F}, v \rangle$ does not have a value. We define an arena $\mathcal{C} = \langle Q, \mathsf{F}, \mathsf{K}, \mathsf{col} \rangle$ such that $Q_{\mathsf{ns}} := \{q_{\mathsf{init}}\}$ (with $q_{\mathsf{init}} \notin \mathsf{O}$, up to a renaming of the outcomes) with $\mathsf{F}(q_{\mathsf{init}}) := \mathcal{F}$. Furthermore, $Q_{\mathsf{s}} := \mathsf{O}$ with, for all $q \in Q_{\mathsf{s}}$, $\mathsf{val}(q) \leftarrow v(q) \in [0, 1]$. Furthermore, K and col are defined arbitrarily. With this construction, we have that, regardless of the payoff function $f : \mathsf{K}^\omega \to [0, 1]$ considered, for $\mathcal{G} := \langle \mathcal{C}, f \rangle$, we have $\chi_{\mathcal{G}}[\mathsf{A}](q_{\mathsf{init}}) = \mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{A}]$ and $\chi_{\mathcal{G}}[\mathsf{B}](q_{\mathsf{init}}) = \mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{B}]$. Therefore, the game $\mathcal{G}$ does not have a value.

Conversely, item 2. of Theorem 2.3 gives that all concurrent game whose local interactions are valuable game forms has a value. □

6.2 Determined game forms

In this section, we are looking for a local-global transfer to ensure that, in infinite win/lose games without stopping states, from every state, either of the player has a winning strategy. Since we are considering winning strategies (recall Definition 1.33), we only consider deterministic standard game forms. Recall that, for such a game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ on some set of outcomes O, we have that for all $a \in \mathsf{Act}_{\mathsf{A}}$ and $b \in \mathsf{Act}_{\mathsf{B}}$, $\varrho(a, b) \in \mathsf{O}$. We have studied this question in [38]. However, the main result of this section, Theorem 6.4 below, that was also central in [38] is now a direct consequence of Theorem 2.3 (in particular, of items 1.a and 2).

Compared to the conditions on game forms we will consider in the next chapter, the condition we consider in this section is rather simple to explain. Our goal is to come up with a definition of game form that ensures the existence of winning strategies. Consider some game forms such as the ones depicted in Figures 6.1, 6.2. A natural way to define win/lose games from these game forms

$$\langle \mathcal{F}_1, v \rangle = \begin{bmatrix} 1 & 1 & 1 \\ 1 & 0 & 0 \\ 1 & 0 & 1 \end{bmatrix}$$

Figure 6.3: The game form $\mathcal{F}_1$ of Figure 6.1 with the valuation $v : \{x, y, z\} \to \{0, 1\}$ such that $v(x) = v(z) := 1, v(y) := 0$.

$$\langle \mathcal{F}_1, v \rangle = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

Figure 6.4: The game form $\mathcal{F}_1$ of Figure 6.1 with the valuation $v : \{x, y, z\} \to \{0, 1\}$ such that $v(x) := 1, v(y) = v(z) := 0$.

is the following: every outcome is mapped to either 1 (i.e. winning for Player A) or 0 (i.e. winning for Player B). Then, a Player-A winning (GF-)strategy in this game is a row on which there are only outcomes winning for her, i.e. mapped to 1. This is the case of the top row in Figure 6.3. Symmetrically, a Player-B winning (GF-)strategy in this game is a column on which there are only outcomes mapped to 0. This is the case of the leftmost column in Figure 6.4. However, neither of the players has a winning (GF-)strategy in Figure 6.5. In fact, game forms for which there are always winning GF-strategies as described above are called determined and are the subject of study of this section. This notion already exists, see for instance [36], where determinacy is referred to as $(0, 1)$-solvable, or [71] where determinacy is referred to as tightness.

Definition 6.1 (Determined game form)**.** *Consider a set of outcomes* O *and a standard deterministic game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$*. It is* determined *if, for all* $w : \mathsf{O} \to \{0, 1\}$*:*

- *either Player* A *has a winning (*GF*-)strategy, i.e. there is some* $a \in \mathsf{Act}_\mathsf{A}$ *such that* $w \circ \varrho(a, \mathsf{Act}_\mathsf{B}) = \{1\}$*;*
- *or Player* B *has a winning (*GF*-)strategy, i.e.there is some* $b \in \mathsf{Act}_\mathsf{B}$ *such that* $w \circ \varrho(\mathsf{Act}_\mathsf{A}, b) = \{0\}$*.*

Concerning the game forms of Figure 6.1 and 6.2, one can realize that $\mathcal{F}_1$ is determined. One can check it by looking at all the possible valuations of the outcomes. (The winning players is the one for whom at least two elements in $\{x, y, z\}$ are mapped to her winning outcome.) However, the game form $\mathcal{F}_2$ is not determined, as witnessed in Figure 6.5.

Furthermore, it is straightforward that all turn-based deterministic game forms are determined.

Proposition 6.2. *All turn-based deterministic game forms are determined.*

Proof. Consider a turn-based deterministic game form $\mathcal{F}$. Assume Player A plays alone, the other case being symmetrical. Consider any map $\mathsf{O} \to \{0, 1\}$. If all outcomes are mapped to 0, then clearly Player B wins. Otherwise, there is

$$\langle \mathcal{F}_2, v \rangle = \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \end{bmatrix}$$

Figure 6.5: The game form $\mathcal{F}_2$ of Figure 6.2 with the valuation $v : \{x, y, z\} \to \{0, 1\}$ such that $v(x) = v(y) := 1, v(z) := 0$.

some outcome in O mapped to 1 that Player A can enforce (as the game form is turn-based). In any case, either of the players has a winning (GF-)strategy. $\square$

It is rather straightforward that any game form that is not determined is unsafe w.r.t. the existence of winning strategies in win/lose concurrent games. This is stated below.

Proposition 6.3. *Consider a set of outcomes O and a standard deterministic game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$. If the game form $\mathcal{F}$ is not determined, then there is a simple win/lose game without stopping states built from $\mathcal{F}$ in which no player has a winning strategy from the state q_{init}.*

Proof. Let $w : \mathsf{O} \to \{0, 1\}$ be a valuation of the outcomes witnessing that the game form $\mathcal{F}$ is not determined. We define an arena $\mathcal{C} = \langle Q, \mathsf{F}, \mathsf{K}, \mathsf{col} \rangle$ such that $Q = Q_{\mathsf{ns}} := \{q_{\mathsf{init}}, 1, 0\}$ with $\mathsf{K} := \{0, 1\}$ and $\mathsf{col}(q_{\mathsf{init}}) = \mathsf{col}(0) := 0$ and $\mathsf{col}(1) := 1$. In addition, denoting $\mathcal{F} = \langle \mathsf{Act_A}, \mathsf{Act_B}, \mathsf{O}, \varrho \rangle$, the local interaction at state q_{init} is equal to $\mathsf{F}(q_{\mathsf{init}}) := \langle \mathsf{Act_A}, \mathsf{Act_B}, \{0, 1\}, \mathbb{E}_v(\varrho) \rangle$. That is, given a pair of actions $(a, b) \in \mathsf{Act_A} \times \mathsf{Act_B}$, the next state reached after q_{init} if (a, b) is played is equal to $v \circ \varrho(a, b)$. Both states 1 and 0 are trivial and loop back on q_{init}. We consider the win/lose objective $W_{\mathsf{A}} \subseteq \mathsf{K}^\omega$ such that $W_{\mathsf{A}} := \{0 \cdot 1 \cdot \mathsf{K}^\omega\}$. That is, from q_0, Player A wins if and only if the next state seen is of color 1, i.e. if the next state is 1. Note that the game $\mathcal{G}$ is indeed a simple game built from $\mathcal{F}$, recall Definition 5.1.

Let us now show that in the game $\mathcal{G} = \langle \mathcal{C}, W_{\mathsf{A}} \rangle$, neither of the player has a (deterministic) winning strategy from q_{init}. Consider any deterministic Player-A strategy $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$. By choice of the valuation v, there is an action $b \in \mathsf{Act_B}$ such that $w \circ \varrho(\mathsf{s_A}(q_{\mathsf{init}}), b) = 0$. A deterministic Player-B strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$ such that $\mathsf{s_B}(q_{\mathsf{init}}) := b$ wins surely against the Player-A strategy $\mathsf{s_A}$ from q_{init}. In fact, Player A has no winning strategy from q_{init}. We can show similarly that Player B has no winning strategy from q_{init}. $\square$

The question now is, are there always winning strategies in win/lose games where all local interactions are determined. It is in fact the case, as stated in the theorem below.

Theorem 6.4. *Consider a concurrent win/lose game $\mathcal{G}$ without stopping states such that, for all $q \in Q$, the local interaction $\mathsf{F}(q)$ is determined. Then, from every state $q \in Q$, either of the players have a winning strategy.*

This statement corresponds to [38, Theorem 17]. In that paper, we have proved this theorem by using the notions of parallelization and sequentialization of strategies that we introduced in Section 3.4 from Chapter 3. The idea in that paper is to consider winning strategies in the sequentialized version of the game — which are ensured to exist by Theorem 2.1 (the determinacy of Borel games) — and translate them back into the original concurrent game by using the determinacy of the local interactions.

However, with the help of Theorem 2.3, it can be proved in a much quicker fashion. Specifically, we want to use item 1.a of Theorem 2.3. However, to do so, we need to express the fact that a game form is determined with the notion of sets of GF-strategies supremizing game forms. This is done in the proposition below where we give an equivalent definition of determined game forms.

Proposition 6.5. *Consider a set of outcomes* O *and a standard deterministic game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$. *It is determined if and only if it is valuable and supremized w.r.t. Player* A *by* Act_A *and w.r.t. Player* B *by* Act_B *(i.e.* ε*-optimal* GF*-strategies can be found among deterministic* GF*-strategies).*

Note that, to prove Theorem 6.4, we do not need an equivalence only the implication assuming the game form is determined.

Proof. Assume that the game form $\mathcal{F}$ is determined. Consider any valuations of the outcome $v : \mathsf{O} \to [0,1]$. For all $u \in [0,1]$, we let $\mathsf{O}_{\geq u} := \{o \in \mathsf{O} \mid v(o) \geq u\}$ and $v_{\geq u} : \mathsf{O} \to \{0,1\}$ such that $v_{\geq u}^{-1}[1] := \mathsf{O}_{\geq u}$. We then let $\mathsf{Win}_\mathsf{A} := \{u \in [0,1] \mid \exists a \in \mathsf{Act}_\mathsf{A},\ v_{\geq u}[\varrho(a, \mathsf{Act}_\mathsf{B})] = \{1\}\}$. Note that $\mathsf{Win}_\mathsf{A} \neq \emptyset$ since $0 \in \mathsf{Win}_\mathsf{A}$.

Now, we let $x := \sup \mathsf{Win}_\mathsf{A}$ and we claim that $\mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{A}] = x = \mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{B}]$. Let $\varepsilon > 0$. There is $u \in \mathsf{Win}_\mathsf{A}$ such that $u \geq x - \varepsilon$. Consider some $a \in \mathsf{Act}_\mathsf{A}$ such that $v_{\geq u}[\varrho(a, \mathsf{Act}_\mathsf{B})] = \{1\}$. Then, we have $\mathsf{val}[\langle \mathcal{F}, v \rangle][a] \geq u$. Indeed, for all Player-B actions $b \in \mathsf{Act}_\mathsf{B}$, we have $v \circ \varrho(a, b) \geq u$. As this holds for all $\varepsilon > 0$, it follows that $\mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{A}] \geq x$ and approaching the value x can be done with deterministic GF-strategies for Player A.

If $x = 1$, we indeed have $\mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{A}] = x = \mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{B}]$. Assume now that $x < 1$. Then, for all $0 < \varepsilon \leq 1 - x$, we have $x + \varepsilon \in [0,1] \setminus \mathsf{Win}_\mathsf{A}$. Therefore, since the game form $\mathcal{F}$ is determined, there is some $b \in \mathsf{Act}_\mathsf{B}$ such that $v_{\geq x+\varepsilon}[\varrho(\mathsf{Act}_\mathsf{A}, b)] = \{0\}$. Then, we have $\mathsf{val}[\langle \mathcal{F}, v \rangle][b] \leq x + \varepsilon$. Indeed, for all Player-A actions $a \in \mathsf{Act}_\mathsf{A}$, we have $v \circ \varrho(a, b) < x + \varepsilon$. As this holds for all $\varepsilon > 0$, it follows that $\mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{B}] \leq x$ and approaching the value x can be done with deterministic GF-strategies for Player B. Hence, $\mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{A}] = x = \mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{B}]$.

Assume now that $\mathcal{F}$ is valuable and supremized w.r.t. Player A by $\mathsf{Act_A}$ and w.r.t. Player B by $\mathsf{Act_B}$. Consider any valuation $v : \mathsf{O} \to \{0,1\}$. Since the game form $\mathcal{F}$ is valuable, we let $x := \mathsf{val}[\langle \mathcal{F}, v \rangle]$. Let us show that $x \in \{0,1\}$. Assume towards a contradiction that it is not the case and let $\varepsilon := \frac{\min(x,1-x)}{2} > 0$. By assumption, Player A (resp. B) has a deterministic GF-strategy $a \in \mathsf{Act_A}$ (resp. $b \in \mathsf{Act_B}$) that is ε-optimal in the game in normal form $\langle \mathcal{F}, v \rangle$. However, we have $v \circ \varrho(a,b) \in \{0,1\}$. Hence, by definition of ε, either a or b is not ε-optimal. Hence the contradiction. In fact, $x \in \{0,1\}$. Assume for instance that $x = 1$ and consider a Player A deterministic GF-strategy $a \in \mathsf{Act_A}$ that is $\frac{1}{2}$-optimal in the game in normal form $\langle \mathcal{F}, v \rangle$. Then, for all $b \in \mathsf{Act_B}$, we have $|v \circ \varrho(a,b) - 1| \leq \frac{1}{2}$. Since $v[\varrho] \subseteq \{0,1\}$, it follows that $v \circ \varrho(a,b) = 1$. In other words, we have $v[\varrho(a, \mathsf{Act_B})] = \{1\}$. Symmetrically, if $x = 0$, we can show that there is some $b \in \mathsf{Act_B}$ such that $v[\varrho(\mathsf{Act_A}, b)] = \{0\}$. □

We can now prove Theorem 6.4.

Proof. This proof is actually somewhat close to the second part of the proof of Proposition 6.5 but adapted to the case of win/lose graph games.

Consider such a game $\mathcal{G}$ and a state $q \in Q$. Since all local interactions are valuable (since they are determined, by Proposition 6.5), by Theorem 2.3, the game $\mathcal{G}$ has a value. Let $x := \chi_{\mathcal{G}}(q)$. Assume towards a contradiction that $x \notin \{0,1\}$. Let $\varepsilon := \frac{\min(x,1-x)}{2} > 0$. By Theorem 2.3, since the local interactions are supremized, w.r.t. both players, by deterministic GF-strategies by Proposition 6.5, Player A (resp. B) has a deterministic GF-strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ (resp. $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$) that is ε-optimal in the game $\mathcal{G}$ from q. Since all local interactions are deterministic, and both strategies $\mathsf{s_A}$ and $\mathsf{s_B}$ are deterministic, it follows that there is a unique path $\rho \in Q^\omega$ such that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[\rho] = 1$. Since the game $\mathcal{G}$ has no stopping states and is win/lose, it follows that we have $\mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[(f_\mathcal{C})^q] \in \{0,1\}$. Hence, by definition of ε, either $\mathsf{s_A}$ (if $\mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[(f_\mathcal{C})^q] = 0$) or $\mathsf{s_B}$ (if $\mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[(f_\mathcal{C})^q] = 1$) is not ε-optimal. Hence the contradiction. In fact, $x \in \{0,1\}$. Assume for instance that $x = 1$ and consider a Player A deterministic strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ that is $\frac{1}{2}$-optimal in the game $\mathcal{G}$ from q. Then, for all deterministic Player-B strategies $\mathsf{s_B} \in \mathsf{Act_B}$, we have $|\mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[(f_\mathcal{C})^q] - 1| \leq \frac{1}{2}$. Since $\mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[(f_\mathcal{C})^q] \in \{0,1\}$, it follows that $\mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[(f_\mathcal{C})^q] = 1$. In other words, for all Player-B deterministic strategies $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$, the only path compatible with $\mathsf{s_A}$ and $\mathsf{s_B}$ has value 1. That is, the Player-A deterministic strategy $\mathsf{s_A}$ is winning. Symmetrically, if $x = 0$, we can show that Player B has a deterministic winning strategy. □

Overall, determined game forms ensure the following.

Theorem 6.6. *Among standard deterministic game forms, being determined is an NSC-transfer for the existence of winning strategies in infinite win/lose games without stopping states.*

Proof. This a consequence of Proposition 6.3 and Theorem 6.4. □

Memory Transfer. In [38] where we have originally considered the use of determined game forms in concurrent games, in addition to proving Theorem 6.4, we have also proved results on the memory required to play winning strategies. We recall what was done in [38] here, without giving the formal proof (which can be found in the arXiv version [72] of [38]).

In [51], the authors proved an equivalence between the shape of a winning objective and the existence of winning strategies that can be implemented with a given memory skeleton $\mathcal{M}$ in turn-based games. They defined the properties of $\mathcal{M}$-selectivity and $\mathcal{M}$-monotony (which we recall in Section 6.6) and proved that for $\mathcal{M}$ a memory skeleton and $W \subseteq \mathsf{K}^\omega$, we have that W and $\mathsf{K}^\omega \setminus W$ are $\mathcal{M}$-monotone and $\mathcal{M}$-selective is equivalent to every deterministic finite-state turn-based game with W as winning objective, winning strategies for both players that can be found among strategies implemented with memory skeleton $\mathcal{M}$ (see Theorem 6.18 in Section 6.6). This also holds in finite-state concurrent games with local interactions that are determined and finite. This is stated Theorem 6.7 below, which amounts to [38, Theorem 18].

Theorem 6.7. *Let K be a non-empty set of colors, $\mathcal{M}$ be a memory skeleton on K and $W \subseteq \mathsf{K}^\omega$ be an objective. The following two assertions are equivalent:*

1. *in every finite-state concurrent game with local interactions that are finite and determined, winning strategies for both players that can be found among strategies implemented with memory skeleton $\mathcal{M}$;*

2. *W and $\mathsf{K}^\omega \setminus W$ are $\mathcal{M}$-monotone and $\mathcal{M}$-selective.*

6.2.1 . Retrieving Theorem 2.1

We want to point out an important fact about that Theorem 6.4.

Informal Statement 6.1. *Borel determinacy, that is the existence of winning strategies, from every state, in all deterministic turn-based games without stopping states (stated in Theorem 2.1) is a logical consequence of Theorem 6.4.*

Proof. This comes from Proposition 6.2: deterministic turn-based game forms are determined. □

Remark 6.1. *Note that Borel determinacy that we have stated in Theorem 2.1 holds even for a set of states uncountable. On the other hand, the games we consider in this dissertation have a countable set of states. However, Theorem 2.3 and Proposition 6.2 would also hold with a set of states that is not countable. (Note that, even if the set of states is not countable, from any starting state, once both players have chosen a strategy, the set of states that can be visited with positive probability is countable. This is because the* *distributions over states that we consider in local interactions always have a countable support.)*

However, it is also very important to note the following. We have used Theorem 2.3 to prove Theorem 6.4 and we have used Theorem 2.1 to prove Theorem 2.3. Therefore, we have **not** given a new proof of Theorem 2.1. The only thing we can say is that Theorem 2.3 is a strengthening of Theorem 2.1.

In addition, note that we have also proved Theorem 6.4 in [38]. However, in that paper we directly used Theorem 2.1 to prove it — since we transferred already existing results in deterministic turn-based games to concurrent games with determined local interactions. Hence, in that case too, we have **not** given a new proof of Theorem 2.1.

6.2.2 . Application: discrete-bidding games

In this subsection, we would like to quickly present an application of determined game forms. More precisely, determined game forms actually appear in concurrent games of the literature. We discuss this on discrete-bidding games. These games were initially introduced in [73]. Here, we consider part of what is done in [74]. We would like to thank Guy Avni (one of the authors of [74]) for the fruitful discussion we had in Highlights 2022 on the subject of this subsection.

The games studied in that paper are finite-state concurrent two-player antagonistic games with specific local interactions. Let us explain exactly how the players interact at each state. The two players start the game with an initial budget (i.e. an integer). Then, at each step of the game, both players bid concurrently some amount of their current budget. The highest bidder pays the other player what she has bid and gets to choose the next state. The process repeats indefinitely, thus creating an infinite sequence of states. The games considered are win/lose and without stopping states.

Contrary to continuous bidding, which is another kind of concurrent games studied in the literature, how discrete-bidding games behave heavily depends on the tie-breaking mechanism used in the game. In [74], the authors present three kinds of tie-breaking mechanisms. All these mechanisms induce different local interactions in the games. Let us first consider deterministic tie-breaking mechanisms, we will discuss briefly stochastic ones at the end of this subsection. There are two deterministic tie-breaking mechanisms studied in [74]:

- Transducer-based: the game is given a transducer that, as a function of the states visited, the winners of the previous ties, the number of ties that have already occurred and past winning bids choose who is the winner of the tie.

- Advantaged based: at the start of the game, one player holds the advantage. Then, whenever there is a tie, the player holding the advantage

$$\mathcal{F}_{\mathsf{bid}}^{3,3} = \begin{bmatrix} v_{0,0} & v_{0,1} & v_{0,2} & v_{0,3} \\ v_{1,0} & v_{1,1} & v_{1,2} & v_{1,3} \\ v_{2,0} & v_{2,1} & v_{2,2} & v_{2,3} \\ v_{3,0} & v_{3,1} & v_{3,2} & v_{3,3} \end{bmatrix}$$

Figure 6.6: A bidding interaction where both players have budget 3.

may decide to either win and give the advantage to the other player or lose and keep the advantage.

Let us first discuss transducer-based tie-breaking. In this setting, we are interested in the existence of winning strategies. As stated in [74, Theorem 4.1], if the transducer has the information of whether or not a tie has occurred, then there are reachability games where neither player has winning strategies. Hence, the authors focus on transducer unaware of ties. In this setting, the authors show in all Muller games (a generalization of parity games), either player has a winning strategy [74, Theorem 4.5].

In fact, in such a setting, we believe that the local interactions involved are determined. This would imply these bidding games inherit all the nice properties that concurrent games with determined local interactions have, see Theorem 6.4 and Theorems 6.7. In particular, in all win/lose games (with Borel winning objectives), one of the players has a winning strategy.

Let us give the intuition of why we believe that the local interactions involved are determined. To do so, consider the bidding interaction $\mathcal{F}_{\mathsf{bid}}$ depicted in Figure 6.6 in the case where both players have budget 3. An outcome $v_{i,j}$ corresponds to a situation where Player A has bid i and Player B has bid j. The diagonal outcomes of the shape $v_{i,i}$ correspond to ties. Furthermore, because of how bidding games are played, some of these outcomes are in fact equal. For instance, $v_{2,0}$ and $v_{2,1}$ are the same since in both cases Player A wins and pays 2 to Player B.

In fact, the bidding interaction has the shape depicted in Figure 6.7 where the outcome $v_{i,\mathsf{A}}$ refers to the fact that Player A has won the bid and has paid 2 to Player B. Let us now consider the ties. This is where we use the unaware of ties assumption. Assume that, given the history of the game, Player A wins the next bid. Then, because the transducer considered is unaware of ties, an outcome $v_{i,i}$ is in fact equal to the outcome $v_{i,\mathsf{A}}$. If the transducer were not unaware of ties, the outcome $v_{i,i}$ could not be compared, a priori, with either $v_{i,\mathsf{A}}$ or $v_{i,\mathsf{B}}$ (this can be seen in the example discussed in [74, Theorem 4.1]). With the unaware of ties assumption, the bidding interaction has the shape depicted in Figure 6.8.

Now, it is actually rather straightforward to show that the game form of

$$\mathcal{F}^{3,3}_{\mathsf{bid}} = \begin{bmatrix} v_{0,0} & v_{1,\mathsf{B}} & v_{2,\mathsf{B}} & v_{3,\mathsf{B}} \\ v_{1,\mathsf{A}} & v_{1,1} & v_{2,\mathsf{B}} & v_{3,\mathsf{B}} \\ v_{2,\mathsf{A}} & v_{2,\mathsf{A}} & v_{2,2} & v_{3,\mathsf{B}} \\ v_{3,\mathsf{A}} & v_{3,\mathsf{A}} & v_{3,\mathsf{A}} & v_{3,3} \end{bmatrix}$$

Figure 6.7: The same bidding interaction where the relations between outcomes are made explicit.

$$\mathcal{F}^{3,3}_{\mathsf{bid,A}} = \begin{bmatrix} v_{0,\mathsf{A}} & v_{1,\mathsf{B}} & v_{2,\mathsf{B}} & v_{3,\mathsf{B}} \\ v_{1,\mathsf{A}} & v_{1,\mathsf{A}} & v_{2,\mathsf{B}} & v_{3,\mathsf{B}} \\ v_{2,\mathsf{A}} & v_{2,\mathsf{A}} & v_{2,\mathsf{A}} & v_{3,\mathsf{B}} \\ v_{3,\mathsf{A}} & v_{3,\mathsf{A}} & v_{3,\mathsf{A}} & v_{3,\mathsf{A}} \end{bmatrix}$$

Figure 6.8: The bidding interaction in the case where Player A wins ties.

Figure 6.8 is determined. Indeed, recall Definition 6.1 and consider a map of the outcomes into $\{0, 1\}$. For the condition of Definition 6.1 not to be met, it must be that:

- $v_{3,\mathsf{A}}$ is mapped to 0, and $v_{3,\mathsf{B}}$ to 1;
- $v_{2,\mathsf{A}}$ is mapped to 0, and $v_{2,\mathsf{B}}$ to 1;
- $v_{1,\mathsf{A}}$ is mapped to 0, and $v_{1,\mathsf{B}}$ to 1;
- Finally, if $v_{0,\mathsf{A}}$ is mapped to 0, the leftmost column is full of 0, and if $v_{0,\mathsf{A}}$ is mapped to 1, the topmost row is mapped to 1.

Hence, the condition of Definition 6.1 is necessarily met. That is, the game form $\mathcal{F}^{3,3}_{\mathsf{bid,A}}$ is determined. It seems that this reasoning can be generalized to any budget of the players, regardless of who wins the current tie. Note that all the observations and reasoning presented in this subsections are given in [74]. However, they are not used to prove the determinacy of the local interactions (since this notion is not defined). Although, it has to be noted that what is proved by the authors (called "local determinacy") is close to the notion of determinacy.

As mentioned above, there is another type of deterministic tie-breaking mechanism: advantage-based tie-breaking. In this setting, the authors prove similar results than for the transducer-based unaware of ties mechanism. It would be interesting to investigate if, also in this case, we could show that the local interactions are determined.

Finally, who wins ties may also be decided randomly. We discuss this further at the end of Section 6.3 below.

6.2.3 . Semi-determined game forms

In this subsection, we consider a weaker notion of determinacy for game forms that is asymmetric in the players, in the sense that we make an assumption for one player but not for the other. If we consider Definition 6.1, it does not seem possible since the resulting definition would be pointless. However, we may use the characterization of Proposition 6.5. since the notion we want to define takes the point of view of only one player, say Player A, it suffices not to make any assumption w.r.t. ε-optimal GF-strategies for Player B. This is what we do in the definition below.

Definition 6.2 (Semi determined game form). *Consider a set of outcomes* O *and a standard deterministic game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$*. It is* semi determined *w.r.t. Player* A *if it is valuable and supremized w.r.t. Player* A *by* Act_A*.*

This is symmetrical for Player B*.*

In fact, for standard deterministic game forms with at least one set of actions which is finite, being determined is equivalent to being semi determined for Player A. This is not the case when both action sets are infinite.

Proposition 6.8. *Consider a set of outcomes* O *and a standard deterministic game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$*. If either* Act_A *or* Act_B *is finite, then* $\mathcal{F}$ *is determined if and only if it is semi determined w.r.t. Player* A *(or Player* B*).*

There is a standard deterministic game form where both players have an infinite action set that is semi determined w.r.t. Player A *but that is not determined.*

Proof. Consider a game form $\mathcal{F}$ and assume that it is semi determined w.r.t. Player A. We do not make any other assumption on $\mathcal{F}$ for now. Consider any valuation $v : \mathsf{O} \to \{0,1\}$. Since the game form $\mathcal{F}$ is valuable, we may consider $x := \mathsf{val}[\langle \mathcal{F}, v \rangle] \in [0,1]$. Assume towards a contradiction that $x \notin \{0,1\}$. Let $0 < \varepsilon < x$. Consider any Player-A deterministic GF-strategy $a \in \mathsf{Act}_\mathsf{A}$ that is ε-optimal in the game in normal form $\langle \mathcal{F}, v \rangle$. Then, for all $b \in \mathsf{Act}_\mathsf{B}$, it must be that $v \circ \varrho(a,b) \geq x - \varepsilon > 0$. That is, $v[\varrho(a, \mathsf{Act}_\mathsf{B})] = \{1\}$. Hence the contradiction. In fact, it must be that $x \in \{0,1\}$.

Assume that $x = 0$. Let us now make additional assumptions on $\mathcal{F}$. Assume first that the set Act_A is finite. Assume towards a contradiction that, for all $b \in \mathsf{Act}_\mathsf{B}$, there is some $a \in \mathsf{Act}_\mathsf{A}$ such that $v \circ \varrho(a,b) = 1$. Then, consider a Player-A GF-strategy $\sigma_\mathsf{A} \in \mathcal{D}(\mathsf{Act}_\mathsf{A})$ that plays all actions in Act_A with uniform probability: for all $a \in \mathsf{Act}_\mathsf{A}$, $\sigma_\mathsf{A}(a) := \frac{1}{|\mathsf{Act}_\mathsf{A}|}$. Then, for all $\sigma_\mathsf{B} \in \mathcal{D}(\mathsf{Act}_\mathsf{B})$, we have $\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) \geq \frac{1}{|\mathsf{Act}_\mathsf{A}|}$. Therefore, we would have $x \geq \frac{1}{|\mathsf{Act}_\mathsf{A}|} > 0$, which is a contradiction. In fact, there is some $b \in \mathsf{Act}_\mathsf{B}$ such that $v[\varrho(\mathsf{Act}_\mathsf{A}, b)] = \{0\}$.

Assume now that the set Act_B is finite. Assume again towards a contradiction that, for all $b \in \mathsf{Act}_\mathsf{B}$, there is some $a_b \in \mathsf{Act}_\mathsf{A}$ such that $v \circ \varrho(a_b, b) = 1$.

$$\begin{bmatrix} y & y & y & y & \dots \\ x & y & y & y & \dots \\ y & x & y & y & \dots \\ y & y & x & y & \dots \\ y & y & y & x & \dots \\ \dots & \dots & \dots & \dots & \dots \end{bmatrix}$$

Figure 6.9: A game form semi determined w.r.t. Player A that is not determined.

Then, consider a Player-A GF-strategy $\sigma_A \in \mathcal{D}(\mathsf{Act}_A)$ that plays all actions in $X_B := \{a_b \mid b \in \mathsf{Act}_B\}$ with uniform probability: for all $b \in X_B$, $\sigma_A(a_b) := \frac{1}{|X_B|}$. Then, for all $\sigma_B \in \mathcal{D}(\mathsf{Act}_B)$, we have $\mathsf{out}[\langle \mathcal{F}, v\rangle](\sigma_A, \sigma_B) \geq \frac{1}{|X_B|}$. Therefore, we would have $x \geq \frac{1}{|X_B|} > 0$, which is a contradiction. In fact, there is some $b \in \mathsf{Act}_B$ such that $v[\varrho(\mathsf{Act}_A, b)] = \{0\}$.

Hence, if either Act_A or Act_B is finite, then we can conclude that the game form $\mathcal{F}$ is determined.

Consider now the game form $\mathcal{F}$ of Figure 6.9. We claim that it is semi determined w.r.t. Player A. Consider any valuation $v : \{x, y\} \to [0, 1]$. If $v(y) \geq v(x)$, we have $\mathsf{val}[\langle \mathcal{F}, v\rangle] = v(y)$ and playing deterministically the top row is optimal for Player A. Assume that $v(y) < v(x)$. Then, we still have $\mathsf{val}[\langle \mathcal{F}, v\rangle] = v(y)$. Indeed, for all $n \in \mathbb{N}$, a Player-B GF-strategy $\sigma_B \in \mathcal{D}(\mathsf{Act}_B)$ that plays uniformly over the first n columns has value $\frac{v(x)+(n-1)\cdot v(y)}{n} \to_{n\to\infty} v(y)$. In that case, any Player-A GF-strategy is optimal. Hence, the game form $\mathcal{F}$ is semi determined w.r.t. Player A. However, it is not determined. Indeed, for $v : \{x, y\} \to \{0, 1\}$ such that $v(x) := 1$ and $v(y) := 0$, there is clearly no rows of 1 nor any column of 0. □

Let us now consider what happens when we use semi determined game forms w.r.t. Player A as local interactions in win/lose concurrent games without stopping states. It still holds that the value of any state is 0 or 1. Furthermore Player A has still winning strategies from every state of value 1. However, Player B does not necessarily have some from states of value 0[1].

Proposition 6.9. *Consider a concurrent win/lose game $\mathcal{G}$ without stopping states such that, for all $q \in Q$, the local interaction $\mathsf{F}(q)$ is semi determined w.r.t. Player A. Then, from every state $q \in Q$, we have $\chi_{\mathcal{G}}(q) \in \{0, 1\}$ and if $\chi_{\mathcal{G}}(q) = 1$, then Player A has a winning strategy from q.*

[1]This is witnessed by the game form of Figure 6.9: if it placed in a game at a state q where reaching x is winning for Player A and reaching y is winning for Player B, then the state q has value 0, but Player B has no winning strategy.

Proof. This proof is almost identical to the proof of Theorem 6.4. Consider such a game $\mathcal{G}$ and a state $q \in Q$. Since all local interactions are valuable, by Theorem 2.3, the game $\mathcal{G}$ has a value. Let $x := \chi_{\mathcal{G}}(q)$. Assume towards a contradiction that $x \notin \{0, 1\}$. Let $0 < \varepsilon < x$. By Theorem 2.3, since the local interactions are supremized, w.r.t. Player A, by deterministic GF-strategies, Player A has a deterministic GF-strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ that is ε-optimal in the game $\mathcal{G}$ from q. Consider any deterministic Player-B strategy $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$. Since all local interactions are deterministic, and both strategies $\mathsf{s_A}$ and $\mathsf{s_B}$ are deterministic, it follows that there is a unique path $\rho \in Q^\omega$ such that $\mathbb{P}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[\rho] = 1$. Since the game $\mathcal{G}$ has no stopping states and is win/lose, it follows that we have $\mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[(f_\mathcal{C})^q] \in \{0, 1\}$. As this holds for all Player-B deterministic strategies, there is a contradiction. Indeed, either for all Player-B deterministic strategies $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$, we have $\mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[(f_\mathcal{C})^q] = 1$, and in that case the Player-A strategy $\mathsf{s_A}$ has value 1. Or, there is some Player-B deterministic strategy $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$ such that we have $\mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[(f_\mathcal{C})^q] = 0$, and in that case the Player-A strategy $\mathsf{s_A}$ has value 0. In fact, $x \in \{0, 1\}$. Assume now that $x = 1$ and consider a Player-A deterministic strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ that is $\frac{1}{2}$-optimal in the game $\mathcal{G}$ from q. Then, for all deterministic Player-B strategies $\mathsf{s_B} \in \mathsf{Act_B}$, we have $|\mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[(f_\mathcal{C})^q] - 1| \leq \frac{1}{2}$. Since $\mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[(f_\mathcal{C})^q] \in \{0, 1\}$, it follows that $\mathbb{E}_{\mathcal{C},q}^{\mathsf{s_A},\mathsf{s_B}}[(f_\mathcal{C})^q] = 1$. In other words, for all Player-B deterministic strategies $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$, the only path compatible with $\mathsf{s_A}$ and $\mathsf{s_B}$ has value 1. That is, the Player-A deterministic strategy $\mathsf{s_A}$ is winning. □

6.3 Finitely-maximizable game forms

In this section, we consider another notion on standard finite game forms — not necessarily on deterministic ones — that is weaker than determinacy on standard finite game forms. Hence, on standard finite game forms, this can be seen as a generalization of the notion of determinacy. Note that, however, contrary to what we did in Theorem 6.6, we will not state an NSC-transfer in this section.

To gain an intuition behind this notion, let us consider again the characterization of determined game forms via Proposition 6.5. In a standard finite game form $\mathcal{F}$ — which is necessarily valuable, recall Theorem 1.11 — this characterization, from Player A's point of view, amounts to: the set $\mathsf{Act_A}$ supremizes the game form $\mathcal{F}$. However, the set $\mathsf{Act_A}$, besides being the set of Player-A deterministic GF-strategies, can simply be seen as a finite set of Player-A GF-strategies. However, in this section, we are not particularly interested in deterministic strategies — which was the case in the previous section since we considered winning strategies. Hence, there is no reason to limit ourselves to that specific finite set of Player-A GF-strategies. This suggests the definition below of finitely maximizable game forms for one player.

Definition 6.3 (Finitely maximizable game forms). *Consider a set of outcomes* O *and a standard finite game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$. *It is* finitely maximizable w.r.t. Player A *if there is finite set* $S_\mathsf{A} \subseteq \mathcal{D}(\mathsf{Act}_\mathsf{A})$ *of Player-*A GF*-strategies that maximizes it.*

The definition is analogous for Player B.

Remark 6.2. *Consider a game form that is finitely maximizable w.r.t. Player* A. *This game form is valuable by Theorem 1.11. Furthermore, since the set* $S_\mathsf{A} \subseteq \mathcal{D}(\mathsf{Act}_\mathsf{A})$ *is finite, it maximizes the game form* $\mathcal{F}$ *if and only if it supremizes it, recall Observation 1.1.*

We will give below two applications of finitely maximizable game forms in finite games, but before that we state that any standard game forms with at most two outcomes is finitely maximizable. It is also the case of determined game forms.

Proposition 6.10. *Consider a set of outcomes* O *and a standard finite game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$. *If* $|\mathsf{O}| \leq 2$ *or if the game form* $\mathcal{F}$ *is determined (it may be turn-based), it is finitely maximizable w.r.t. Player* A *and Player* B.

Proof. If $\mathcal{F}$ is determined, this is a direct consequence of Proposition 6.5.

Assume now that $|\mathsf{O}| \leq 2$. Let us show that it is finitely maximizable w.r.t. Player A, this is similar for Player B. If $|\mathsf{O}| = 1$, this is obvious since any Player-A GF-strategy is optimal in games in normal form that can be obtained from $\mathcal{F}$. Assume now that $|\mathsf{O}| = 2$. We write $\mathsf{O} = \{x, y\}$. We let $v_x : \{x, y\} \to [0, 1]$ be such that $v_x(x) := 1$ and $v_x(y) := 0$ and symmetrically, we let $v_y : \{x, y\} \to [0, 1]$ be such that $v_y(x) := 0$ and $v_y(y) := 1$. Let $\sigma_\mathsf{A}^x \in \mathcal{D}(\mathsf{Act}_\mathsf{A})$ (resp. $\sigma_\mathsf{A}^y \in \mathcal{D}(\mathsf{Act}_\mathsf{A})$) be a Player-$\mathsf{A}$ GF-strategy that is optimal in the game in normal form $\langle \mathcal{F}, v_x \rangle$ (resp. $\langle \mathcal{F}, v_y \rangle$). We claim that the set $\{\sigma_\mathsf{A}^x, \sigma_\mathsf{A}^y\}$ maximizes the game form $\mathcal{F}$. Consider any valuation $v : \{x, y\} \to [0, 1]$. If $v(x) = v(y)$, any Player-A GF-strategy is optimal in the game in normal form $\langle \mathcal{F}, v \rangle$. Assume now that $v(x) > v(y)$. We let $w := v - v(y) : \{x, y\} \to [0, 1]$. We have:

- $v = v(x) \cdot v_x + v(y) \cdot v_y$
- $w = (v(x) - v(y)) \cdot v_x$

Therefore, by Lemma 1.10 and since σ_A^x is optimal in $\langle \mathcal{F}, v_x \rangle$:

$$\begin{aligned}
\mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_\mathsf{A}^x) &= \mathsf{val}[\langle \mathcal{F}, w \rangle](\sigma_\mathsf{A}^x) - v(y) \\
&= (v(x) - v(y)) \cdot \mathsf{val}[\langle \mathcal{F}, v_x \rangle](\sigma_\mathsf{A}^x) - v(y) \\
&= (v(x) - v(y)) \cdot \mathsf{val}[\langle \mathcal{F}, v_x \rangle] - v(y) \\
&= \mathsf{val}[\langle \mathcal{F}, w \rangle] - v(y) \\
&= \mathsf{val}[\langle \mathcal{F}, v \rangle]
\end{aligned}$$

That is, the Player-A GF-strategy σ_{A}^x is optimal in the game in normal form $\langle \mathcal{F}, v \rangle$. Similarly, if $v(x) < v(y)$, we would have the Player-A GF-strategy σ_{A}^y is optimal in the game in normal form $\langle \mathcal{F}, v \rangle$. In fact, finite set $\{\sigma_{\mathsf{A}}^x, \sigma_{\mathsf{A}}^y\} \subseteq \mathcal{D}(\mathsf{Act}_{\mathsf{A}})$ maximizes the game form $\mathcal{F}$ for Player A. □

With finitely maximizable local interactions in finite concurrent games, we can make use of the results of (the first subsection of) Section 3.4 — regarding the sequentialization and parallelization of strategies — and of Theorem 2.3. That is, consider a finite concurrent game $\mathcal{G}$ where all local interactions are finitely maximizable w.r.t. Player A. In particular, this game is B-finite, recall Definition 3.25. The informal idea is the following: consider Definition 3.18 and the turn-based game $\mathcal{G}(\Lambda, \eta)$ that is the sequentialized version of the concurrent game $\mathcal{G}$ where, for all $q \in Q$, we have $\Lambda_q := S_{\mathsf{A}}^q$ for $S_{\mathsf{A}}^q \subseteq \mathcal{D}(\mathsf{Act}_{\mathsf{A}}^q)$ a finite set supremizing the game form $\mathsf{F}(q)$ (the function η is defined arbitrarily for now). Since S_{A}^q is finite, the turn-based game $\mathcal{G}(\Lambda, \eta)$ is finite. By using Theorem 2.3 and Proposition 3.32, we can then show that, for all $q \in Q$, we have $\chi_{\mathcal{G}}(q) = \chi_{\mathcal{G}(\Lambda,\eta)}(q)$. This means informally that what happens in the turn-based game $\mathcal{G}(\Lambda, \eta)$ is the same as what happens in the game $\mathcal{G}$, from every state. Stated differently, the finite concurrent game $\mathcal{G}$ behaves like a finite turn-based games, and therefore enjoys (some of) the nice properties that finite turn-based games enjoy. We present two such properties below.

First, just like in finite turn-based games, there are always subgame optimal strategies in finite concurrent games with local interactions which are finitely maximizable w.r.t. Player A when the payoff function is PI upward well-founded.

Theorem 6.11. *Consider a finite concurrent game $\mathcal{G}$ with a PI upward well founded payoff function where all local interactions are finitely maximizable w.r.t. Player A. For all $q \in Q$, we let $S_{\mathsf{A}}^q \subseteq \mathcal{D}(\mathsf{Act}_{\mathsf{A}}^q)$ be a finite set of Player-A GF-strategies supremizing the game form $\mathsf{F}(q)$. Then, Player A has a subgame optimal strategy in $\mathcal{G}$ generated by $(S_{\mathsf{A}}^q)_{q \in Q}$.*

Remark 6.3. *Before proving this theorem, we want to make two quick remarks. First, this result generalizes Corollary 3.25, though we use this corollary to prove this theorem. Second, since for all $q \in Q$, the set S_{A}^q is finite, this theorem shows the existence of a finite-choice subgame optimal strategy in $\mathcal{G}$ (recall Definition 3.22).*

Proof. Consider a fresh color $k \notin \mathsf{K}$ and let $\mathsf{K}' := \{k\}$ and $\eta : \mathsf{K} \to \{k\}$. For all $q \in Q$, we let $\Lambda_q := S_{\mathsf{A}}^q$ and $\Lambda := (\Lambda_q)_{q \in Q}$. Consider now the turn-based game $\mathcal{G}(\Lambda, \eta)$ from Definition 3.18. Since the payoff function of the game $\mathcal{G}$ is PI upward well-founded, so is the payoff function of the turn-based game $\mathcal{G}(\Lambda, \eta)$. Hence, by Corollary 3.25, Player A has a subgame optimal deterministic strategy $\mathsf{s}_{\mathsf{A}} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}(\Lambda,\eta)}$ in the game $\mathcal{G}(\Lambda, \eta)$.

Furthermore, by Proposition 3.32, for all $q \in Q$, we have $\chi_{\mathcal{G}(\Lambda,\eta)}(q) = \sup_{\mathsf{t_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda)} \chi_\mathcal{G}[\mathsf{t_A}](q)$ where $\mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda)$ refers to the set of Player-A strategies generated by Λ. By definition of Λ and Theorem 2.3, we have $\sup_{\mathsf{t_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}(\Lambda)} \chi_\mathcal{G}[\mathsf{t_A}](q) = \chi_\mathcal{G}(q)$. Hence, for all $q \in Q$, we have $\chi_{\mathcal{G}(\Lambda,\eta)}(q) = \chi_\mathcal{G}(q)$. In addition, Proposition 3.32 also gives that the Player-A strategy $\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}) \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ (from Definition 3.19) is generated by Λ and ensures, for all $q \in Q$, that $\chi_{\mathcal{G}(\Lambda,\eta)}(q) = \chi_{\mathcal{G}(\Lambda,\eta)}(\mathsf{s_A})(q) \leq \chi_\mathcal{G}(\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A}))(q)$. Therefore, the strategy $\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A})$ is optimal in $\mathcal{G}$. We can now conclude by applying Theorem 3.28: since $\mathsf{Pr}_\mathsf{A}^\Lambda(\mathsf{s_A})$ is generated by Λ and is positively bounded (since for all $q \in Q$, Λ_q is finite), there is a Player-A strategy in the game $\mathcal{G}$ that is subgame optimal and generated by Λ. □

Second, let us consider the memory necessary and sufficient to be (subgame) optimal in such concurrent games. In [38], when dealing with determined game forms, in addition to proving Theorem 6.4 (from the previous subsection), we also established memory transfer from turn-based games to concurrent games with determined local interactions (see [38, Theorem 12, Corollary 16]). Since this is the main focus in this dissertation, we will state a similar theorem only for parity objectives. However, note that we could have similar statements for other objectives, well-behaved in turn-based games.

Corollary 6.12. *Consider a finite concurrent parity game $\mathcal{G}$ where all local interactions are finitely maximizable w.r.t. Player* A*. Then, Player* A *has a positional strategy that is (subgame) optimal in $\mathcal{G}$.*

Proof. By Theorem 6.11, there is a finite-choice subgame optimal strategy in $\mathcal{G}$. Corollary 3.38 then gives that there is positional subgame optimal strategy in $\mathcal{G}$. □

Discrete-bidding games. We have discussed earlier in Section 6.2 the fact that determined game forms appear in the literature. We discussed it by using part of what is done in [74]. We have not considered yet the case of random-based tie-breaking mechanism since determined game forms are, by definition, deterministic. In this setting, the authors of [74] show that all reachability (finite-state) games have a value when restricted to deterministic strategies. (More precisely, they show that the games they consider have a value, and in the games they consider, the players can only play deterministic strategies.)

In fact, we make the conjecture below about the local interactions occurring in this setting.

Conjecture 6.13. *The local interactions occurring in discrete-bidding games with random-based tie-breaking mechanisms are supremized by deterministic* GF*-strategies. (Since these game forms are standard finite, it would follow that these game forms are finitely maximizable.)*

$$\begin{bmatrix} x & y & z & t \\ t & x & y & z \\ z & t & x & y \\ y & z & t & x \end{bmatrix}$$

Figure 6.10: A circular game form ($n = 4$).

$$\begin{bmatrix} x & y \\ y & x \end{bmatrix}$$

Figure 6.11: A circular game form ($n = 2$).

$$\begin{bmatrix} x & y & u \\ z & t & u \\ u & u & u \end{bmatrix}$$

Figure 6.12: Another uniquely maximizable game form.

This conjectures implies that all games using random-based tie-breaking mechanisms have a value when restricted to deterministic strategies (regardless of the Borel objectives or measurable payoff function considered).

6.4 Uniquely maximizable game forms

We can strengthen the notion of finite maximizability of local interactions to ensure very strong properties on concurrent games, though it will apply to (much) fewer games. For a game form to be finitely maximizable, we require that there is a finite set that maximizes it. A natural way to strengthen this property is to require that this set is not only finite but a singleton. This defines uniquely maximizable game forms. Note that this is defined in arbitrary game forms, not only standard ones. Furthermore, as in the previous section, we will not state any NSC-transfer in this section.

Definition 6.4 (Uniquely maximizable game forms). *Consider a set of outcomes* O *and an arbitrary game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$. *It is* uniquely maximizable w.r.t. Player A *if there is a Player-*A GF*-strategy* $\sigma_\mathsf{A} \in \mathcal{D}(\mathsf{Act}_\mathsf{A})$ *such that the singleton* $\{\sigma_\mathsf{A}\}$ *maximizes the game form* $\mathcal{F}$.

The definition is analogous for Player B.

Being uniquely maximizable is a strong property on game forms. However, we present below a class (namely, circular game forms) of standard finite game forms that are uniquely maximizable w.r.t. both players.

Definition 6.5 (Circular game forms). *Consider a finite set of outcomes* O. *Let us denote* O *by* $\mathsf{O} = \{o_0, o_1, \ldots, o_{n-1}\}$ *for* $n := |\mathsf{O}| \in \mathbb{N}$. *A game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ *is* circular *if* $\mathsf{Act}_\mathsf{A} = \mathsf{Act}_\mathsf{B} := [\![0, n-1]\!]$ *and for all* $(i, j) \in \mathsf{Act}_\mathsf{A} \times \mathsf{Act}_\mathsf{B}$, *we have:* $\varrho(i, j) := o_{j-i \mod n}$.

In Figure 6.10, we have depicted a circular game form for $n = 4$, and in Figure 6.11, we have depicted a circular game form for $n = 2$. Note that this game form is also known as the matching pennies interaction. In fact, all circular game forms are uniquely maximizable w.r.t. both players.

Proposition 6.14. *Consider a finite set of outcomes* O *and a game form* $\mathcal{F} \in \mathsf{Form}[\mathsf{O}]$. *If* $\mathcal{F}$ *is circular, for both players, a* GF*-strategy* σ *that plays uniformly at random all actions is such that the set* $\{\sigma\}$ *maximizes the game form* $\mathcal{F}$. *Hence, the game form* $\mathcal{F}$ *is uniquely maximizable w.r.t. both players.*

Proof. Consider $v : \mathsf{O} \to [0,1]$. Let $n := |\mathsf{O}|$ and $u := \frac{\sum_{i=0}^{n-1} v(o_i)}{n} \in [0,1]$. We claim that $u = \mathsf{val}[\langle \mathcal{F}, v \rangle]$.

Let $\sigma_{\mathsf{A}} \in \mathcal{D}(\mathsf{Act}_{\mathsf{A}})$ be the Player-A GF-strategy that plays uniformly over all actions. Consider any Player-B action $j \in [\![0, n-1]\!]$. We have:

$$\mathsf{out}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}, v) = \sum_{i=0}^{n-1} \sigma_{\mathsf{A}}(i) \cdot v(o_{j-i \mod n}) = \frac{\sum_{i=0}^{n-1} v(o_i)}{n} = u$$

Therefore, $u = \mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{A}}) \leq \mathsf{val}[\langle \mathcal{F}, v \rangle]$. Symmetrically, denoting $\sigma_{\mathsf{B}} \in \mathcal{D}(\mathsf{Act}_{\mathsf{B}})$ the Player-B GF-strategy that plays uniformly over all actions, we have $u = \mathsf{val}[\langle \mathcal{F}, v \rangle](\sigma_{\mathsf{B}}) \geq \mathsf{val}[\langle \mathcal{F}, v \rangle]$. Therefore, $u = \mathsf{val}[\langle \mathcal{F}, v \rangle]$ and σ_{A} and σ_{B} are optimal GF-strategies in the game in normal form $\langle \mathcal{F}, v \rangle$. □

We want to mention that not all uniquely maximizable game forms are circular. This is for instance the case of the game form of Figure 6.12.

When all game forms are uniquely maximizable w.r.t. Player A in an arena, it becomes very easy for Player A to play optimally. Indeed, she has a positional strategy that is subgame optimal regardless of the payoff function considered.

Theorem 6.15. *Consider an arbitrary concurrent arena* $\mathcal{C}$ *(that need not be finite) and assume that all local interactions are uniquely maximizable w.r.t. Player* A. *Then, there is a Player-*A *positional strategy* $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ *such that, for all payoff functions* $f : \mathsf{K}^\omega \to [0,1]$, *the strategy* $\mathsf{s_A}$ *is subgame optimal in the game* $\langle \mathcal{C}, f \rangle$.

Proof. For all $q \in Q$, we let $\sigma_{\mathsf{A}}^q \in \Sigma_{\mathsf{A}}^q$ be a Player-A GF-strategy such that the set $\{\sigma_{\mathsf{A}}^q\}$ maximizes the game form $\mathsf{F}(q)$. Let $\mathsf{s_A} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ be a positional Player-A strategy such that, for all $q \in Q$, we have $\mathsf{s_A}(q) := \sigma_{\mathsf{A}}^q$.

Consider any payoff function $f : \mathsf{K}^\omega \to [0,1]$. Let us first show that the strategy $\mathsf{s_A}$ is optimal in the game $\mathcal{G} = \langle \mathcal{C}, f \rangle$. First, note that the strategy $\mathsf{s_A}$ is the only Player-A strategy generated by $(\{\sigma_{\mathsf{A}}^q\})_{q \in Q}$. Furthermore, by Theorem 2.3, for all $\varepsilon > 0$, for all $q \in Q$, there is a Player-A strategy generated by $(\{\sigma_{\mathsf{A}}^q\})_{q \in Q}$ that is ε-optimal from q in $\mathcal{G}$. In other words, for all $\varepsilon > 0$, the strategy $\mathsf{s_A}$ is ε-optimal from all states $q \in Q$. That is, the strategy $\mathsf{s_A}$ is optimal in the game $\mathcal{G}$. This holds for all payoff functions $f : \mathsf{K}^\omega \to [0,1]$.

Consider now any payoff function $f : \mathsf{K}^\omega \to [0,1]$ and let us show that the strategy $\mathsf{s_A}$ is subgame optimal in the game $\mathcal{G} = \langle \mathcal{C}, f \rangle$. Let $\rho \in (Q_{\mathsf{ns}})^+$. We let $\gamma := \mathsf{col}^*(\mathsf{tl}(\rho)) \in \mathsf{K}^*$. Then, the strategy $\mathsf{s_A}$ is optimal from the state ρ_{lt} in the game $\mathcal{G} = \langle \mathcal{C}, f^\gamma \rangle$. Since this holds for all $\rho \in (Q_{\mathsf{ns}})^+$, it follows that the strategy $\mathsf{s_A}$ is subgame optimal in $\mathcal{G} = \langle \mathcal{C}, f \rangle$. □

Another benefit of uniquely maximizable game forms is that they behave well in games, even if not all local interactions are uniquely maximizable. Indeed, whenever there are uniquely maximizable local interactions in an arena, Player A should always play, positionally in the corresponding states, a GF-strategy supremizing the game form. Let us define formally this change of strategy.

Definition 6.6. *Consider a concurrent arena $\mathcal{C}$. Let $S \subseteq Q$ be any set of states. Consider any Player-A strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$. We say that another Player-A strategy s'_A S-trivialize $\mathsf{s_A}$ if:*

- *for all $q \in Q \setminus S$ and $\rho \in Q^+$, we have $\mathsf{s}'_\mathsf{A}(\rho \cdot q) = \mathsf{s_A}(\rho \cdot q)$;*
- *for all $q \in S$, there is $\sigma_\mathsf{A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ such that the set $\{\sigma_\mathsf{A}\}$ maximizes the game form $\mathsf{F}(q)$. and for all $\rho \in Q^+$, we have: $\mathsf{s}'_\mathsf{A}(\rho \cdot q) = \sigma_\mathsf{A}$.*

We have the proposition below.

Proposition 6.16. *Consider a concurrent arena $\mathcal{C}$ and a subset $S \subseteq Q$ of states. Assume that, for all $q \in S$, the game form $\mathsf{F}(q)$ is uniquely maximizable w.r.t. Player A. Let $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ be any Player-A strategy in the arena $\mathcal{C}$. Then, there are Player-A strategies S-trivializing the strategy $\mathsf{s_A}$. For all such Player-A strategies $\mathsf{s}'_\mathsf{A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$, for all payoff functions $f : \mathsf{K}^\omega \to [0, 1]$ and for all finite paths $\rho \in (Q_\mathsf{ns})^+$, we have $\chi_{\langle \mathcal{C}, f \rangle}[\mathsf{s_A}](\rho) \leq \chi_{\langle \mathcal{C}, f \rangle}[\mathsf{s}'_\mathsf{A}](\rho)$ (recall Definition 3.3).*

We only give an informal proof of this statement, as formalizing it properly would be somewhat lengthy.

Proof. Consider a payoff function $f : \mathsf{K}^\omega \to [0, 1]$, a finite path $\rho \in (Q_\mathsf{ns})^+$ and a Player-B strategy $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$. We let $\gamma := \mathsf{col}^*(\mathsf{tl}(\rho)) \in \mathsf{K}^*$. We build an arena $\mathcal{C}^\mathsf{Unfold}$ such that $\mathcal{C}^\mathsf{Unfold} := \langle Q', \mathsf{F}', \mathsf{K}, \mathsf{col}' \rangle$ with:

- $Q' := \rho \cdot Q^*$ with $Q'_\mathsf{ns} := \{\pi \in Q' \mid \pi_\mathsf{lt} \in Q_\mathsf{ns}\}$;
- for all $\pi \in Q'_\mathsf{ns}$:

$$\mathsf{F}'(\pi) := \begin{cases} \langle \Sigma_\mathsf{A}^q, \{\mathsf{s_B}(\pi)\}, Q, \varrho^{\pi_\mathsf{lt}} \rangle & \text{if } \pi_\mathsf{lt} \in S \\ \langle \{\mathsf{s_A}(\pi)\}, \{\mathsf{s_A}(\pi)\}, Q, \varrho^{\pi_\mathsf{lt}} \rangle & \text{otherwise} \end{cases}$$

- for all $\pi \in Q'_\mathsf{ns}$, $\mathsf{col}'(\pi) := \mathsf{col}(\pi_\mathsf{lt})$.

We can make several observations about the arena $\mathcal{C}^\mathsf{Unfold}$. First, this arena is in fact an MDP where Player A plays alone. We denote by $\mathsf{t_B}$ the only Player-B strategy in that arena. Second, for the Player-A positional strategy $\mathsf{t_A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}^\mathsf{unfold}}$ (resp. $\mathsf{t}'_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}^\mathsf{unfold}}$) such that for all $\pi \in Q'$, we have $\mathsf{t_A}(\pi) := \mathsf{s_A}(\pi)$ (resp. $\mathsf{t}'_\mathsf{A}(\pi) := \mathsf{s}'_\mathsf{A}(\pi)$), we have[2]:

$$\mathbb{E}_{\mathcal{C},\rho}^{\mathsf{s_A},\mathsf{s_B}}[f^\gamma] = \mathbb{E}_{\mathcal{C}^\mathsf{unfold},\rho}^{\mathsf{t_A},\mathsf{t_B}}[f^\gamma]$$

[2]This is where the proof is informal, since this would require a formal proof, for instance by applying Lemma 1.2.

and

$$\mathbb{E}^{\mathsf{s}'_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}[f^\gamma] = \mathbb{E}^{\mathsf{t}'_\mathsf{A},\mathsf{t}_\mathsf{B}}_{\mathcal{C}^{\mathsf{unfold}},\rho}[f^\gamma]$$

Third, all the local interactions in the arena $\mathcal{C}^{\mathsf{Unfold}}$ are uniquely maximizable. Hence, by Theorem 6.15, the Player-A positional strategy $\mathsf{t}'_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}^{\mathsf{unfold}}}$ is subgame optimal in the game $\langle \mathcal{C}^{\mathsf{unfold}}, f \rangle$. In particular, it is optimal from ρ. Hence, $\mathbb{E}^{\mathsf{t}_\mathsf{A},\mathsf{t}_\mathsf{B}}_{\mathcal{C}^{\mathsf{unfold}},\rho}[f^\gamma] \leq \mathbb{E}^{\mathsf{t}'_\mathsf{A},\mathsf{t}_\mathsf{B}}_{\mathcal{C}^{\mathsf{unfold}},\rho}[f^\gamma]$. Therefore, $\mathbb{E}^{\mathsf{s}_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}[f^\gamma] \leq \mathbb{E}^{\mathsf{s}'_\mathsf{A},\mathsf{s}_\mathsf{B}}_{\mathcal{C},\rho}[f^\gamma]$. Since this holds for all Player-B strategies $\mathsf{s}_\mathsf{B} \in \mathsf{S}_\mathsf{B}^{\mathcal{C}}$, it follows that $\chi_{\langle \mathcal{C}, f\rangle}[\mathsf{s}_\mathsf{A}](\rho) \leq \chi_{\langle \mathcal{C}, f\rangle}[\mathsf{s}'_\mathsf{A}](\rho)$. □

6.5 Discussion, open questions and future work

In this chapter, we have given several applications of item $(1.a)$ of Theorem 2.3. That is, we have provided restrictions on game forms such that, when all local interactions of a game satisfy these restrictions, the game satisfies desirable properties.

In Subsection 6.2.2, we have discussed the use of the notion of determined game forms in discrete-bedding games. We have also briefly discussed discrete-bidding games in Section 6.3 with random tie-breaking mechanisms. It would be interesting to formally prove both what we have informally explained in Subsection 6.2.2 (along with exploring the case of advantage-based tie-breaking) and Conjecture 6.13 we made in Section 6.3.

Furthermore, we have not stated an NSC-transfer (we have only stated a sufficiency result) for game forms that are semi-determined, finitely maximizable and uniquely maximizable. We believe that it should not be too difficult to retrieve some kind of NSC-transfer for semi-determined game forms and uniquely maximizable game forms. However, the case of finitely maximizable game forms (for Player A) seems more complicated. That is, we do not know if being finitely maximizable w.r.t. Player A is an NSC-transfer for either of the two applications that we presented in this chapter. This is stated as an open question below.

Open Question 6.1. *Does Theorem 6.11 and/or Corollary 6.12 still hold if the finitely maximizable assumption is weakened?*

Finally, we would like to mention that we believe that Proposition 6.5 stated in this chapter, that gives an alternate definition of determined game forms, can be stated at a graph game level. This constitutes the conjecture below.

Conjecture 6.17. *Consider a standard deterministic concurrent arena $\mathcal{C}$ without stopping states. Then, for all states $q \in Q$, the two propositions below are equivalent:*

- *for all Borel sets $W \in \mathsf{Borel}(Q)$, in the win/lose game $\mathcal{G} = \langle \mathcal{C}, W \rangle$, either* *of the players has a (deterministic) winning strategy from q;*

- *for all payoff functions $f : Q^\omega \to [0,1]$, the game $\mathcal{G} = \langle \mathcal{C}, f \rangle$ has a value from q. Furthermore, from q, for all $\varepsilon > 0$, ε-optimal strategies can be found among deterministic strategies and at least one of the players has an optimal strategy.*

6.6 Appendix

We recall here the condition for the existence of finite-memory strategy in turn-based games established in [51]. This comes from the Appendix of [72].

First, in [51], the authors do not consider a winning objective $W \subseteq \mathsf{K}^\omega$ but rather a preference relation $\preceq: \mathsf{K}^\omega \times \mathsf{K}^\omega$ for Player A and the antagonistic preference $\preceq^{-1}$ for Player B. Hence, we need to translate a winning objective into a preference relation. For $W \subseteq \mathsf{K}^\omega$, we consider the preference relation $\preceq_W: \mathsf{K}^\omega \times \mathsf{K}^\omega$ defined by $\rho \prec_W \rho'$ for all $\rho \notin W$ and $\rho' \in W$.

Let us now focus on the condition stated in [51] for the existence of optimal finite-memory strategies. Let us first recall a few definitions. Consider a non-empty set of colors K and a language $L \subseteq \mathsf{K}^*$. The language $[L] := \{\rho \in \mathsf{K}^\omega \mid \forall n \in \mathbb{N},\ \exists \pi \in L,\ \rho_{\leq n} \sqsubset \pi\}$ refers to the set of infinite words whose prefixes are also the prefixes of a word in L. Furthermore, the notation $\mathcal{R}(\mathsf{K})$ refers to the regular languages on a finite subset of K. In addition, for a preference $\preceq \subseteq \mathsf{K}^\omega \times \mathsf{K}^\omega$, and two languages $L, L' \subseteq \mathsf{K}^\omega$, $L \preceq L'$ refers to $\forall \rho \in L,\ \exists \rho' \in L',\ \rho \preceq \rho'$ and $L \prec L'$ refers to $\exists \rho' \in L',\ \forall \rho \in L,\ \rho \prec \rho'$. Note that $L \prec L' \Leftrightarrow \neg(L' \preceq L)$. Finally, for a memory skeleton $\langle M, m_{init}, \mu \rangle$ on K, and two memory states $m, m' \in M$, we denote $L^{\mathcal{M}}_{m,m'} := \{\rho \in \mathsf{K}^* \mid \overline{\mu}(m, \rho) = m'\}$. Let us consider the definitions of $\mathcal{M}$-monotony and $\mathcal{M}$-selectivity from [51]:

Definition 6.7 ($\mathcal{M}$-monotone preference)**.** *Let $\mathcal{M} = \langle M, m_{init}, \mu \rangle$ be a memory skeleton. A preferences $\preceq \subseteq \mathsf{K}^\omega \times \mathsf{K}^\omega$ is $\mathcal{M}$-monotone if, for all $m \in M$ and $L_1, L_2 \in \mathcal{R}(\mathsf{K})$: $(\exists \rho \in L^{\mathcal{M}}_{m_{init},m},\ [\rho \cdot L_1] \prec [\rho \cdot L_2]) \Rightarrow (\forall \rho' \in L^{\mathcal{M}}_{m_{init},m}, [\rho' \cdot L_1] \preceq [\rho' \cdot L_2])$.*

Definition 6.8 ($\mathcal{M}$-selective preference)**.** *Let $\mathcal{M} = \langle M, m_{init}, \mu \rangle$ be a memory skeleton. A preference $\preceq \subseteq \mathsf{K}^\omega \times \mathsf{K}^\omega$ is $\mathcal{M}$-selective if, for all $\rho \in \mathsf{K}^*$, $m = \overline{\mu}(m_{init}, \rho) \in M$, for all $L_1, L_2 \in \mathcal{R}(\mathsf{K})$ such that $L_1, L_2 \subseteq L^{\mathcal{M}}_{m,m}$, for all $L_3 \in \mathcal{R}(\mathsf{K})$, $[\rho \cdot (L_1 \cup L_2)^* \cdot L_3] \preceq [\rho \cdot L_1^*] \cup [\rho \cdot L_2^*] \cup [\rho \cdot L_3]$.*

By extension, we say that an objective W is $\mathcal{M}$-monotone and $\mathcal{M}$-selective if the preference relation $\preceq_W$ is. We have the theorem below.

Theorem 6.18 (Theorem 9 in [51])**.** *Let $\mathcal{M}$ be a memory skeleton and $W \subseteq \mathsf{K}^\omega$. The two following assertions are equivalent:*

1. *in every deterministic turn-based game with finitely many actions at each state and W winning objective, winning strategies for both players*

can be found among strategies implemented with memory skeleton $\mathcal{M}$*;*

2. W *and* $\mathsf{K}^\omega \setminus W$ *are* $\mathcal{M}$*-monotone and* $\mathcal{M}$*-selective.*

7 - Arbitrary game forms for Safety, Reachability and Büchi objectives

In this chapter and the next, contrary to the previous Chapter 6, we focus on the following question. Which classes of game forms is allowed to build finite-state parity games in which there are positional optimal strategies for Player A? Hence, the restrictions that we define are tailored for the very objective that we consider. In particular, we will see that depending on the parity objective that is considered (i.e. on the number of colors involved), the restrictions may differ. In each of these cases, the restrictions will be not only sufficient but also necessary. However, contrary to the next chapter where we restrict the setting to standard finite game forms, in this chapter, we consider arbitrary game forms. (Note that, since we consider finite-state parity games, we will always assume that there are finitely many outcomes.) As can be seen in Table 3.1, without any assumption on the local interactions occurring in a parity game, it is not always possible to play optimally, as it is already the case for Büchi games. Even when it is possible to play optimally, it may require infinite choice, and therefore infinite memory. As hinted above, the goal of this chapter is to characterize — and therefore to establish NSC-transfers — the arbitrary game forms ensuring the existence of positional optimal strategies when used in finite-state parity games. However, manipulating strategies in arbitrary finite-state concurrent games is quite involved, as it can be seen in Appendix 4.4.2 already in the case of reachability games. Hence, in this chapter, we decide to handle only three special cases. On the other hand, parity objectives with arbitrarily many colors will be handled in the next chapter, but only with standard game forms.

First, we consider safety games. In this case, it is straightforward to characterize the game forms ensuring the existence of positional optimal strategies. Indeed, these exactly correspond to the game forms maximizable w.r.t. Player A. (This holds even for infinite games.) This is stated in Proposition 7.1.

Second, we focus on reachability games. We use the results stated in Section 4.2 with standard finite game forms and proved in Appendix 4.4.2 with arbitrary game forms, and in particular Lemmas 4.9 and 4.10. With the help of these results, we characterize the game forms ensuring the existence of optimal strategies in reachability games. They are called reach-maximizable (RM for short) game forms. This is stated in Theorem 7.5. It turns out that they are strictly included in game forms maximizable w.r.t. Player A. This generalizes what we did in [39] (Theorem 36) where we only dealt with standard finite game forms.

Finally, we consider Büchi games. The characterization of the game forms ensuring the existence of optimal strategies in Büchi games is rather straightforward by combining what we have done already in this chapter and what we have shown in Section 5.1, in particular Lemma 5.4. We first obtain an NSC-transfer in Proposition 7.6. We then refine it into a more subtle NSC-transfer in Proposition 7.7. In that proposition, the kind of local interactions considered at each state of the game depends on the colors of the states. Recall, in a Büchi game, all states are colored with 0 or 1 and Player A wants to see infinitely often the color 1. Then, Proposition 7.7 states that for states of color 1, the game forms should be maximizable w.r.t. Player A: they ensure the existence of optimal strategies in safety games. However, for states of color 0, the game forms should be RM: they ensure the existence of optimal strategies in reachability games. This also generalizes what we did in [40] (Theorem 19) where we only dealt with standard finite game forms.

Note that, in Section 7.4, we explain why the NSC-transfers that we proved in [40] are not handled in this dissertation.

7.1 Safety objectives

First, as in Section 4.1, we only consider safety games without stopping states. Then, we seek those (arbitrary) game forms ensuring the existence of optimal strategies in finite-state safety games. It is rather straightforward to realize that, from a game form $\mathcal{F}$ that is not maximizable w.r.t. Player A, we can build a simple safety game where Player A does not have an optimal strategy. Furthermore, if all local interactions of a safety game are maximizable w.r.t. Player A, then by Corollary 4.3, Player A has a positional optimal strategy. Since this holds for infinite games, we may consider arbitrary game forms with infinitely many outcomes. We obtain the NSC-transfer below.

Proposition 7.1. *Among arbitrary game forms, being maximizable w.r.t. Player A is an NSC-transfer for the existence of Player-A positional optimal strategies in all infinite safety games without stopping states.*

Proof. Consider a set of outcomes O and a game form $\mathcal{F} = \langle \mathsf{Act_A}, \mathsf{Act_B}, \mathsf{O}, \varrho \rangle \in \mathsf{Form}(\mathsf{O})$ that is not maximizable w.r.t. Player A. Consider a valuation $v : \mathsf{O} \to [0, 1]$ such that Player A has no optimal GF-strategy in the game in normal form $\langle \mathcal{F}, v \rangle$. We build a simple game $\mathcal{G} = \langle \mathcal{C}, \mathsf{Safe} \rangle$ on $\mathcal{F}$ as follows: the set of states is equal to $\{q_{\mathsf{init}}, \top, \bot\} \cup v[\mathsf{O}]$. The local interaction at q_{init} is equal to $\langle \mathsf{Act_A}, \mathsf{Act_B}, v[\mathsf{O}], v \circ \varrho \rangle$. Furthermore all other states are trivial: all states $x \in v[\mathsf{O}]$ have as outcome $d_x \in \mathcal{D}(\{\top, \bot\})$ where $d_x := \{\top \mapsto x,\ \bot \mapsto 1 - x\}$. Both states $\top$ and $\bot$ are self-looping sinks. Finally, the only state of color 1, that Player A wants to avoid, is $\bot$. This concludes the definition of the game $\mathcal{G}$[1].

[1]Note that the simple game that we have built from $\mathcal{F}$ does not exactly fit Definition 5.1. However, we would have obtained an equivalent game if all trivial states $x \in v[\mathsf{O}]$ were replaced by stopping states of value x.

Clearly, all states $x \in v[\mathsf{O}]$ have value x. Consider any Player-A strategy $\mathsf{s_A} \in \mathcal{C}$. By Lemma 3.10, we have, $\chi_{\mathcal{G}}[\mathsf{s_A}](q_{\mathsf{init}}) \leq \mathsf{val}[\langle \mathsf{F}(q_{\mathsf{init}}), \chi_{\mathcal{G}}[\mathsf{A}] \rangle](\mathsf{s_A}(q_{\mathsf{init}})) = \mathsf{val}[\langle \mathcal{F}, v \rangle](\mathsf{s_A}(q_{\mathsf{init}})) < \mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{A}]$. This inequality comes from the fact that, by assumption, Player A has no optimal GF-strategy in the game in normal form $\langle \mathcal{F}, v \rangle$. Furthermore, we have by Proposition 3.9 for the last equality, $\mathsf{val}[\langle \mathcal{F}, v \rangle][\mathsf{A}] = \mathsf{val}[\langle \mathsf{F}(q_{\mathsf{init}}), \chi_{\mathcal{G}}[\mathsf{A}] \rangle][\mathsf{A}] = \chi_{\mathcal{G}}[\mathsf{A}](q_{\mathsf{init}})$. That is, the Player-$\mathsf{A}$ strategy $\mathsf{s_A}$ is not optimal from q_{init}, and this holds for Player-A strategy $\mathsf{s_A}$.

Consider now a safety game without stopping states where all local interactions are maximizable w.r.t. Player A. Then, by Corollary 4.3, Player A has a positional optimal strategy. □

7.2 Reachability objectives

Let us now focus on reachability games. As for safety games, the games we consider are without stopping states. However, contrary to safety games we will only consider finite-state games (since there does not always exist optimal strategies in infinite games, even turn-based ones). The proofs being quite technical for this objective, they are provided in appendix, we only give proof sketches in this section.

The goal is to characterize the game forms (with finitely many outcomes) ensuring the existence of positional optimal strategies in finite-state reachability games. The first step consists in defining the simple reachability games that we will use to properly define the well-behaved game forms. The simple games we consider are as follows: there is a central state q_{init} that is not in the target. From that central state, we can either loop, or reach a trivial state with a given probability to win, i.e. to reach the target. The color of this central state is 0, therefore looping indefinitely on it is loosing for Player A. This is formally defined below.

Definition 7.1 (Simple reachability games)**.** *Consider a finite set of outcomes* O *and a game form* $\mathcal{F} = \langle \Sigma_{\mathsf{A}}, \Sigma_{\mathsf{B}}, \mathsf{O}, \varrho \rangle \in \mathsf{Form}(\mathsf{O})$ *on that set of outcomes* O*. Consider some function* $m : \mathsf{O} \to \{q_{\mathsf{init}}\} \cup [0,1]$*. We define the reachability game* $\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m} := \langle \mathcal{C}_{\mathcal{F},m}, \mathsf{Reach} \rangle$ *such that* $\mathcal{C}_{\mathcal{F},m} := \langle Q_{\mathsf{O},m}, \mathsf{F}_{\mathcal{F},m}, \{0,1\}, \mathsf{col} \rangle$ *with:*

- $Q_{\mathsf{O},m} := \mathsf{Succ}_{\mathsf{O},m} \cup \{\top, \bot\}$ *with* $\mathsf{Succ}_{\mathsf{O},m} := \{q_{\mathsf{init}}\} \cup m[\mathsf{O}]$*;*
- *All states* $x \in Q_{\mathsf{O},m} \cap [0,1]$ *are trivial states with* $d_x \in \mathcal{D}(\{\top, \bot\})$ *as only outcome where* $d_x := \{\top \mapsto x,\ \bot \mapsto 1 - x\}$*;*
- $\mathsf{F}_{\mathcal{F},m}(q_{\mathsf{init}}) := \mathcal{F}^m = \langle \Sigma_{\mathsf{A}}, \Sigma_{\mathsf{B}}, \mathsf{Succ}_{\mathsf{O},m}, \mathbb{E}_m(\varrho) \rangle$*;*
- $\mathsf{col}(\bot) = \mathsf{col}(q_{\mathsf{init}}) := 0$, $\mathsf{col}(\top) := 1$.

We let $\alpha_{\mathcal{F},m} := \chi_{\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}}(q_{\mathsf{init}})$ *and, for all* $u \in [0,1]$*,* $v^u_{\mathcal{F},m} : Q_{\mathsf{O},m} \to [0,1]$ *such that* $v^u_{\mathcal{F},m}(q_{\mathsf{init}}) := u$*, for all* $x \in Q_{\mathsf{O},m} \cap [0,1]$*,* $v^u_{\mathcal{F},m}(x) := x$ *and* $v^u_{\mathcal{F},m}(\top) = 1$ *and* $v^u_{\mathcal{F},m}(\bot) = 0$*. When* $u = \alpha_{\mathcal{F},m}$*, it is omitted.*

$$\begin{bmatrix} x & y \\ y & z \end{bmatrix}$$

Figure 7.1: A standard finite game form $\mathcal{F}$.

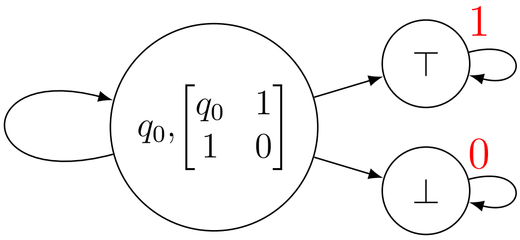

Figure 7.2: The reachability game $\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}$ without the states 1 and 0.

A simple reachability game $\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}$ is depicted in Figure 7.2 for $m(x) := q_{\mathsf{init}}$, $m(y) := 1$ and $m(z) := 0$, though we have not depicted the intermediate states 1 and 0. It is built on the game form $\mathcal{F}$ of Figure 7.1.

We consider the existence of Player-A positional optimal strategies in reachability games. A positional strategy in a simple reachability game $\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}$ is entirely defined by a Player-A GF-strategy in $\mathcal{F}$. We define below the optimal Player-A GF-strategies in simple reachability games.

Definition 7.2 (Optimal GF-strategies). *Consider a finite set of outcomes O, a game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ on that set of outcomes O and some function $m : Q \to \{q_{\mathsf{init}}\} \cup [0,1]$. For all Player-$\mathsf{A}$ GF-strategies $\sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}}$, the Player-$\mathsf{A}$ positional strategy $\mathsf{s}_{\mathsf{A}} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}_{\mathcal{F},m}}$ is* defined by *σ_{A} if $\mathsf{s}_{\mathsf{A}}(q_{\mathsf{init}}) := \sigma_{\mathsf{A}}$. A Player-$\mathsf{A}$ GF-strategies $\sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}}$ is* optimal *w.r.t. $(\mathcal{F}, m)$ if the Player-A positional strategy $\mathsf{s}_{\mathsf{A}} \in \mathsf{S}_{\mathsf{A}}^{\mathcal{C}}$ defined by σ_{A} is (subgame) optimal in the game $\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}$.*

The safe game forms for the existence of positional optimal strategies are then the game forms $\mathcal{F}$ for which, for all functions $m : \mathsf{O} \to \{q_{\mathsf{init}}\} \cup [0,1]$, there is a Player-$\mathsf{A}$ optimal GF-strategy w.r.t. $(\mathcal{F}, m)$. Such game forms are said to be reach maximizable (RM for short). This is defined below.

Definition 7.3 (Reach-maximizable game forms). *Consider a finite set of outcomes O and a game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ on that set of outcomes O. It is* reach-maximizable *(RM for short) if for all functions $m : Q \to \{q_{\mathsf{init}}\} \cup [0,1]$, there is a Player-$\mathsf{A}$ GF-strategy $\sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}}$ that is optimal w.r.t. $(\mathcal{F}, m)$.*

The definition, given in Definition 7.2, of optimal Player-A GF-strategies is not very practical in the sense that we do not know exactly how such GF-strategy behaves against Player-B GF-strategies. In particular, we would like to express how the GF-strategy behaves with similar notions than the ones used in Subsection 4.4.2. We give below a necessary and sufficient condition for a Player-A GF-strategy to be optimal that we will manipulate in the following.

Proposition 7.2 (Proof 7.5.1). *Consider a finite set of outcomes O, a game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ on that set of outcomes O and some function $m : \mathsf{O} \to \{q_{\mathsf{init}}\} \cup [0,1]$. A Player-$\mathsf{A}$ GF-strategy $\sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}}$ is optimal w.r.t. $(\mathcal{F}, m)$ if and* *only if we have $\alpha_{\mathcal{F},m} = 0$ or:*

1. *The* GF*-strategy σ_A is optimal in $\langle \mathcal{F}^m, v_{\mathcal{F},m}\rangle$; and*
2. *Mimicking the notations of Definition 4.13 and letting, for all $\sigma_\mathsf{B} \in \Sigma_\mathsf{B}$, $p_{\mathcal{F},m}(\sigma_\mathsf{A}, \sigma_\mathsf{B}) := \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{m^{-1}[[0,1]]}\rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B})$, the* GF*-strategy σ_A ensures that:*

$$\inf_{\sigma_\mathsf{B} \in \Sigma_\mathsf{B}} p_{\mathcal{F},m}(\sigma_\mathsf{A}, \sigma_\mathsf{B}) > 0$$

Proof sketch. Assume that the GF-strategy σ_A is optimal w.r.t. $(\mathcal{F}, m)$ and that $\alpha_{\mathcal{F},m} > 0$. Then, the positional Player-A strategy $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}_{\mathcal{F},m}}$ defined by σ_A is subgame optimal in the game $\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}$. Therefore, it is locally optimal by Theorem 3.12[2]. This implies that the GF-strategy σ_A satisfies item 1. Furthermore, if it does not satisfy item 2., it means that, against the strategy s_A, Player B has strategies to ensure that the probability to ever see a state in $[0, 1]$ (and therefore the target) is arbitrarily close to 0. Hence, the value of the strategy s_A would be 0, which is not possible since $\alpha_{\mathcal{F},m} > 0$.

Assume now that $\alpha_{\mathcal{F},m} > 0$ and that the GF-strategy σ_A satisfies items 1. and 2. Let $\delta := \inf_{\sigma_\mathsf{B} \in \Sigma_\mathsf{B}} p_{\mathcal{F},m}(\sigma_\mathsf{A}, \sigma_\mathsf{B}) > 0$. Consider any Player-B strategy $\mathsf{s}_\mathsf{B} \in \mathsf{S}_\mathsf{B}^{\mathcal{C}_{\mathcal{F},m}}$ against the Player-A strategy $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}_{\mathcal{F},m}}$ generated by the GF-strategy σ_A. As long as the game loops on q_init, there is probability at least δ to exit to a state in $[0, 1]$, by definition of δ. Therefore, the game loops indefinitely on q_init with probability 0. Furthermore, whenever there is a positive probability to exit to a state in $[0, 1]$, the expected value of the states reached is at least $\alpha_{\mathcal{F},m}$, by item 1. Hence, the Player-A strategy s_A has value $\alpha_{\mathcal{F},m}$, it is therefore optimal in $\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}$. □

Let us now consider what happens when RM game forms are used in finite-state reachability games.

Theorem 7.3. *Consider a finite-state reachability game $\mathcal{G} = \langle \mathcal{C}, \mathsf{Reach}\rangle$, and assume that, for all $q \in Q$ such that $\mathsf{col}(q) = 0$, the game form $\mathsf{F}(q)$ is* RM*. Then, Player* A *has a positional (subgame) optimal strategy.*

To prove this theorem, we are going to extract simple reachability games from the global game $\mathcal{G}$. Let us define formally this operation.

Definition 7.4. *Consider a finite-state reachability game $\mathcal{G} = \langle \mathcal{C}, \mathsf{Reach}\rangle$. Consider a subset of states $S \subseteq Q$. Then, for all $q \in Q$, we let $m_q^S : Q \to \{q_\mathsf{init}\} \cup [0, 1]$ such that, for all $q' \in Q$:*

$$m_q^S(q') := \begin{cases} \chi_\mathcal{G}[\mathsf{A}](q') \in [0, 1] & \text{if } q' \in Q \setminus S \\ q_\mathsf{init} & \text{otherwise} \end{cases}$$

[2] Since the target in the reachability game $\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}$ is self-looping, the objective can be seen as being a Büchi objective, and therefore as being PI.

This extraction satisfies a crucial property: when all the state in S have the same Player-A value $u \in [0,1]$ and are not in the target, then the maximum of the values of the simple games is at least u. This is formally stated in the lemma below.

Lemma 7.4 (Proof 7.5.2). *Consider a finite-state reachability game $\mathcal{G} = \langle \mathcal{C}, \mathsf{Reach} \rangle$. Consider some $u \in [0,1]$ and a non-empty subset of states $\emptyset \neq S \subseteq Q$ such that, for all $q \in Q$, we have $\chi_{\mathcal{G}}[\mathsf{A}](q) = u$. If $S \cap T = \emptyset$, we have:*

$$\max_{q \in Q} \alpha_{\mathsf{F}(q), m_q^S} \geq u$$

Proof sketch. Assume towards a contradiction that it is not the case, i.e. that $x := \max_{q \in Q} \alpha_{\mathsf{F}(q), m_q^S} < u$. Then, we can show that for all states $q \in Q$ and Player-A strategies $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}}$, the value of this Player-A strategy from state q is at most x. Indeed, consider some $\varepsilon > 0$. Then, for all $\rho \in S^+$, we have $\mathsf{s}_\mathsf{A}(\rho) \in \Sigma_\mathsf{A}^{\rho_{\mathsf{lt}}}$. Then, at state ρ_{lt}, either Player B can ensure to loop on states in S with probability arbitrarily close to 1. This occurs when $\inf_{\sigma_\mathsf{B} \in \Sigma_\mathsf{B}} p_{\mathcal{F},m}(\sigma_\mathsf{A}, \sigma_\mathsf{B}) = 0$ (from Proposition 7.2). Or, there is some $\delta > 0$ such that, regardless of what Player B plays in $\mathsf{F}(\rho_{\mathsf{lt}})$, the probability to see states outside S is at least δ. But in that case, Player B has a GF-strategy to ensure that, against $\mathsf{s}_\mathsf{A}(\rho)$, the expected Player-$\mathsf{A}$ value of the states seen (outside of S) is at most $x + \varepsilon$, by definition of m. Furthermore, from all states $q \in Q \setminus S$ outside of S, Player B can play against the strategy s_A to ensure that the value from q is ε-close to the Player-A value of q. Since looping indefinitely in S is losing for Player A, it follows that the value of the Player-A strategy s_A from q is at most x. Hence the contradiction. □

We can now give an informal proof of Theorem 7.3. The formal proof is given in Subsection 7.5.3.

Proof sketch. Let us assume that all states in the target are self-looping sinks. This does not change the game since there are no stopping states, once the target is reached, the game has value 1. Consider Definition 4.11: we want to show that $\mathsf{Sec}(\emptyset) = Q$. In turn, with Lemma 4.10, this will show that there is a Player-A positional optimal strategy in $\mathcal{G}$.

Assume towards a contradiction that we have $\mathsf{Sec}(\emptyset) \neq Q$. Let $B := Q \setminus \mathsf{Sec}(\emptyset) \neq \emptyset$. Note that $\mathsf{col}[B] = \{0\}$. Consider the greatest $u \in (0,1]$ such that $B_u := \{q \in B \mid \chi_{\mathcal{G}}[\mathsf{A}](q) = u\} \neq \emptyset$. Let us apply Lemma 7.4 to the set B_u: there is some state $q \in B_u$ such that $\alpha_{\mathsf{F}(q), m_q^{B_u}} \geq u$. By assumption, the local interaction $\mathsf{F}(q)$ is RM. Therefore, there is a Player-A GF-strategy $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}^q$ that is optimal w.r.t. $(\mathsf{F}(q), m_q^{B_u})$. Let us argue that this GF-strategy is progressive w.r.t. $\mathsf{Sec}(\emptyset)$ (recall Definition 4.14). First, it is indeed optimal in the game in normal form $\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle$ by item 1. of Proposition 7.2 and by definition of the function $m_q^{B_u}$. Furthermore, by item 2. of that same Proposition 7.2, letting

$\delta := \inf_{\sigma_B \in \Sigma_B} p_{\mathsf{F}(q), m_q^{B u}}(\sigma_A, \sigma_B)$, we have $\delta > 0$. Furthermore, it can be shown (since σ_A satisfies item 1. of Proposition 7.2), that with the GF-strategy σ_A and for all Player-B GF-strategies in the game form $\mathsf{F}(q)$, if there is probability p to see states of values different from u, then there is probability at least $p \cdot y$ to see states of value more than u, for $y := \min_{x \in \chi_{\mathcal{G}}[A][Q],\, x < u} \frac{u-x}{1-x} > 0$. Since all states of value u outside of S are in $\mathsf{Sec}(\emptyset)$ and, by definition of u, all states of value more than u are also in $\mathsf{Sec}(\emptyset)$, it follows that for all Player-B GF-strategies in $\mathsf{F}(q)$, with the GF-strategy σ_A, there is probability at least $\delta \cdot y > 0$ to see a state in $\mathsf{Sec}(\emptyset)$. Therefore the Player-A GF-strategy σ_A is progressive w.r.t. $\mathsf{Sec}(\emptyset)$. This is in contradiction with the fact that the state q is not in $\mathsf{Sec}(\emptyset)$. □

Overall, we obtain the following NSC-transfer:

Theorem 7.5. *Among arbitrary game forms with finitely many outcomes, being* RM *is an NSC-transfer for the existence of Player-*A *positional optimal strategies in finite-state reachability games without stopping states.*

Proof. This is a direct consequence of Proposition 7.2 and Theorem 7.3. □

7.3 Büchi objectives

Let us now consider finite-state Büchi games. In fact, for the Büchi objective, it suffices to use RM game forms. This is stated as an NSC-transfer in the proposition below.

Proposition 7.6. *Among arbitrary game forms with finitely many outcomes, being* RM *is an NSC-transfer for the existence of Player-*A *positional optimal strategies in finite-state Büchi games without stopping states.*

We do not provide a proof of this statement for now since we state and prove a slightly more subtle (and stronger) statement. Indeed, we can combine both the RM game forms characterized for reachability games and the game forms maximizable w.r.t. Player A characterized for safety games. We obtain straightforwardly an NSC-transfer by applying Lemma 5.4 that is different from the ones we have stated so far. Indeed, the class of game forms to be used depends on the colors of the states considered. Let us formally state the result, we will explain afterwards exactly what it means.

Proposition 7.7. *Among arbitrary game forms with finitely many outcomes, being:*

- *maximizable w.r.t. Player* A *when colored with 1;*
- RM *when colored with 0;*

*is an NSC-transfer for the existence of Player-*A *positional optimal strategies in finite-state Büchi games.*

What this theorem means is the following: given any game form that is not maximizable w.r.t. Player A, we can build a simple Büchi game where the central state q_{init} is colored with 1 and where Player A has no positional optimal strategy. Similarly, from any game form that is not RM, we can build a simple Büchi game where the central state q_{init} is colored with 0 and where Player A has no positional optimal strategy. Conversely, in any finite-state Büchi game where all local interactions at states colored with 1 (resp. 0) are maximizable w.r.t. Player A (resp. RM), Player A has a positional optimal strategy.

Note in particular that Proposition 7.6 is a direct corollary of Proposition 7.7 because all RM game forms are maximizable w.r.t. Player A.

Proof. Consider a game form $\mathcal{F}$ that is not maximizable w.r.t. Player A. Then, the simple game built on $\mathcal{F}$ to prove Proposition 7.1 can be seen as a Büchi game up to reversing the roles of the colors 0 and 1. In that game, Player A has no positional optimal strategy.

Consider now a game form $\mathcal{F}$ that is not RM. Then, the simple game built on $\mathcal{F}$ to prove Theorem 7.5 from Definition 7.1 can be seen as a Büchi game (since the target is self-looping). In that game, Player A has no positional optimal strategy.

Now, consider a finite-state Büchi game $\mathcal{G}$ where all local interactions at states colored with 1 (resp. 0) are maximizable w.r.t. Player A (resp. RM). We want to apply Lemma 5.4. First, in all non-trivial states of the reachability game $\mathcal{G}^{\mathsf{Reach}}$, all local interactions are RM, therefore by Theorem 7.3, there is an optimal strategy in the game $\mathcal{G}^{\mathsf{Reach}}$. Second, all local interactions at states of color 1 are maximizable w.r.t. Player A. It follows that we can indeed apply Lemma 5.4. We obtain that Player A has a positional optimal strategy in $\mathcal{G}$. □

7.4 Discussion and future work

In this chapter, we have proved NSC-transfers, among arbitrary game forms, for the existence of positional optimal strategies in finite-state safety, reachability and Büchi games. This extends, to the case of arbitrary game forms, what we have done previously: [39, Theorem 36] for reachability games and [40, Theorem 19] for Büchi games. However, we have not extended, or even stated, arguably the two main results we have shown in [40]: Theorem 22 and 25 which give NSC-transfers (though this terminology is not used) for the existence of positional almost-optimal strategies in finite Büchi games and for the existence of positional optimal strategies in finite co-Büchi games, respectively, restricting to standard finite game forms.

The ideas behind the proofs of these theorems lie in how the value is obtained in standard finite parity games, namely fixed points. Recall, we proved in Proposition 4.7 that the Player-A value of reachability games can be computed with a least fixed point, even with non-standard game forms. In fact, with standard finite local interactions, it was shown in [32] that the value of parity games could be computed with nested fixed points, i.e. a greatest fixed point, nested with a least fixed point, etc. Furthermore, if the highest color appearing in the parity game is n, then the number of nested fixed points considered is n. In [40, Theorem 22, 25], the definitions of the safe game forms considered is based on the fixed points needed to compute the values in Büchi and co-Büchi games, i.e. nested least and greatest fixed points, the order depending on the objective. The proofs of these theorems were quite intricate as it involves linking the fixed points on the concurrent game and on the local interactions.

As a future work, we want to extend [40, Theorem 22, 25] to the case of arbitrary game forms, while avoiding the nested fixed points arguments. One of the main difficulty that will arise is to properly state a proposition detailing what being safe amounts to (w.r.t. the objective/type of strategies considered), as we did in Proposition 7.2 for the reachability objective. This needs to hold even with arbitrary game forms. As witnessed in Definitions 4.8, 4.14 and Definitions 4.9, 4.15, definitions compatible with arbitrary game forms are more intricate compared to their counterparts well-suited for the case of standard finite game forms, that we gave in [40].

7.5 Appendix

7.5.1 . Proof of Proposition 7.2

First, we make a quick remark: $\alpha_{\mathcal{F},m}$ is equal to the Player-A value of the game in normal form $\langle \mathcal{F}^m, v_{\mathcal{F},m} \rangle$.

Lemma 7.8. *Consider a finite set of outcomes O, a game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ on that set of outcomes O and some function $m : \mathsf{O} \to \{q_{\mathsf{init}}\} \cup [0,1]$. We have:*

$$\alpha_{\mathcal{F},m} = \mathsf{val}[\langle \mathcal{F}^m, v_{\mathcal{F},m} \rangle][\mathsf{A}]$$

Proof. By definition, we have $\alpha_{\mathcal{F},m} = \chi_{\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}}[\mathsf{A}](q_{\mathsf{init}})$, $\mathsf{F}_{\mathcal{F},m}(q_{\mathsf{init}}) = \mathcal{F}^m$ and $v_{\mathcal{F},m} = \chi_{\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}}[\mathsf{A}]^{q_{\mathsf{init}}}$. The result is therefore given by Proposition 3.9. $\square$

Let us now express exactly what is the value of a Player-A positional strategy generated by a GF-strategy in a simple reachability game.

Lemma 7.9. *Consider a finite set of outcomes O, a game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ on that set of outcomes O and some function $m : \mathsf{O} \to \{q_{\mathsf{init}}\} \cup [0,1]$. Let $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}_{\mathcal{F},m}}$ be the Player-A positional strategy generated by some Player-A GF-strategy $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}$. If:*

$$\inf_{\sigma_\mathsf{B} \in \Sigma_\mathsf{B}} p_{\mathcal{F},m}(\sigma_\mathsf{A}, \sigma_\mathsf{B}) = 0$$

then $\chi_{\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}}[\mathsf{s_A}](q_{\mathsf{init}}) = 0$. *Otherwise:*

$$\chi_{\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}}[\mathsf{s_A}](q_{\mathsf{init}}) = \sup\{u \in [0,1] \mid \mathsf{val}[\langle \mathcal{F}^m, v^u_{\mathcal{F},m}\rangle][\sigma_{\mathsf{A}}] \geq u\}$$

Proof. Let $W := \mathsf{col}^{-1}[\mathsf{Reach}]$ be the Player-A winning set. First, assume that $\inf_{\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}} p_{\mathcal{F},m}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) = 0$ and let $\varepsilon > 0$. Let us define a Player-B strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}_{\mathcal{F},m}}$ such that, for all $n \in \mathbb{N}$, we have $\mathsf{s_B}(q^n_{\mathsf{init}}) \in \Sigma_{\mathsf{B}}$ such that $p_{\mathcal{F},m}(\sigma_{\mathsf{A}}, \mathsf{s_B}(q^n_{\mathsf{init}})) \leq \frac{\varepsilon}{2^{n+1}}$. Then, we have by Definition 1.28 and 1.29 for the second equality:

$$\begin{aligned}
\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_{\mathcal{F},m},q_{\mathsf{init}}}[W] &\leq \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_{\mathcal{F},m},q_{\mathsf{init}}}[(q_{\mathsf{init}})^* \cdot [0,1]] = \sum_{n \in \mathbb{N}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_{\mathcal{F},m},q_{\mathsf{init}}}[(q_{\mathsf{init}})^n \cdot [0,1]] \\
&\leq \sum_{n \in \mathbb{N}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_{\mathcal{F},m},q^n_{\mathsf{init}}}[[0,1]] = \sum_{n \in \mathbb{N}} \mathsf{out}[\langle \mathsf{F}_{\mathcal{F},m}(q_{\mathsf{init}}), \mathbb{1}_{[0,1]}\rangle](\mathsf{s_A}(q^n_{\mathsf{init}}), \mathsf{s_B}(q^n_{\mathsf{init}})) \\
&= \sum_{n \in \mathbb{N}} \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{m^{-1}[[0,1]]}\rangle](\sigma_{\mathsf{A}}, \mathsf{s_B}(q^n_{\mathsf{init}})) = \sum_{n \in \mathbb{N}} p_{\mathcal{F},m}(\sigma_{\mathsf{A}}, \mathsf{s_B}(q^n_{\mathsf{init}})) \\
&\leq \sum_{n \in \mathbb{N}} \frac{\varepsilon}{2^{n+1}} = \varepsilon
\end{aligned}$$

As this holds for all $\varepsilon > 0$, it follows that the value of the Player-A strategy $\mathsf{s_A}$ is 0 from q_{init}.

Assume now it is not the case, i.e that $\delta := \inf_{\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}} p_{\mathcal{F},m}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})$ is positive: $\delta > 0$. Then, regardless of the Player-B strategy, almost-surely the game does not loop indefinitely on q_{init} by definition of $\delta > 0$: each time the game loops on q_{init} there is probability at least δ to exit to a trivial state. Therefore, almost-surely, the game ends up in $\{\top, \bot\}$.

Let $x := \sup\{u \in [0,1] \mid \mathsf{val}[\langle \mathcal{F}, v^u_{\mathcal{F},m}\rangle][\sigma_{\mathsf{A}}] \geq u\}$. Consider some $\varepsilon > 0$ and let $u \geq x - \varepsilon$ such that $\mathsf{val}[\langle \mathcal{F}, v^u_{\mathcal{F},m}\rangle][\sigma_{\mathsf{A}}] \geq u$. Let us apply Corollary 3.14[3] to show that the Player-A strategy $\mathsf{s_A}$ guarantees the valuation $v^u_{\mathcal{F},m}$. It satisfies the first condition of that corollary by assumption. Furthermore, as mentioned above, almost-surely, the game loops on $\top$ (and Player A wins), or loops on $\bot$ of value 0. Hence, it also satisfies the second condition of Corollary 3.14. Therefore, it dominates the valuation $v^u_{\mathcal{F},m}$, with $v^u_{\mathcal{F},m}(q_{\mathsf{init}}) = u$. Hence the value of the Player-A strategy $\mathsf{s_A}$ from q_{init} is at least $u \geq x - \varepsilon$.

Consider some $\varepsilon > 0$ and let $u \leq x + \varepsilon$ such that $\mathsf{val}[\langle \mathcal{F}, v^u_{\mathcal{F},m}\rangle][\sigma_{\mathsf{A}}] < u$. Consider then a Player-B positional strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}_{\mathcal{F},m}}$ such that:

$$\mathsf{out}[\langle \mathcal{F}, v^u_{\mathcal{F},m}\rangle](\mathsf{s_A}(q_{\mathsf{init}}), \mathsf{s_B}(q_{\mathsf{init}})) = \mathsf{out}[\langle \mathcal{F}, v^u_{\mathcal{F},m}\rangle](\sigma_{\mathsf{A}}, \mathsf{s_B}(q_{\mathsf{init}})) \leq u$$

Then, in the stochastic tree $\mathcal{T}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_{\mathcal{F},m},q_{\mathsf{init}}}$, the valuation $v^u_{\mathcal{F},m}$ is non-increasing (recall Definition 2.3). Hence, by Proposition 2.9, since the game almost-surely

[3]This only applies to PI games, however since the target is a self-looping sink, the game $\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}$ can be seen as a Büchi game.

ends up in $\{\top, \bot\}$, it follows that:

$$u = v^u_{\mathcal{F},m}(q_{\mathsf{init}}) \geq \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_{\mathcal{F},m},q_{\mathsf{init}}}(q^*_{\mathsf{init}} \cdot [0,1] \cdot \top) \cdot v^u_{\mathcal{F},m}(\top) + \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_{\mathcal{F},m},q_{\mathsf{init}}}(q^*_{\mathsf{init}} \cdot [0,1] \cdot \bot) \cdot v^u_{\mathcal{F},m}(\bot)$$
$$= \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_{\mathcal{F},m},q_{\mathsf{init}}}(q^*_{\mathsf{init}} \cdot [0,1] \cdot \top) = \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_{\mathcal{F},m},q_{\mathsf{init}}}[W]$$

Hence the value of the Player-A strategy $\mathsf{s_A}$ from q_{init} is at most $u \leq x + \varepsilon$. □

The proof of Proposition 7.2 is now direct.

Proof. Assume that the Player-A GF-strategy σ_{A} is optimal w.r.t. $(\mathcal{F}, m)$ and that $\alpha_{\mathcal{F},m} > 0$. Then, by Lemma 7.9, it must be that $\inf_{\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}} p_{\mathcal{F},m}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) > 0$. Furthermore, we have $\alpha_{\mathcal{F},m} = \chi_{\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}}[\mathsf{s_A}](q_{\mathsf{init}}) = \sup\{u \in [0,1] \mid \mathsf{val}[\langle \mathcal{F}^m, v^u_{\mathcal{F},m}\rangle][\sigma_{\mathsf{A}}] \geq u\}$. Hence, for all $u < \alpha_{\mathcal{F},m}$, we have $u \leq \mathsf{val}[\langle \mathcal{F}^m, v^u_{\mathcal{F},m}\rangle][\sigma_{\mathsf{A}}] \leq \mathsf{val}[\langle \mathcal{F}^m, v_{\mathcal{F},m}\rangle][\sigma_{\mathsf{A}}]$. Therefore, $\alpha_{\mathcal{F},m} \leq \mathsf{val}[\langle \mathcal{F}^m, v_{\mathcal{F},m}\rangle][\sigma_{\mathsf{A}}]$. In addition, $\alpha_{\mathcal{F},m} = \mathsf{val}[\langle \mathcal{F}^m, v_{\mathcal{F},m}\rangle][\mathsf{A}]$ by Lemma 7.8. Therefore, the Player-A GF-strategy σ_{A} is optimal in the game in normal form $\langle \mathcal{F}^m, v_{\mathcal{F},m}\rangle$.

Conversely, if $\alpha_{\mathcal{F},m} = 0$ then any Player-A strategy is optimal from q_{init} in the game $\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}$ and therefore the GF-strategy σ_{A} is optimal w.r.t. $(\mathcal{F}, m)$. Assume now that the Player-A GF-strategy σ_{A} satisfies items 1. and 2. Since $\alpha_{\mathcal{F},m} = \mathsf{val}[\langle \mathcal{F}^m, v_{\mathcal{F},m}\rangle][\mathsf{A}]$ and σ_{A} satisfies item 1., it follows that $\mathsf{val}[\langle \mathcal{F}^m, v_{\mathcal{F},m}\rangle][\sigma_{\mathsf{A}}] = \alpha_{\mathcal{F},m}$. Then, by Lemma 7.9 and since σ_{A} satisfies item 2., we obtain that $\chi_{\mathcal{G}^{\mathsf{Reach}}_{\mathcal{F},m}}[\mathsf{s_A}](q_{\mathsf{init}}) \geq \alpha_{\mathcal{F},m}$. Therefore, it is optimal w.r.t. $(\mathcal{F}, m)$. □

7.5.2 . Proof of Lemma 7.4

Proof. We let $W := (\mathsf{col}^\omega)^{-1}[\mathsf{Reach}] \subseteq Q^\omega$ be the Player-A winning set. Let $x := \max_{q \in S} \alpha_{\mathsf{F}(q),m^S_q} \in [0,1]$ and assume towards a contradiction that $x < u$. Consider any state $q \in S$ and Player-A strategy $\mathsf{s_A} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{A}}$. Let us show that $\chi_{\mathcal{G}}[\mathsf{s_A}](q) \leq x$. For all $q \in S$, we let p_q denote the function $p_{\mathsf{F}(q),m^S_q}$ from Proposition 7.2. Let $\varepsilon > 0$.

Consider some $\rho \in q \cdot S^*$. Since $\alpha_{\mathsf{F}(\rho_{\mathsf{lt}}),m^S_{\rho_{\mathsf{lt}}}} \leq x$, it follows that, by Lemma 7.9, that:

- $\inf_{\sigma_{\mathsf{B}} \in \Sigma^{\rho_{\mathsf{lt}}}_{\mathsf{B}}} p_{\rho_{\mathsf{lt}}}(\mathsf{s_A}(\rho), \sigma_{\mathsf{B}}) = 0$; or
- for all $x < z$, we have:

$$\mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}})^{m^S_{\rho_{\mathsf{lt}}}}, v^x_{\mathsf{F}(\rho_{\mathsf{lt}}),m^S_{\rho_{\mathsf{lt}}}}\rangle][\mathsf{s_A}(\rho)] \leq \mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}})^{m^S_{\rho_{\mathsf{lt}}}}, v^z_{\mathsf{F}(\rho_{\mathsf{lt}}),m^S_{\rho_{\mathsf{lt}}}}\rangle][\mathsf{s_A}(\rho)] < z$$

  Therefore: $\mathsf{val}[\langle \mathsf{F}(\rho_{\mathsf{lt}})^{m^S_{\rho_{\mathsf{lt}}}}, v^x_{\mathsf{F}(\rho_{\mathsf{lt}}),m^S_{\rho_{\mathsf{lt}}}}\rangle][\mathsf{s_A}(\rho)] \leq x$.

We let $S_{\mathsf{Lp}} := \{\rho \in S^* \mid \inf_{\sigma_{\mathsf{B}} \in \Sigma^{q \cdot \rho_{\mathsf{lt}}}_{\mathsf{B}}} p_{\rho_{\mathsf{lt}}}(\mathsf{s_A}(q \cdot \rho), \sigma_{\mathsf{B}}) = 0\}$ and $S_{\leq x} := S^* \setminus S_{\mathsf{Lp}}$.

Let us now define a Player-B strategy $\mathsf{s_B} \in \mathsf{S}^{\mathcal{C}}_{\mathsf{B}}$:

- For all $\rho \in S^+$, we define $\mathsf{s}^\rho_{\mathsf{B}}$ such that, for all $q' \in Q \setminus S$: $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho \cdot q'}[W] \leq \chi_{\mathcal{G}}[\mathsf{A}](q') + \frac{\varepsilon}{3}$;

- For all $\rho \in q \cdot S_{\mathsf{Lp}}$, we let $\mathsf{s_B}(\rho) \in \Sigma_{\mathsf{B}}^{\rho_{\mathsf{lt}}}$ be such that $p_{\rho_{\mathsf{lt}}}(\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \leq \frac{\varepsilon}{3 \cdot |S|^{|\rho|-1} \cdot 2^{|\rho|}}$;

- For all $\rho \in q \cdot S_{\leq x}$, we let $\mathsf{s_B}(\rho) \in \Sigma_{\mathsf{B}}^{\rho_{\mathsf{lt}}}$ be such that $\mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}})^{m^S_{\rho_{\mathsf{lt}}}}, v^x_{\mathsf{F}(\rho_{\mathsf{lt}}), m^S_{\rho_{\mathsf{lt}}}} \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \leq x + \frac{\varepsilon}{3 \cdot |S|^{|\rho|-1} \cdot 2^{|\rho|}}$.

Then, consider some $\rho \in q \cdot S_{\mathsf{Lp}}$. We have, by Definition 1.28 and Definition 7.4:

$$\begin{aligned}
\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[Q \setminus S] &= \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \mathbb{1}_{Q \setminus S} \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \\
&= \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \mathbb{1}_{(m^S_{\rho_{\mathsf{lt}}})^{-1}[0,1]} \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \\
&= p_{\rho_{\mathsf{lt}}}(\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \leq \frac{\varepsilon}{3 \cdot |S|^{|\rho|-1} \cdot 2^{|\rho|}}
\end{aligned}$$

We obtain:

$$\begin{aligned}
\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[S_{\mathsf{Lp}} \cdot Q \setminus S] &= \sum_{n \in \mathbb{N}} \sum_{\rho \in S^n \cap S_{\mathsf{Lp}}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}(\rho) \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q \cdot \rho}[Q \setminus S] \\
&\leq \sum_{n \in \mathbb{N}} \sum_{\rho \in S^n \cap S_{\mathsf{Lp}}} \frac{\varepsilon}{3 \cdot |S|^n \cdot 2^{n+1}} \leq \sum_{n \in \mathbb{N}} \frac{\varepsilon}{3 \cdot 2^{n+1}} = \frac{\varepsilon}{3}
\end{aligned}$$

Furthermore, for all $z \in [0,1] \setminus \{u\}$, we let $Q_z := \chi_{\mathcal{G}}[\mathsf{A}]^{-1}[\{z\}]$ and $Q_u := \chi_{\mathcal{G}}[\mathsf{A}]^{-1}[\{u\}] \setminus S$. That way, $Q \setminus S = \cup_{z \in [0,1]} Q_z$. Then, for all $\rho \in q \cdot S_{\leq x}$, we have:

$$\begin{aligned}
x + \frac{\varepsilon}{3 \cdot |S|^{|\rho|-1} \cdot 2^{|\rho|}} &\geq \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}})^{m^S_{\rho_{\mathsf{lt}}}}, v^x_{\mathsf{F}(\rho_{\mathsf{lt}}), m^S_{\rho_{\mathsf{lt}}}} \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \\
&= \sum_{z \in [0,1]} \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}})^{m^S_{\rho_{\mathsf{lt}}}}, \mathbb{1}_{(v^x_{\mathsf{F}(\rho_{\mathsf{lt}}), m^S_{\rho_{\mathsf{lt}}}})^{-1}[\{z\}]} \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \cdot z \\
&= x \cdot \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \mathbb{1}_S \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) + \sum_{z \in [0,1]} \mathsf{out}[\langle \mathsf{F}(\rho_{\mathsf{lt}}), \mathbb{1}_{Q_z} \rangle](\mathsf{s_A}(\rho), \mathsf{s_B}(\rho)) \cdot z \\
&= x \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[S] + \sum_{z \in [0,1]} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[Q_z] \cdot z
\end{aligned}$$

Therefore, we have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[Q \setminus S] \cdot x + \frac{\varepsilon}{3 \cdot |S|^{|\rho|-1} \cdot 2^{|\rho|}} \geq \sum_{z \in [0,1]} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},\rho}[Q_z] \cdot z$$

Hence:

$$\begin{aligned}
\sum_{z \in [0,1]} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[S_{\leq x} \cdot Q_z] \cdot z &\leq \sum_{n \in \mathbb{N}} \sum_{\rho \in S^n \cap S_{\leq x}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}(\rho) \cdot \left(\sum_{z \in [0,1]} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q \cdot \rho}[Q_z] \cdot z \right) \\
&\leq \sum_{n \in \mathbb{N}} \sum_{\rho \in S^n \cap S_{\leq x}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}(\rho) \cdot \left(\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q \cdot \rho}[Q \setminus S] \cdot x + \frac{\varepsilon}{3 \cdot |S|^n \cdot 2^{n+1}} \right) \\
&\leq \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[S_{\leq x} \cdot (Q \setminus S)] \cdot x + \sum_{n \in \mathbb{N}} \sum_{\rho \in S^n \cap S_{\leq x}} \frac{\varepsilon}{3 \cdot |S|^n \cdot 2^{n+1}} \\
&\leq \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[S_{\leq x} \cdot (Q \setminus S)] \cdot x + \frac{\varepsilon}{3}
\end{aligned}$$

Overall, since we have reachability objective, staying indefinitely in S is losing for Player A. Hence, we have:

$$\begin{aligned}
\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[W] &= \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[W \cap S^\omega] + \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[W \cap S_{\mathsf{Lp}} \cdot (Q \setminus S)] + \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[W \cap S_{\leq x} \cdot (Q \setminus S)] \\
&\leq \frac{\varepsilon}{3} + \sum_{z \in [0,1]} \sum_{q' \in Q_z} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[W \cap S_{\leq x} \cdot q'] \\
&= \frac{\varepsilon}{3} + \sum_{z \in [0,1]} \sum_{q' \in Q_z} \sum_{\rho \in S_{\leq x}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}(\rho \cdot q') \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q \cdot \rho \cdot q'}[W] \\
&\leq \frac{\varepsilon}{3} + \sum_{z \in [0,1]} \sum_{q' \in Q_z} \sum_{\rho \in S_{\leq x}} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}(\rho \cdot q') \cdot (\chi_{\mathcal{G}}[\mathsf{A}](q') + \frac{\varepsilon}{3}) \\
&= \frac{\varepsilon}{3} + \sum_{z \in [0,1]} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[S_{\leq x} \cdot Q_z] \cdot z + \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[S^* \cdot (Q \setminus S)] \cdot \frac{\varepsilon}{3} \\
&\leq \frac{2 \cdot \varepsilon}{3} + \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C},q}[S_{\leq x} \cdot (Q \setminus S)] \cdot x + \frac{\varepsilon}{3} \\
&\leq x + \varepsilon
\end{aligned}$$

As this holds for all $\varepsilon > 0$, the value, from q, of the Player-A strategy $\mathsf{s_A}$ is at most x, hence the contradiction. □

7.5.3 . Proof of Theorem 7.3

Proof. Let us assume that all states in the target are self-looping sinks. This does not change the game since there are no stopping states, hence once the target is reached, the game has value 1.

We proceed as in the proof sketch. That is, considering Definition 4.11: we want to show that $\mathsf{Sec}(\emptyset) = Q$, which in turn, with Lemma 4.10, will show that there is a Player-A positional optimal strategy in $\mathcal{G}$. Hence, assume towards a contradiction that we have $\mathsf{Sec}(\emptyset) \neq Q$. Let $B := Q \setminus \mathsf{Sec}(\emptyset) \neq \emptyset$. Note that $\mathsf{col}[B] = \{0\}$. Consider the greatest $u \in [0,1]$ such that $B_u := \{q \in B \mid \chi_{\mathcal{G}}[\mathsf{A}](q) = u\} \neq \emptyset$. It must be that $u > 0$ since no state in B is of value 0. Let us apply Lemma 7.4 to the set B_u: there is some state $q \in B_u$ such that $\alpha_{\mathsf{F}(q),m_q^{B_u}} \geq u$. By assumption, the local interaction $\mathsf{F}(q)$ is RM. Therefore, there is a Player-A GF-strategy $\sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}}^q$ that is optimal w.r.t. $(\mathsf{F}(q), m_q^{B_u})$.

Let us show that this GF-strategy is progressive w.r.t. $\mathsf{Sec}(\emptyset)$ (recall Definition 4.14). First, we show that σ_{A} is optimal in the game in normal form $\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle$. For all $q' \in Q$, we have:

- If $q' \in B_u$, then $m_q^{B_u}(q') = q_{\mathsf{init}}$ and $v_{\mathsf{F}(q), m_q^{B_u}}(q_{\mathsf{init}}) = \alpha_{\mathsf{F}(q), m_q^{B_u}} \geq u = \chi_{\mathcal{G}}[\mathsf{A}](q')$;
- Otherwise, we have $\chi_{\mathcal{G}}[\mathsf{A}](q') = v_{\mathsf{F}(q), m_q^{B_u}} \circ m_q^{B_u}(q')$.

For all $z \in [0,1] \setminus \{u\}$, we let $Q_z := \chi_{\mathcal{G}}[\mathsf{A}]^{-1}[\{z\}]$ and $Q_u := \chi_{\mathcal{G}}[\mathsf{A}]^{-1}[\{u\}] \setminus B_u$. That way, $Q \setminus B_u = \cup_{z \in [0,1]} Q_z$. Then, for all Player-$\mathsf{B}$ GF-strategies $\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}^q$, we have, denoting $y := \mathsf{out}[\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})$:

$$
\begin{aligned}
y &= \sum_{q' \in Q} \chi_{\mathcal{G}}[\mathsf{A}](q') \cdot \mathsf{out}[\langle \mathsf{F}(q), q' \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \\
&= u \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{B_u} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) + \sum_{z \in [0,1]} \sum_{q' \in Q_z} \chi_{\mathcal{G}}[\mathsf{A}](q') \cdot \mathsf{out}[\langle \mathsf{F}(q), q' \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \\
&\geq (u - \alpha_{\mathsf{F}(q), m_q^{B_u}}) + \alpha_{\mathsf{F}(q), m_q^{B_u}} \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{B_u} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \\
&+ \sum_{z \in [0,1]} \sum_{q' \in Q_z} v_{\mathsf{F}(q), m_q^{B_u}} \circ m_q^{B_u}(q') \cdot \mathsf{out}[\langle \mathsf{F}(q), q' \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \\
&= (u - \alpha_{\mathsf{F}(q), m_q^{B_u}}) + v_{\mathsf{F}(q), m_q^{B_u}}(q_{\mathsf{init}}) \cdot \mathsf{out}[\langle \mathsf{F}_{\mathsf{F}(q), m_q^{B_u}}(q_{\mathsf{init}}), q_{\mathsf{init}} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \\
&+ \sum_{z \in [0,1]} v_{\mathsf{F}(q), m_q^{B_u}}(z) \cdot \mathsf{out}[\langle \mathsf{F}_{\mathsf{F}(q), m_q^{B_u}}(q_{\mathsf{init}}), z \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \\
&= u - \alpha_{\mathsf{F}(q), m_q^{B_u}} + \mathsf{out}[\langle \mathsf{F}_{\mathsf{F}(q), m_q^{B_u}}(q_{\mathsf{init}}), v_{\mathsf{F}(q), m_q^{B_u}} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \\
&\geq u - \alpha_{\mathsf{F}(q), m_q^{B_u}} + \mathsf{val}[\langle \mathsf{F}_{\mathsf{F}(q), m_q^{B_u}}(q_{\mathsf{init}}), v_{\mathsf{F}(q), m_q^{B_u}} \rangle](\sigma_{\mathsf{A}}) \\
&= u - \alpha_{\mathsf{F}(q), m_q^{B_u}} + \mathsf{val}[\langle \mathsf{F}_{\mathsf{F}(q), m_q^{B_u}}(q_{\mathsf{init}}), v_{\mathsf{F}(q), m_q^{B_u}} \rangle] \\
&= u - \alpha_{\mathsf{F}(q), m_q^{B_u}} + \alpha_{\mathsf{F}(q), m_q^{B_u}} = u
\end{aligned}
$$

This second to last equality comes from the fact that the Player-A GF-strategy σ_{A} satisfies item 1. of Proposition 7.2. The last equality comes from Lemma 7.8 and the fact that $\mathsf{F}_{\mathsf{F}(q), m_q^{B_u}}(q_{\mathsf{init}}) = \mathsf{F}(q)^{m_q^{B_u}}$. Hence, the Player-$\mathsf{A}$ GF-strategy σ_{A} is optimal in the game in normal form $\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle$, since $\mathsf{val}[\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle][\mathsf{A}] = \chi_{\mathcal{G}}[\mathsf{A}](q) = u$ by Proposition 3.9.

Furthermore, since the Player-A GF-strategy also ensures item 2. of that same Proposition 7.2, letting $\delta := \inf_{\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}^q} p_{\mathsf{F}(q), m_q^{B_u}}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})$, we have $\delta > 0$. Recall that, for all $\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}^q$, we have

$$p_{\mathsf{F}(q), m_q^{B_u}}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) = \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q \setminus B_u} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})$$

We let $V := \chi_{\mathcal{G}}[\mathsf{A}][Q] \subseteq [0,1]$, $V_{<u} := \{z \in V \mid z < u\}$ and $V_{\geq u} := V \setminus V_{<u}$. We also let $x_u := \max V_{<u} < u$ and $y := \frac{u - x_u}{1 - x_u} > 0$. Consider any Player-$\mathsf{B}$

GF-strategy $\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}^q$. We have:

$$\begin{aligned} u &\leq \mathsf{out}[\langle \mathsf{F}(q), \chi_{\mathcal{G}}[\mathsf{A}] \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \\ &= u \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{B_u} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) + \sum_{z \in [0,1]} z \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q_z} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \end{aligned}$$

Therefore:

$$u \cdot p_{\mathsf{F}(q), m_q^{B_u}}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \leq \sum_{z \in [0,1]} z \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q_z} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})$$

We then have:

$$\begin{aligned} u \cdot p_{\mathsf{F}(q), m_q^{B_u}}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) &\leq \sum_{z \in [0,1]} z \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q_z} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \\ &= \sum_{z \in V_{<u}} z \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q_z} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) + \sum_{z \in V_{\geq u}} z \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q_z} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \\ &\leq \sum_{z \in V_{<u}} x_u \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q_z} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) + \sum_{z \in V_{\geq u}} \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q_z} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \\ &= x_u \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{\cup_{z \in V_{<u}} Q_z} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) + \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{\cup_{z \in V_{\geq u}} Q_z} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \end{aligned}$$

Hence, letting $p_{\geq u} := \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{\cup_{z \in V_{\geq u}} Q_z} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}})$, we have $\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{\cup_{z \in V_{<u}} Q_z} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) = p_{\mathsf{F}(q), m_q^{B_u}}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) - p_{\geq u}$. Therefore:

$$p_{\geq u} \geq p_{\mathsf{F}(q), m_q^{B_u}}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \cdot \frac{u - x_u}{1 - x_u} \geq \delta \cdot y$$

In addition, note that, for all $z \in V_{\geq u}$, we have $Q_z \subseteq \mathsf{Sec}(\emptyset)$ by definition of u. Hence $\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{\mathsf{Sec}(\emptyset)} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \geq p_{\geq u}$. It follows that we have:

$$\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{\mathsf{Sec}(\emptyset)} \rangle](\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) \geq \delta \cdot y$$

This holds for all $\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}^q$. Hence, recalling Definition 4.16, it follows that the GF-strategy σ_{A} is progressive w.r.t. $\mathsf{Sec}(\emptyset)$ (i.e. $\sigma_{\mathsf{A}} \in \mathsf{Prog}_q(\mathsf{Sec}(\emptyset))$) since $\inf_{\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}^q} p^q_{\mathsf{Sec}(\emptyset)}(\sigma_{\mathsf{A}}, \sigma_{\mathsf{B}}) = \delta \cdot y > 0$. Hence the contradiction with the fact that $q \notin \mathsf{Sec}(\emptyset)$, recall Definition 4.11. □

8 - Standard game forms for parity objectives

In this chapter, we characterize the standard finite game forms ensuring the existence of positional optimal strategies for both players in finite parity games. This is stated as an NSC-transfer in Corollary 8.10. The main difficulty consists in proving Theorem 8.3, which states that if all standard finite local interactions of a finite concurrent parity game are individually well-behaved (in a specific sense, the formal definition is given in Definition 8.7), then both players have positional optimal strategies. This chapter is almost entirely devoted to the proof of this theorem.

Contrary to what we did in the previous chapter, the result we show deals with both players at the same time. Indeed, as mentioned above, we characterize the game forms ensuring the existence of positional optimal strategies for both players. The reason why we do that is the same reason why we only consider standard finite game forms instead of arbitrary game forms: we want to manipulate only positional strategies for the players. If we only had an assumption for one player (as in the previous chapter), we would also need to consider infinite-choice strategies for the other player, a priori. Since positional strategies suffice, the difficulty of the proof lies mainly on the intricate nature of the parity objective itself.

The core of this chapter gives informal explanations of the definitions and statements (with proof sketches) used to prove Theorem 8.3. Additionally, we illustrate the definitions and statements on examples and, once all the relevant definitions are given, provide a big picture of the proof in Subsection 8.4.1. We proceed that way because the technical details are quite heavy and may obfuscate the underlying ideas behind the proof. We give all technical details in Section 8.6.

Finally, note that since we only consider standard finite local interactions, the concurrent parity games we consider in this chapter all have a value (by Theorems 2.3 and 1.11).

This work is not yet published, but will be resubmitted soon.

8.1 Dominating and guaranteeing a valuation

As mentioned above, in this section, we will only manipulate positional strategies. Furthermore, as stated in Observation 1.3, in a concurrent game, two positional strategies induce a Markov chain. Furthermore, in the following, we will be especially interested in the BSCCs that can occur in the Markov chain induced by two positional strategies. We define the relevant notions for us below.

Definition 8.1 (BSCCs compatible with a strategy). *Consider a standard finite parity game $\mathcal{G} = \langle \mathcal{C}, \mathsf{col} \rangle$. Let $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ be a positional Player-*A *strategy. We let $H_{\mathsf{s_A}}$ denote the set of BSCCs* compatible *with $\mathsf{s_A}$, i.e. the BSCCs of some Markov chain $\mathcal{T}^{\mathsf{s_A},\mathsf{s_B}}$, where $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$ ranges over Player-*B *positional deterministic strategies.*

*This is analogous for a Player-*B *strategy $\mathsf{s_B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$.*

A BSCC $H \in H_{\mathsf{s_A}}$ is even-colored *if $\max \mathsf{col}[H]$ is even. Otherwise, it is* odd-colored.

A subset of states $S \subseteq Q$ occurs *in a BSCC H if $H \cap S \neq \emptyset$. A state $q \in Q$* occurs *in a BSCC H if $q \in H$.*

The reason why we only consider positional deterministic strategy for Player B given a Player-A positional strategy is because, once such a strategy is fixed, we obtain a finite MDP. Furthermore, positional deterministic strategies are enough to play optimally in finite MDPs with parity objectives [27].

Recall Definition 3.7: we defined the notion of strategy dominating a valuation $v : Q^+ \to [0, 1]$. In the context of this chapter, we only consider positional strategies and valuations $v : Q \to [0, 1]$. Hence, we recall the notion of domination in this context, along with the notion of guaranteeing a valuation (recall Definition 3.2). We also define a stronger notion that dominating a valuation: parity dominating a valuation.

Definition 8.2 (Parity dominating a valuation). *Let $\mathcal{G}$ be a standard finite concurrent parity game and $v : Q \to [0, 1]$ be a valuation over its states. Consider a positional Player-*A *strategy $\mathsf{s_A}$ (resp. Player* B *strategy $\mathsf{s_B}$). The strategy $\mathsf{s_A}$ (resp. $\mathsf{s_B}$):*

- dominates *the valuation v if for all $q \in Q$, it holds that $v(q) \leq \mathsf{val}[\langle \mathsf{F}(q), v \rangle](\mathsf{s_A}(q))$ (resp. $v(q) \geq \mathsf{val}[\langle \mathsf{F}(q), v \rangle](\mathsf{s_B}(q))$);*
- parity dominates *the valuation v if it dominates v and all BSCCs H compatible with $\mathsf{s_A}$ (resp. $\mathsf{s_B}$) such that $\min v[H] > 0$ (resp. $\max v[H] < 1$) are even-colored (resp. odd-colored);*
- guarantees *the valuation v if, for all $q \in Q$, it holds $v(q) \leq \chi_\mathcal{G}[\mathsf{s_A}](q)$ (resp. $v(q) \geq \chi_\mathcal{G}[\mathsf{s_B}](q)$).*

Parity dominating a valuation implies guaranteeing it. This is a direct consequence of Corollary 3.16.

Proposition 8.1 (Proof 8.6.2). *Consider a standard finite concurrent game $\mathcal{G}$, a Player-*A *positional strategy $\mathsf{s_A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ and a valuation $v : Q \to [0, 1]$. If the strategy $\mathsf{s_A}$ dominates v, then for all BSCCs $H \in H_{\mathsf{s_A}}$, there is $v_H \in [0, 1]$ such that $v[H] = \{v_H\}$. If in addition $\mathsf{s_A}$ parity dominates v, it also guarantees v.*

Finally, let us introduce a few notations we will use throughout because we are tackling the parity objective with arbitrarily many colors.

Definition 8.3. *As stated in Section 1.1, for all* $(i, j) \in \mathbb{N}^2$*, we denote by* $[\![i, j]\!] := \{k \in \mathbb{N} \mid i \leq k \leq j\}$ *the set of integers between* i *and* j*. However, in this chapter only, for convenience, we assume that this set is typed in the sense that the integers in* $[\![i, j]\!]$ *are not seen as real numbers. In particular,* $[\![0, 1]\!] \cap [0, 1] = \emptyset$*. Furthermore, for all* $e \in \mathbb{N}$*, we let* $K_e := \{k_i \mid i \in [\![0, e]\!]\}$*.*

In addition, for all finite subsets $S \subseteq \mathbb{N}$*, we let* $\mathsf{Even}(S)$ *(resp.* $\mathsf{Odd}(S)$*) be the smallest even (resp. odd) integer that is greater than or equal to all elements in* S*. Similarly, for all* $n \subseteq \mathbb{N}$*, we let* $\mathsf{Even}(n)$ *(resp.* $\mathsf{Odd}(n)$*) be the smallest even (resp. odd) integer that is greater than or equal to* n*.*

8.2 Local Environment

The goal of this section is to define simple parity games with a single non-trivial local interaction. This will allow us to define exactly the game forms that should be used in parity games if one requires positional optimal strategies for both players. We first define what a (parity) environment on a given set of outcomes is. We can then define the parity game induced by such an environment (along with a game form).

Definition 8.4 (Parity environment)**.** *Consider a non-empty finite set of outcomes* O*. An* environment E *on* O *is a tuple* $E := \langle c, e, p \rangle$ *where* $c, e \in \mathbb{N}$ *with* $c \leq e$ *and* $p : \mathsf{O} \to \{q_{\mathsf{init}}\} \uplus K_e \uplus [0, 1]$*. We let* $p_{[0,1]} := p[\mathsf{O}] \cap [0, 1]$*. The* size *w.r.t. Player* A *(resp.* B*)* $\mathsf{Sz}_{\mathsf{A}}(E)$ *(resp.* $\mathsf{Sz}_{\mathsf{B}}(E)$*) of the environment* E *is equal to* $\mathsf{Sz}_{\mathsf{A}}(E) := \mathsf{Even}(e) - c$ *(resp.* $\mathsf{Sz}_{\mathsf{B}}(E) := \mathsf{Odd}(e) - c$*). We denote by* $\mathsf{Env}(\mathsf{O})$ *the set of all environments on the set of outcomes* O*.*

Remark 8.1. *A quick note on the size of an environment. From Player* A*'s perspective, the size of an environment* $E := \langle c, e, p \rangle$*, assuming that* $c = 0$*, is equal to: 0 if* $e = 0$*, 2 if* $e = 1$ *or* $e = 2$*, etc. Informally, when* $c = 0$*, if* $e = 0$ *this corresponds to safety game, if* $e = 1$ *this corresponds to a co-Büchi game, etc. This will become more apparent with Definition 8.5 below.*

We can then define the simple parity game corresponding to a parity environment.

Definition 8.5 (Parity game induced by an environment)**.** *Consider a non-empty finite set of outcomes* O*, a standard finite game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ *and an environment* $E = \langle c, e, p \rangle \in \mathsf{Env}(\mathsf{O})$*. Let* $Y := (\mathcal{F}, E)$*. The local arena* $\mathcal{C}_Y = \langle Q, \mathsf{F}, \mathsf{K}, \mathsf{col} \rangle$ *induced by* Y *is such that:*

- $Q := \{q_{\mathsf{init}}\} \cup K_e \cup p_{[0,1]}$*,* $Q_{\mathsf{s}} := p_{[0,1]}$*, and for all* $u \in p_{[0,1]}$*, we set the value of the stopping state* u *to be* u *itself:* $\mathsf{val}(u) \leftarrow u$*;*

- $\mathsf{F}(q_{\mathsf{init}}) := \mathcal{F}^p = \langle \mathsf{Act}_\mathsf{A}, \mathsf{Act}_\mathsf{B}, Q, \mathbb{E}_p(\varrho) \rangle_\mathsf{s}$, *and for all* $i \in [\![0, e]\!]$, *the set* k_i *is trivial, its only outcome is the state* q_{init}*;*

- $\mathsf{K} := [\![0, e]\!]$, $\mathsf{col}(q_{\mathsf{init}}) := c$ *and for all* $i \in [\![0, e]\!]$, *we have* $\mathsf{col}(k_i) := i$.

For all $u \in [0, 1]$, *we denote by* $v_Y^u : Q \to [0, 1]$ *the valuation such that:* $v_Y^u(q_{\mathsf{init}}) = v_Y^u(k_i) := u$ *for all* $i \in [\![0, e]\!]$ *and* $v_Y^u(x) := x$ *for all* $x \in p_{[0,1]}$. *Furthermore, for all Player-*A GF*-strategies* $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathcal{F})$, *we denote by* $\mathsf{s}_\mathsf{A}^Y(\sigma_\mathsf{A}) \in \mathsf{S}_\mathsf{A}^{\mathcal{C}_Y}$ *the Player-*A *positional strategy defined by* σ_A *in the arena* $\mathcal{C}_Y$.

The game $\mathcal{G}_Y$ *is then equal to* $\mathcal{G}_Y := \langle \mathcal{C}_Y, \mathsf{Parity}_\mathsf{K} \rangle$.

Example 8.1. *Definition 8.5 above is illustrated in Figures 8.1. Note that the colors of the non-stopping states are depicted in red next to the states.*

As in the previous chapter, we can define the notion of Player-A GF-strategy being optimal in a parity environment.

Definition 8.6 (Optimal GF-strategies). *Consider a non-empty finite set of outcomes* O*, a game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$*, an environment* $E = \langle c, e, p \rangle \in \mathsf{Env}(\mathsf{O})$*, and let* $Y := (\mathcal{F}, E)$*. A Player-*A GF*-strategy* $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathcal{F})$ *is said to be* optimal *w.r.t.* Y *if the Player-*A *positional strategy* $\mathsf{s}_\mathsf{A}^Y(\sigma_\mathsf{A})$ *is optimal in* $\mathcal{G}_Y$*. The definition is analogous for Player* B.

Given a finite set of outcomes O, we can now define the game forms on O ensuring the existence of optimal strategies w.r.t. all environments.

Definition 8.7 (Game forms with optimal strategies). *Consider a non-empty finite set of outcomes* O*, a standard finite game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ *and some* $n \in \mathbb{N}$.

Consider a Player $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$*. The game form* $\mathcal{F}$ *is said to be* positionally maximizable up to n *w.r.t. Player* C *if, for each environment* $E \in \mathsf{Env}(\mathsf{O})$ *with* $\mathsf{Sz}_\mathsf{C}(E) \leq n$*, there is an optimal* GF*-strategy for Player* C *w.r.t.* $(\mathcal{F}, E)$.

When this holds for both Players, $\mathcal{F}$ *is said to be* positionally optimizable up to n*. The corresponding set of game forms is denoted* $\mathsf{ParO}(n)$*. If this holds for all* $n \in \mathbb{N}$*,* $\mathcal{F}$ *is simply said to be* positionally optimizable, *and the corresponding set of game forms is denoted* ParO.

Remark 8.2. *First, note that all standard finite game forms are positionally optimizable up to 0. This is because environement of size 0 induce safety games. However, there are some standard finite game forms that are not positionally maximizable w.r.t. any player up to 1. This is for instance the case of the game form of Figure 7.1.*

Furthermore, by definition, from a game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ *that is not positionally optimizable up to some* $n \in \mathbb{N}$*, there exists an environment* $E \in \mathsf{Env}(\mathsf{O})$ *such that either of the players has no positional optimal strategy in the simple parity game* $\mathcal{G}_{(\mathcal{F},E)}$ *where the difference between* $\mathsf{col}(q_{\mathsf{init}})$ *and the maximum of the colors appearing in* $\mathcal{G}_{(\mathcal{F},E)}$ *is at most* n.

In the game $\mathcal{G}_{(\mathcal{F},E)}$ depicted on the right of Figure 8.1, Player A has positional optimal strategies: it suffices to play both rows with positive probability. (This is similar for Player B.) As a side remark, the game form in the left of Figure 8.1 is positionally optimizable.

In Lemma 8.2 below, we formulate more explicitly (using the notion of parity domination from Definition 8.2) what optimal GF-strategies are.

Lemma 8.2 (Proof 8.6.3)**.** *Consider a non-empty finite set of outcomes* O*, a standard finite game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$*, an environment* $E = \langle c, e, p \rangle \in \mathsf{Env}(\mathsf{O})$ *and* $Y := (\mathcal{F}, E)$*. A Player-*A GF*-strategy* $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathcal{F})$ *is* optimal *w.r.t.* Y *if and only if, letting* $u := \chi_{\mathcal{G}_Y}(q_{\mathsf{init}})$*, either (i)* $u = 0$*, or (ii) the positional Player-*A *strategy* $\mathsf{s}_\mathsf{A}^Y(\sigma_\mathsf{A})$ *parity dominates the valuation* v_Y^u*.*

*Furthermore (ii) is equivalent to: (1) the Player-*A *positional strategy* $\mathsf{s}_\mathsf{A}^Y(\sigma_\mathsf{A})$ *dominates the valuation* v_Y^u*, i.e.* σ_A *is optimal in the game in normal form* $\langle \mathcal{F}, v_Y^u \circ p \rangle$ *and (2) for all* $b \in \mathsf{Act_B}$*, if* $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[0,1]} \rangle](\sigma_\mathsf{A}, b) = 0$ *(i.e. the probability under* σ_A *and* b *to reach a stopping state is null), then* $\max(\mathsf{Color}(\mathcal{F}, p, \sigma_\mathsf{A}, b) \cup \{c\})$ *is even where* $\mathsf{Color}(\mathcal{F}, p, \sigma_\mathsf{A}, b) := \{i \in [\![0, e]\!] \mid \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[k_i]} \rangle](\sigma_\mathsf{A}, b) > 0\}$ *is the set of colors that can be seen with positive probability under* σ_A *and* b*. This is symmetrical for Player* B*.*

Remark 8.3. *Informally, this proposition states that for a Player-*A GF*-strategy* σ_A *to be optimal in a simple game* $\mathcal{G}_Y$ *with positive value, it must be the case that for every Player-*B *action* b*: either there is a positive probability (w.r.t.* σ_A *and* b*) to exit* q_{init} *and the expected value of the stopping states visited is at least* u*; or the game loops on* q_{init} *with probability* 1*, and the maximum of the colors that can be seen with positive probability (w.r.t.* σ_A *and* b*) is even. In particular, if* $c \leq \max \mathsf{Color}(\mathcal{F}, p, \sigma_\mathsf{A}, b)$ *or if* c *is odd, then* $\max \mathsf{Color}(\mathcal{F}, p, \sigma_\mathsf{A}, b)$ *is even.*

8.3 The main theorem

The goal of this section is to formally state the main theorem of this chapter, Theorem 8.3 below. Informally, this theorem consists in extracting, for every state of a game and for each Player, a local environment which will summarize the context of the state to the Player, and tell her how to play optimally (and positionally).

Before stating this theorem, let us recall and define below some useful notations, in particular we recall the notation for value slices, a notion we have already used in Chapter 3.

Definition 8.8 (Value slice)**.** *Consider a standard finite parity game* $\mathcal{G}$*. For all subsets of states* $S \subseteq Q$*, we denote by* $V_S := \{u \in [0,1] \mid \exists q \in S,\ \chi_\mathcal{G}(q) = u\}$ *the finite set of values of states in* S*. Furthermore, for all* $u \in V_Q$*, we let* $Q_u := \{q \in Q \mid \chi_\mathcal{G}(q) = u\}$ *be the set of states whose value is* u*: it is*

$$\mathcal{F} := \begin{bmatrix} y & x \\ x & y \end{bmatrix}$$

- $E := \langle 2, 4, p \rangle$;
- $p(x) := 1/2$;
- $p(y) := k_3$.

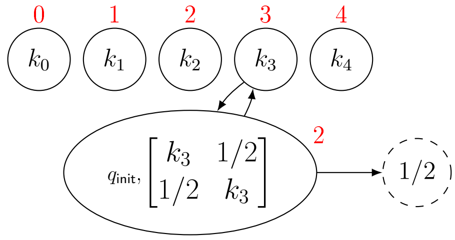

Figure 8.1: On the left, a game form on the set of outcomes $\mathsf{O} := \{x, y\}$, in the middle the description of an environment on O and on the right the parity game $\mathcal{G}_{(\mathcal{F},E)}$ obtained from what is depicted on the left. The dashed state is a stopping of value 1/2.

the u-slice of $\mathcal{G}$. Finally, for all $u \in V_Q$, we let $e_u := \mathsf{Even}(\mathsf{col}[Q_u])$ and $o_u := \mathsf{Odd}(\mathsf{col}[Q_u])$.

We also introduce the notion of positional strategies generated by an environment function: this is a strategy that, at each state, plays a GF-strategy that is optimal in the corresponding environment.

Definition 8.9 (Positional strategy generated by an environment function)**.** *For all environment functions $\mathsf{Ev} : Q \to \mathsf{Env}(\mathsf{O})$, a Player-$\mathsf{A}$ positional strategy $\mathsf{s_A}$ is* generated *by Ev if for all $q \in Q$, the GF-strategy $\mathsf{s_A}(q)$ is optimal w.r.t. $(\mathsf{F}(q), \mathsf{Ev}(q))$ (and similarly for Player B).*

We can now state the main result of this chapter: Theorem 8.3.

Theorem 8.3. *Let $\mathcal{G} = \langle \mathcal{C}, \mathsf{Parity_K} \rangle$ be a standard finite parity game. Assume that for all $q \in Q$, the game form $\mathsf{F}(q)$ is positionally maximizable up to $e_{\chi_\mathcal{G}(q)} - \mathsf{col}(q)$ w.r.t. Player A and positionally maximizable up to $o_{\chi_\mathcal{G}(q)} - \mathsf{col}(q)$ w.r.t. Player B. Then, there is a function $\mathsf{Ev_A} : Q \to \mathsf{Env}(\mathsf{O})$ (resp. $\mathsf{Ev_B}$) such that all Player-A (resp. Player-B) positional strategies $\mathsf{s_A}$ (resp. $\mathsf{s_B}$) generated by $\mathsf{Ev_A}$ (resp. $\mathsf{Ev_B}$) are optimal in $\mathcal{G}$; and such strategies exist.*

Remark 8.4. *Given some $u \in V_Q$, one can realize that the requirement at states $q, q' \in Q_u$ changes depending on the color of q and q'. More specifically, if $\mathsf{col}(q) < \mathsf{col}(q')$, then the requirement at state q is (a priori) stronger than the requirement at state q' since the game form $\mathsf{F}(q)$ should behave well for environments of larger size than the game form $\mathsf{F}(q')$.*

The remainder of this chapter is almost-entirely devoted to the explanation of the construction of the environment function $\mathsf{Ev_A}$ (the construction being similar for Player B), hence we are taking the point of view of Player A. First, let us argue that we can restrict ourselves to a specific u-slice Q_u for some $u \in V_Q$. Such a restriction is properly defined (using stopping states) in Definition 8.10 below.

Definition 8.10 (Game restricted to a u-slice). *For all $u \in V_Q$, we let $\mathcal{G}^u$ be the concurrent game $\mathcal{G}$ where all states outside Q_u are made stopping states: for every $q \in Q \setminus Q_u$, we set $\mathsf{val}(q) \leftarrow \chi_{\mathcal{G}}(q)$. The states, game forms and coloring function on Q_u are left unchanged.*

Interestingly, a Player-A positional strategy optimal in $\mathcal{G}$ can be obtained by merging appropriate positional strategies $\mathsf{s}_{\mathsf{A}}^u$ in the games $\mathcal{G}^u$ for all $u \in V_Q \setminus \{0\}$, which is actually a straightforward consequence of Proposition 8.1.

Lemma 8.4 (Proof 8.6.4). *For all $u \in V_Q \setminus \{0\}$, consider a positional Player-A strategy $\mathsf{s}_{\mathsf{A}}^u$ that parity dominates the valuation $\chi_{\mathcal{G}}$ in $\mathcal{G}^u$. Then, the Player-A positional strategy s_{A} such that $\mathsf{s}_{\mathsf{A}}(q) := \mathsf{s}_{\mathsf{A}}^u(q)$ for all $u \in V_Q \setminus \{0\}$ and $q \in Q_u$ guarantees the valuation $\chi_{\mathcal{G}}$ in $\mathcal{G}$ (i.e. it is optimal).*

8.4 The proof

In Appendix 8.6.5, we give a quick overview of the technical properties and lemmas we show in the appendix that we will use to prove the lemmas (we will state later in this chapter) leading to the proof of Theorem 8.3.

In this section, we fix a standard finite parity game $\mathcal{G} = \langle \mathcal{C}, \mathsf{Parity}_{\mathsf{K}} \rangle$. In particular, the set of states Q is fixed and recall that Q is also the set of outcomes of all game forms occurring in $\mathcal{G}$. Also, Lemma 8.4 above justifies that we focus on a given u-slice Q_u for some positive $u \in (0, 1]$. We also let $e := e_u$, $o := o_u$ and $K := K_e = \{k_i \mid i \in [\![0, e]\!]\}$ and for all $n \in [\![0, e]\!]$, we let $K^n := \{k_i^n \mid i \in [\![0, e]\!]\}$.

8.4.1 . Big picture of the proof

In order to give an idea of the steps we take to prove Theorem 8.3, let us first consider the very simple case of finite turn-based deterministic reachability games. Computing the area L_{A} from which Player A wins can be done inductively. That is, initially we set $L_{\mathsf{A}} := T$ where T denotes the target that Player A wants to reach. Then, the inductive step is handled with a (deterministic) attractor: we add to L_{A} any Player-A state with a successor in L_{A} and any Player-B state with all successors in L_{A}. After finitely many steps, there is no more state to add in L_{A}: this exactly corresponds to the states from which Player A has a winning strategy.

Computing a single attractor is not merely enough to take into account the intricate behavior of parity objectives, which is what Theorem 8.3 deals with. Therefore, we are going to iteratively compute several layers of (virtual) colors, with a local update to change the (virtual) color (and therefore the layer it belongs to) of a state. This local update can be seen as an attractor except in a concurrent stochastic setting. Hence, when we update the (virtual) color of a state, we take into account the concurrent interaction of the players at each state along with the probability to see stopping states or states with

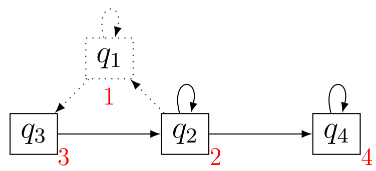

Figure 8.2: A (deterministic turn-based) game.

different (virtual) colors. We define this local update in Subsection 8.4.3. Let us describe below the steps we take to capture the behavior of the parity objective.

We compute layers of successive probabilistic attractors with leaks towards the stopping states. Although we compute a strategy, e.g., for Player A, we alternate players to build layers, then move the last non-empty layer into the closest layer with same parity, then backtrack the attractor computation from this layer downwards, and start over again the full attractor computation on the new layer structure. In a more concrete way, let us assume below that the highest color in the u-slice is 6. We proceed as follows:

1. Add the states colored with 6 to layer L_6.

2. Recursively add to L_6 the states where Player A can guarantee that with positive probability (pp) either a leak towards stopping states occurs now and its expected explicit value is u or more ($\mathsf{Leak}_{\geq u}$), or with pp the next state is in L_6, i.e. with pp color 6 or $\mathsf{Leak}_{\geq u}$ will occur.

3. Add the remaining states colored with 5 to layer L_5.

4. Recursively add to L_5 the states where Player B can guarantee that either $\mathsf{Leak}_{<u}$ occurs now with pp, or the next state is surely not in L_6 and with pp in L_5; i.e. if $\mathsf{Leak}_{<u}$ occurs with probability 0, then color 6 will not occur but color 5 will eventually occur with pp.

5. Add the remaining states colored with 4 to layer L_4.

6. Recursively add to L_4 the states where Player A can guarantee that either $\mathsf{Leak}_{\geq u}$ will occur with pp, or the maximal layer index of the next states seen with pp is 4 or 6, i.e. even and at least 4.

7. And so on, from color 3 to 0. The layers so far only give information about what can happen at finite horizon.

8. For instance, from L_2, Player A can guarantee that either $\mathsf{Leak}_{\geq u}$ will occur with pp, or the maximal color that will be seen with pp is in $\{2, 4, 6\}$. Now, if e.g. $L_0 = L_1 = \emptyset$, we merge L_2 into L_4. This is,

arguably, the most surprising step, let us try to give a conceptual intuition behind this step, we will then illustrate it on a concrete example. Consider what happens in the L_2 layer, assuming $L_0 = L_1 = \emptyset$. From states in that layer, either:

- $\mathsf{Leak}_{\geq u}$ occurs with pp;
- otherwise, $L_3 \cup L_4 \cup L_5 \cup L_6$ occurs with pp, and the maximum index seen with pp is even;
- otherwise, the game loops surely in L_2, and with pp the maximum color seen is 2. In that case, Player A wins the (real) parity game almost-surely.

In other words, either the game loops in L_2 and Player A wins, or what happens is alike to what happens in the layer L_4. In Figure 8.2 without dotted state q_1 where Player B plays alone, $L_4 = \{q_4\}$, then $L_3 = \{q_3\}$, then $L_2 = \{q_2\}$, but if the play stays in L_2, Player A wins, so Player B may just as well go to q_4 since there is no L_0, L_1 below to escape. However, in Figure 8.2 with dotted state q_1, Player B can leave L_2 via $L_1 \neq \emptyset$, avoid color 4, and win. There L_2 and L_4 are not alike. Hence, we do not merge them.

9. Earlier, some states of color 1 may have been added to L_3 since L_3 "overruled" L_2 w.r.t indices, but since L_2 has just been merged into L_4, it now overrules L_3, so some states from L_3 may have to go back to L_1. Therefore we reset the layers below L_4 and repeat the above attractor alternation all over again, until all the states are eventually in L_6 as we shall prove.

The key property that is growing throughout the above computation and will hold in the final L_6 involves layer games: the L_n-game is derived from the u-slice by abstracting each L_i with $i \neq n$ via one state k_i^n from which the player who dislikes the parity of n chooses any next state in L_n, making it harder for the other to win. If $i > n$ then k_i^n is i-colored, else $(n-1)$-colored, also making it harder for the other to win. And states in L_n bear their true colors. See for instance Figure 8.5. The L_n game is only seemingly harder to win: it is actually equivalently hard, but its useful properties are easier to prove.

The key growing property is as follows: between two merges, the attractor computation from the top layer down to L_n ensures that Player A has a positional strategy of value at least u in each L_i for even $i \geq n$, and Player B less than u for odd $i \geq n$. In the very end, there is only one even layer with all states bearing their true colors, and no abstract states, so the last layer game equals the u-slice game, for which we have thus computed a positional optimal strategy.

Let us hint at how to show positional optimality in the L_n-games when it holds: we break L_n each into one simple parity game built on $\mathsf{F}(q)$ per state q in L_n, abstracting the other states in L_n into one. Our theorem assumption yields an optimal GF-strategy for Player A or B in the simple parity game. Gluing them does the job.

Now recall the above example of computation, where Player A for even indices, and Player B for odd ones, guarantees (assuming neither $\mathsf{Leak}_{\geq u}$ nor $\mathsf{Leak}_{<u}$ occurs):

- in L_6, that a next state is of color 6 with pp;
- in L_5, that the next state is surely not of color 6, and with pp a next state is of color 5;
- in L_4, that a next state is of color 6 with pp or surely the next state is not of color 5, and with pp a next state is of color 4;
- in L_3, that the next state is surely not of color 6 and that the next state is of color 5 with pp or surely the next state is not of color 4, and with pp a next state is of color 3; etc.

One can realize that the furthest the layer is from the maximal color (that is 6), the more complex the requirement is at that layer. That is why the strength of our assumption on the game form induced at some state increases with the difference between maximal true color in the u-slice and the true color (at most n) of the state, as stated in Remark 8.4. We will discuss this further in the next chapter. In particular, we will show that there is an infinite strict hierarchy between these assumptions.

8.4.2 . Extracting an environment function from a parity game

Once restricted to the game $\mathcal{G}^u$, the method we use to prove Theorem 8.3 consists in iteratively building a pair of (virtual) coloring and environment functions ensuring a nice property (namely faithfulness, defined below in Definition 8.17). For the remainder of this section, we illustrate the definitions and lemmas on the game depicted in Figures 8.3 and 8.4.

Example 8.2. *We explain the notations used to depict this game (it is in fact the same arena in both Figures 8.3 and 8.4, with different coloring functions – real or virtual). On the sides in green are the slices $Q_0, Q_{1/4}, Q_{3/4}$ and Q_1 from left to right. We focus on the central slice $Q_{1/2}$. In $Q_{1/2}$, there are seven states, five of which (the square-shaped ones) are turn-based for Player* B*, that is, Player* A *has only one available action. On the other hand, the two circled-shaped states q_0 and q_5 are "truly" concurrent in the sense that both players have several actions available. Furthermore, note that only one stochastic outcome (of a local interaction) is drafted as a black dot: from q_4, Player* B

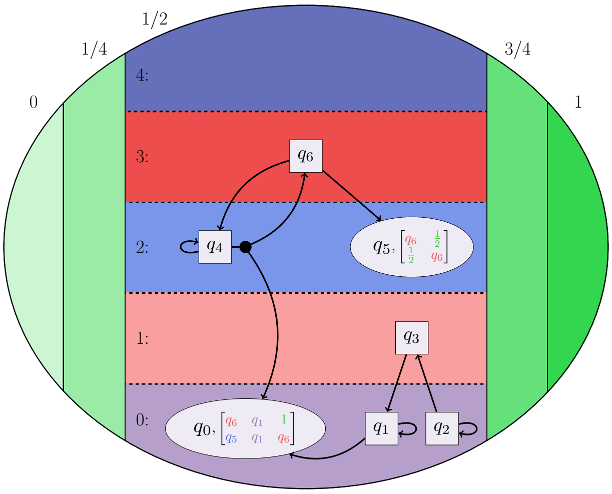

Figure 8.3: The depiction of a game restricted to the 1/2-slice $Q_{1/2}$ with the initial coloring function col.

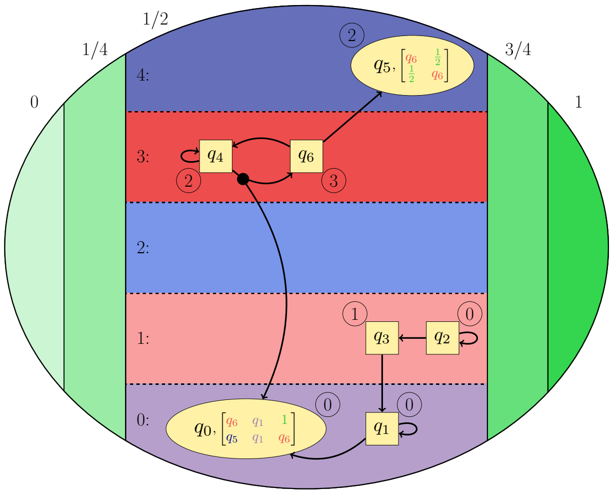

Figure 8.4: The same game restricted to $Q_{1/2}$ with a different coloring function vcol.

may either loop on q_4 or go with equal probability to q_0 and q_6. The other arrows lead to a single state and the outcomes of the game forms in q_0 or q_5 is a single state or a value: 1 or 1/2. These formally refer to a (distribution over) stopping states outside of the 1/2-slice $Q_{1/2}$. The horizontal layers depict the colors of the states. In Figure 8.3, the coloring function considered is the initial one col *whereas in Figure 8.4 we have depicted a (virtual) coloring function* vcol. *For instance,* $\mathsf{col}(q_6) = 3$ *whereas* $\mathsf{col}(q_5) = 2$. *Similarly,* $\mathsf{vcol}(q_6) = 3$ *whereas* $\mathsf{vcol}(q_5) = 4$. *Note that, in Figure 8.4, the real colors (given by* col*) are reminded next to some states with circled numbers. Finally, note that* $e := e_{1/2} = 4$.

Given a (virtual) coloring function, we need to extract local environments from the parity game $\mathcal{G}$, which summarize how the Players see their neighboring states via the virtual coloring function. This is (partly) done in Definition 8.11.

Definition 8.11 (Probability function extracted from an arena and a (virtual) coloring function)**.** *Consider a virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$. *For all* $n \in [\![0, e]\!]$, *we let* $Q_n := \mathsf{vcol}^{-1}[n]$. *We then define the map* $p_{n,\mathsf{vcol}} : Q \to \mathcal{D}(Q_n \uplus K^n \uplus V_{Q \setminus Q_u})$ *such that, for all* $q \in Q$:

- *if* $q \in Q_n$, $p_{n,\mathsf{vcol}}(q) := q$;
- *if* $q \in Q_u \setminus Q_n$, $p_{n,\mathsf{vcol}}(q) := k^n_{\mathsf{vcol}(q)}$;
- *if* $q \in Q \setminus Q_u$, $p_{n,\mathsf{vcol}}(q) := \chi_{\mathcal{G}}(q)$.

Given a (virtual) coloring function $\mathsf{vcol} : Q_u \to [\![0, e]\!]$ and a color $n \in [\![0, e]\!]$, we can now extract a smaller parity game from $\mathcal{G}$ where the states with non-trivial game forms are the states in $\mathsf{vcol}^{-1}[n]$, the states in $Q \setminus Q_u$ are stopping states and the arena loops back to $\mathsf{vcol}^{-1}[n]$ when a state in $Q_u \setminus \mathsf{vcol}^{-1}[n]$ is seen. This is done below in Definition 8.12.

Definition 8.12 (Parity game extracted from the u-slice)**.** *Consider a virtual coloring function $\mathsf{vcol} : Q_u \to [\![0, e]\!]$ and a color $n \in [\![0, e]\!]$. Let $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$ be a Player: A if n is odd and B if n is even. The arena $\mathcal{C}^n_{\mathsf{vcol}} = \langle Q', \mathsf{F}', \mathsf{K}, \mathsf{vcol}_n \rangle$ is such that, denoting $Q_n := \mathsf{vcol}^{-1}[n]$:*

- *$Q' := Q_n \cup K^n \cup V_{Q \setminus Q_u}$ where all $x \in V_{Q \setminus Q_u}$ are stopping states with $\mathsf{val}(x) \leftarrow x$;*
- *for all $q \in Q_n$, $\mathsf{F}'(q) := \mathsf{F}(q)^{p_{n,\mathsf{vcol}}} = \langle \mathsf{Act}^q_\mathsf{A}, \mathsf{Act}^q_\mathsf{B}, Q', \mathbb{E}_{p_{n,\mathsf{vcol}}} \rangle_\mathsf{s}$ and for all $k \in K^n$, we set $\mathsf{F}'(k)$ as a Player-C state where the outcomes are all the states in Q_n;*
- *for all $q \in Q_n$, we let $\mathsf{vcol}_n(q) := \mathsf{col}(q)$ and for all $i \in [\![0, e]\!]$, we have $\mathsf{vcol}_n(k^n_i) := \max(i, n-1)$.*

For $t \in [0, 1]$, we define the valuation $v^t_{n,\mathsf{vcol}} : Q' \to [0, 1]$: $v^t_{n,\mathsf{vcol}}[Q_n \cup K^n] := \{t\}$ and for all $x \in V_{Q \setminus Q_u}$, $v^t_{n,\mathsf{vcol}}(x) := x$.

The game $\mathcal{L}^n_{\mathsf{vcol}}$ is then equal to $\mathcal{L}^n_{\mathsf{vcol}} = \langle \mathcal{C}^n_{\mathsf{vcol}}, \mathsf{Parity}_\mathsf{K} \rangle$.

Remark 8.5. *First, the notation $\mathcal{L}^n_{\mathsf{vcol}}$ comes from the fact that the game is extracted for the n-colored layer w.r.t. the coloring function vcol. The idea behind Definition 8.12 is the following: the states of interest are those of Q_n, that is, those for which the virtual color given by vcol is n. Note however that the colors of these states in $\mathcal{L}^n_{\mathsf{vcol}}$ are given by the real coloring function col. On the other hand, for all $i \in [\![0, e]\!]$, the state k^n_i in $\mathcal{L}^n_{\mathsf{vcol}}$ correspond to the states in $\mathcal{G}^u$ colored with i w.r.t. vcol. In the case where n is even, as formally defined later in Definition 8.16, we will require that any Player-A positional strategy generated by a given environment has value at least u, in the game $\mathcal{L}^n_{\mathsf{vcol}}$, from all states in Q_n. However, all states k^n_i for $i \in [\![0, e]\!]$ are Player-B's, who can then choose to loop back to any state in Q_n. Therefore, given a Player-A positional strategy s_A, if the game cannot exit to any stopping state, for the strategy s_A not to have value 0, the game may loop on some k^n_i only at the condition that the highest color seen with positive probability is even. In addition, note that the color of the state k^n_i for $i \in [\![0, n-1]\!]$ is $n-1$ (which is odd). Hence, all other things being equal, the game is harder for Player A when $n = 4$ than when $n = 2$ or 0. Finally, note that, when n is odd, we will require that any Player-B positional strategy generated by a given environment has value less than u.*

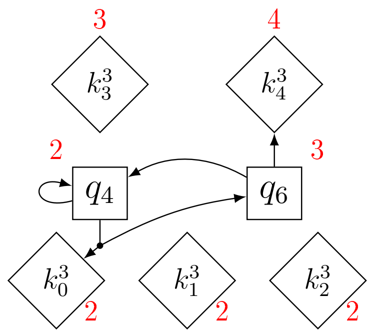

Figure 8.5: The game $\mathcal{L}^3_{\mathsf{vcol}}$. For readability, exiting arrows from $k_0^3, k_1^3, k_2^3, k_3^3$ and k_4^3 are not depicted: they would all loop back to both q_4 and q_6.

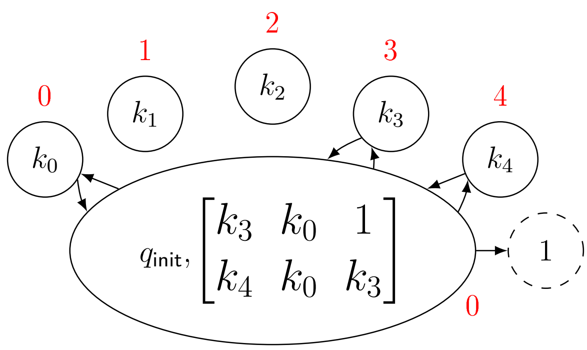

Figure 8.6: The game $\mathcal{G}^0_{q_0,\mathsf{vcol}}$ with vcol the coloring function depicted in Figure 8.4.

Example 8.3. *The game $\mathcal{L}^3_{\mathsf{vcol}}$ is partly depicted in Figure 8.5 (the virtual coloring function* vcol *being the one depicted in Figure 8.4). The colors of the states are depicted in red. Although the arrows are not depicted, from all states $k^3_0, k^3_1, k^3_2, k^3_3$ and k^3_4 Player* A *can decide to which state among $\{q_4, q_6\}$ to loop back (since $n = 3$ is odd). In an even-colored layer, it would have been Player* B *to decide.*

Given a virtual coloring function, we also associate a local environment with each state.

Definition 8.13 (Local environment induced by a virtual coloring function and a color)**.** *Consider a state $q \in Q_u$, a coloring function $\mathsf{vcol} : Q_u \to [\![0, e]\!]$. We define $p_{q,\mathsf{vcol}} : Q \to \{q_{\mathsf{init}}\} \cup K_e \cup [0, 1]$ similarly to how we define $p_{n,\mathsf{vcol}}$ in Definition 8.11. That is, for all $q' \in Q$, we have:*

- *if $q' = q$, $p_{q,\mathsf{vcol}} := q_{\mathsf{init}}$;*
- *if $q' \in Q_u \setminus \{q\}$, $p_{q,\mathsf{vcol}} := k_{\mathsf{vcol}(q')}$;*
- *if $q' \in Q \setminus Q_u$, $p_{q,\mathsf{vcol}}(q) := \chi_{\mathcal{G}}(q)$.*

Then, for all $n \in [\![0, e]\!]$, the environment $E^n_{q,\mathsf{vcol}}$ associated with state q w.r.t. vcol *and n is such that $E^n_{q,\mathsf{vcol}} := \langle \max(c_n, \mathsf{vcol}(q)), e, p_{q,\mathsf{vcol}} \rangle$ where $c_n = n + 1$ if n is odd and $c_n := n - 1$ if n is even. We say that the coloring function* vcol *is* associated *with the environment $E^n_{q,\mathsf{vcol}}$. The corresponding (local) game $\mathcal{G}_{(\mathsf{F}(q),E^n_{q,\mathsf{vcol}})}$ (see Definition 8.5) is denoted $\mathcal{G}^n_{q,\mathsf{vcol}}$. For all $x \in [0, 1]$, we set $v^x_{q,\mathsf{vcol}} := v^x_{(\mathsf{F}(q),E^0_{q,\mathsf{vcol}})}$ (see Definition 8.5).*

The definition of c_n may seem ad hoc. We give an explanation below in Page 303 of this definition.

Example 8.4. *The game $\mathcal{G}^n_{q_5,\mathsf{col}}$ is depicted on the right of Figure 8.1 for $n = 0, 1, 2$. However, if $n = 3$, the color of q_{init} would be 4, and if $n = 4$, it would be 3. The game $\mathcal{G}^n_{q_0,\mathsf{vcol}}$ is depicted in Figure 8.6 for $n = 0$. However, if $n = 1$, the color of q_{init} would be 2, if $n = 2$, the color would be 1, if $n = 3$, the color would be 4 and if $n = 4$ the color would be 3.*

8.4.3 . Local Operator

We want to define a way to update a (virtual) coloring function vcol. This will be done via a local operator mapping a given state q to the best color k for which Player A can achieve the value u in the corresponding local parity game $\mathcal{G}^k_{q,\mathsf{vcol}}$. Note that "best" is to be understood considering an ordering compatible with the parity objective. Specifically, taking the point-of-view of Player A, any even number is better than any odd number, and when they increase, odd numbers get worse whereas even numbers get better. This induces a new ordering.

Definition 8.14 (Parity order). *We define a total strict order relation* $\prec_{\mathsf{par}}$ *on* $\mathbb{N}$ *such that, for all* $m, n \in \mathbb{N}$*, we have* $m \prec_{\mathsf{par}} n$ *if:* m *is odd and* n *is even; or* $m > n$ *and* m *and* n *are odd; or* $m < n$ *and* m *and* n *are even.*

Definition 8.15 (Local operator). *Consider a state* $q \in Q_u$ *and a (possibly virtual) coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$*. The color* $\mathsf{NewCol}(q, \mathsf{vcol}) \in \mathbb{N}$ *induced by* vcol *at* q *is defined by:*

$$\mathsf{NewCol}(q, \mathsf{vcol}) := \max_{\prec_{\mathsf{par}}} \{n \in [\![0, e]\!] \mid \chi_{\mathcal{G}^n_{q,\mathsf{vcol}}}(q_{\mathsf{init}}) \geq u\}$$

The meaning of a new virtual color n assigned to a state q via NewCol is the following: in the game $\mathcal{G}^u$ with the coloring function vcol, from state q and in at most one step, the highest color w.r.t. vcol seen with positive probability when both players play optimally is n (and no stopping state is seen).

Let us now explain the choice of c_n in Definition 8.13. In a local environment parameterized by n, the integer n induces a shifted parity objective for Player A: her objective is that the maximal color seen infinitely often is at least n w.r.t. $\prec_{\mathsf{par}}$; in particular $n = 0$ induces the usual parity objective. The value c_n encodes that winning condition. For instance, if $n = 2$, assuming $\mathsf{vcol}(q) = 0$ for simplicity, then $c_n = 1$, which implies that seeing 0 infinitely often is not enough, but seeing 2 infinitely often is enough to win. Similarly, if $n = 1$, $c_n = 2$, which implies that seeing 1 infinitely often is now enough to win, but seeing 3 infinitely often is still losing.

Remark 8.6. *Assume for instance that* $\mathsf{NewCol}(q, \mathsf{vcol}) = 2$ *for some state* $q \in Q$ *and some (virtual) coloring function* vcol[1]*. In particular, Player* A *has a* GF*-strategy* σ_{A} *optimal w.r.t.* $(\mathcal{F}, E^2_{q,\mathsf{vcol}})$*, which yields a value at least* u*, and Player* B *has a* GF*-strategy* σ_{B} *optimal w.r.t.* $(\mathcal{F}, E^4_{q,\mathsf{vcol}})$*, which yields a value less than* u*. Remember that all games* $\mathcal{G}^k_{q,\mathsf{vcol}}$ *(where* k *ranges over* $[\![0, e]\!]$*) share the same structure and that only the color of* q_{init} *changes. Consider what happens in the game* $\mathcal{G}^0_{q,\mathsf{vcol}}$ *(where the state* q_{init} *is colored by* $\mathsf{vcol}(q)$*) when playing strategies* σ_{A} *and* σ_{B}*. Recalling Remark 8.3, the game cannot exit to any stopping state since the expected value of the stopping states reached would be both at least* u *(since* σ_{A} *is optimal in* $\mathcal{G}^2_{q,\mathsf{vcol}}$*) and less than* u *(since* σ_{B} *is optimal in* $\mathcal{G}^4_{q,\mathsf{vcol}}$*); hence this cannot happen. Furthermore, if some color of value at least 3 is seen with positive probability, then the highest such color must be both even (since* σ_{A} *is optimal in* $\mathcal{G}^2_{q,\mathsf{vcol}}$*) and odd (since* σ_{B} *is optimal in* $\mathcal{G}^4_{q,\mathsf{vcol}}$*). In fact, the highest color seen with positive probability in* $\mathcal{G}^0_{q,\mathsf{vcol}}$ *under* σ_{A} *and* σ_{B} *is 2 and neither of the players can do better (w.r.t. the ordering* $\prec_{\mathsf{par}}$*). From this, we can infer the semantics of a virtual color* n *assigned to a state* q *via operator* NewCol *(i.e. a color given by a virtual coloring function): in the game* $\mathcal{G}^u$ *with coloring function* vcol*, from state* q *and in at most one*

[1]In particular, it must be the case that $\mathsf{vcol}(q) \leq 2$, see Proposition 8.13.

step, the highest color w.r.t. vcol *seen with positive probability is* n *(and no stopping state can be seen).*

Let us exemplify this operator on an example.

Example 8.5. *First, consider Figure 8.1 and let us compute* $\mathsf{NewCol}(q_5, \mathsf{col})$. *We can realize that, regardless of the color of state* q_{init}*, Player* A *can (positionally) play both rows with positive probability and ensure reaching (almost-surely) the stopping state* $1/2$. *In fact, for all* $n \in [\![0, 4]\!]$*, we have* $\chi_{\mathcal{G}^n_{q_5,\mathsf{col}}}(q_{\mathsf{init}}) = 1/2$. *Hence,* $\mathsf{NewCol}(q_5, \mathsf{col}) = 4$.

Consider now Figure 8.6 and let us compute $\mathsf{NewCol}(q_0, \mathsf{vcol})$. *As mentioned in Example 8.4, the game* $\mathcal{G}^4_{q_0,\mathsf{vcol}}$ *corresponds to the game depicted in Figure 8.6 except that* q_{init} *is colored with 3. One can realize that, with this choice (of coloring of the state* q_{init}*), if the highest color* $i \in [\![0, 4]\!]$ *such that* k_i *is seen infinitely often is such that* $i \prec_{\mathsf{par}} 4$*, then Player* A *loses. The value of this game is 0 as Player* B *can ensure looping on* k_0 *and* q_{init} *(by playing, positionally and deterministically, the middle column) thus ensuring that the highest color seen infinitely often is 3. Thus,* $\mathsf{NewCol}(q_0, \mathsf{vcol}) \prec_{\mathsf{par}} 4$. *In the game* $\mathcal{G}^2_{q_0,\mathsf{vcol}}$*,* q_{init} *is colored with 1. Again, with this choice (of coloring of the state* q_{init}*), if the highest color* $i \in [\![0, 4]\!]$ *such that* k_i *is seen infinitely often is such that* $i \prec_{\mathsf{par}} 2$*, then Player* A *loses. The value of this game is also 0 as Player* B *can still play the middle column ensuring that the highest color seen infinitely often is 1. Thus,* $\mathsf{NewCol}(q_0, \mathsf{vcol}) \prec_{\mathsf{par}} 2$. *Consider now the game* $\mathcal{G}^0_{q_0,\mathsf{vcol}}$*, the one depicted in Figure 8.6. The value of the state* q_{init} *is now 1. Indeed, if Player* A *plays the two rows with equal probability, one can see that this strategy parity dominates (see Definition 8.2) the valuation* $v^1_{q_0,\mathsf{vcol}}$ *(recall Definition 8.13). Indeed, the BSCCs compatible with this strategy are* $\{q_{\mathsf{init}}, k_3, k_4\}$ *and* $\{q_{\mathsf{init}}, k_0\}$ *and they are even-colored. Hence, by Proposition 8.1,* $\chi_{\mathcal{G}^0_{q_1,\mathsf{col}}}(q_{\mathsf{init}}) = 1 \geq 1/2$ *and* $\mathsf{NewCol}(q_0, \mathsf{vcol}) \succeq_{\mathsf{par}} 0$. *That is,* $\mathsf{NewCol}(q_0, \mathsf{vcol}) = 0$.

Some of the properties enjoyed by the local operator NewCol are given in the appendix in Page 316.

8.4.4 . Faithful coloring function

To prove Theorem 8.3, we iteratively build a (virtual) coloring function and a local environment. We want to define the desirable property that the pair of coloring and environment functions should satisfy that will be preserved step by step. First, we need to define the notion of an environment function witnessing a color.

Definition 8.16 (Environment witnessing a color). *Consider a coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$*, a color* $n \in [\![0, e]\!]$ *and an environment function* $\mathsf{Ev} : Q_n \to \mathsf{Env}(Q)$ *with* $Q_n := \mathsf{vcol}^{-1}[n]$.

Assume that n *is even. We say that the pair* $(\mathsf{vcol}, \mathsf{Ev})$ witnesses *the color*

n if for all $q \in Q_n$, $\mathsf{Sz}_\mathsf{A}(\mathsf{Ev}(q)) \leq e - \mathsf{col}(q)$ and all positional Player-A strategies $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$ generated by Ev (recall Definition 8.9) in the game $\mathcal{L}^n_{\mathsf{vcol}}$ parity dominate the valuation $v^u_{n,\mathsf{vcol}}$ (recall Definition 8.12).

Assume that n is odd. We say that the pair $(\mathsf{vcol}, \mathsf{Ev})$ witnesses the color n, if for all $q \in Q_n$, $\mathsf{Sz}_\mathsf{B}(\mathsf{Ev}(q)) \leq o - \mathsf{col}(q)$ and for all positional Player-B strategies $\mathsf{s}_\mathsf{B} \in \mathsf{S}_\mathsf{B}^\mathcal{C}$ generated by Ev in the game $\mathcal{L}^n_{\mathsf{vcol}}$, there is some $u' < u$, such that s_B parity dominates the valuation $v^{u'}_{n,\mathsf{vcol}}$.

Remark 8.7. *The condition on the size of the environments considered along with the assumptions of Theorem 8.3 ensures that the quantification over the strategies generated by the environment function are not over the empty set.*

This definition was hinted in Remark 8.5. Informally, it means that, in the (virtual) games given by vcol, in the even-colored layers, Player A can achieve at least what she should be able to achieve in this u-slice (i.e. the value of the states is at least u). Whereas, in the odd-colored layers, Player B can prevent Player A from achieving this.

We can now define the notion of faithful pair of coloring and environment functions.

Definition 8.17 (Faithful pair of coloring and environment functions). *Consider a coloring function $\mathsf{vcol} : Q_u \to [\![0, e]\!]$, some $n \in [\![0, e+1]\!]$ and a partial environment function $\mathsf{Ev} : Q_u \to \mathsf{Env}(Q)$ defined on $\mathsf{vcol}^{-1}[[\![n, e]\!]]$. We say that $(\mathsf{vcol}, \mathsf{Ev})$ is* faithful down to *n if:*

- *for all $k \in [\![n, e]\!]$, the pair $(\mathsf{vcol}, \mathsf{Ev})$ witnesses color k;*
- *for all $q \in Q_u$, if $\mathsf{vcol}(q) < n$, then $\mathsf{col}(q) = \mathsf{vcol}(q)$ and $\mathsf{NewCol}(q, \mathsf{vcol}) < n$;*

If $n = 0$, we say that the pair $(\mathsf{vcol}, \mathsf{Ev})$ is completely faithful.

Only the first condition for faithfulness is really of interest to us. For instance, this first condition suffices to show the crucial proposition below. However, the second condition is used in the proofs. (It is also helpful as it guides us in how to build a completely faithful pair, as discussed below.) Note that, in the proofs, we use an even stronger notion of faithfulness with a third condition. However, we do not present it here in order not to complexify too much the approach, and the two first conditions are sufficient for the (informal) explanation of Theorem 8.3. It can however be found in Appendix 8.6.5.

The benefit of faithful environments and coloring functions lies in the proposition below: if all states are mapped w.r.t. the coloring function to e, then the environment function guarantees the value u in the whole u-slice Q_u.

Proposition 8.5. *For a coloring function $\mathsf{vcol} : Q_u \to [\![0, e]\!]$ and an environment function $\mathsf{Ev} : Q_u \to \mathsf{Env}(Q)$, assume that $(\mathsf{vcol}, \mathsf{Ev})$ is completely faithful* *and that* $\mathsf{col}_u[Q_u] = \{e\}$. *Then, all Player-*$\mathsf{A}$ *positional strategies generated by the environment function* Ev *parity dominate the valuation* $\chi_{\mathcal{G}}$ *in the game* $\mathcal{G}^u$.

Proof. This is direct from the definitions. Indeed, as $(\mathsf{vcol}, \mathsf{Ev})$ is completely faithful, it follows that $(\mathsf{vcol}, \mathsf{Ev})$ witnesses the color e (see Definition 8.17). That is, all Player-A positional strategies $\mathsf{s_A}$ generated by Ev in the game $\mathcal{L}^e_{\mathsf{vcol}}$ parity dominate the valuation $v^u_{e,\mathsf{vcol}}$ (see Definition 8.16). Since $\mathsf{vcol}[Q_u] = \{e\}$, both games $\mathcal{L}^e_{\mathsf{vcol}}$ and $\mathcal{G}^u$ are identical (see Definitions 8.10 and 8.12). Similarly, the valuation $v^u_{e,\mathsf{vcol}}$ is equal to the valuation $\chi_{\mathcal{G}}$ in the game $\mathcal{G}^u$ (also see Definition 8.16). □

8.4.5 . Computing a completely faithful pair

Given Lemma 8.5, our goal is to come up with a pair of an environment function and a coloring function completely faithful such that all states are colored with e. Let us first consider how to obtain a completely faithful pair from the initial coloring function and the empty environment function (i.e. no state is mapped to an environment). Note that this initial pair of coloring and environment functions is faithful down to $e+1$. Hence, our goal is, given a pair $(\mathsf{vcol}, \mathsf{Ev})$ faithful down to some $n \in [\![1, e+1]\!]$, to build a new pair that is faithful down to $n-1$. To do so, let us be guided by the second property for faithfulness: to be faithful down to $n-1$, no state $q \in Q_u$ such that $\mathsf{vcol}(q) \leq n-2$ should be such that $\mathsf{NewCol}(q, \mathsf{vcol}) = n-1$. Hence, the idea is, for all such states $q \in Q_u$, to change their colors to $n-1$ until no state $q \in Q_u$ with $\mathsf{vcol}(q) \leq n-2$ satisfies[2] $\mathsf{NewCol}(q, \mathsf{vcol}) = n-1$. The environment associated to each such state q newly colored by $n-1$ will be given by the coloring function vcol for which $\mathsf{NewCol}(q, \mathsf{vcol}) = n-1$ for the first time (crucially, this is done before the color of q is updated to $n-1$). The procedure we have described is formally given in the Appendix as Algorithm 8.12. Interestingly, the update done in the algorithm preserves the faithfulness of environment and coloring functions.

Lemma 8.6 (Proof Page 328). *Consider a coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$, $n \in [\![1, e+1]\!]$, *and a partial environment function* $\mathsf{Ev} : Q_u \to \mathsf{Env}(Q)$ *defined on* $\mathsf{vcol}^{-1}[[\![n, e]\!]]$. *Assume that* $(\mathsf{vcol}, \mathsf{Ev})$ *is faithful down to* n. *Let* $(\mathsf{vcol}', \mathsf{Ev}') \leftarrow \mathsf{UpdateColEnv}(n-1, \mathsf{vcol}, \mathsf{Ev})$ *be the pair computed by Algorithm 8.12 for index* $n-1$. *(Only states* q *such that* $\mathsf{vcol}'(q) = n-1$ *may have changed their colors and be newly mapped to an environment.) Then,* $(\mathsf{vcol}', \mathsf{Ev}')$ *is faithful down to* $n-1$.

Before giving a proof sketch of this lemma, let us illustrate it on an example.

[2]This can seen as computing a "probabilistic attractor with leaks towards the stopping states" that we mentioned above.

Example 8.6. *Let us illustrate this algorithm on Figures 8.3 and 8.4. The first step is to build a pair that is faithful down to $e = 4$. As mentioned above in Example 8.5, we have* $\mathsf{NewCol}(q_5, \mathsf{col}) = 4$*. Hence, the color of this state is changed to 4 (we obtain a virtual coloring function* vcol^{q_5}*) and we set* $\mathsf{Ev}(q_5) := E^4_{q_5,\mathsf{col}}$*. Note that a Player-*$\mathsf{A}$ GF*-strategy* σ_A *is optimal in this environment if and only if it plays both rows with positive probability. Furthermore, note that, in the extracted game* $\mathcal{L}^4_{\mathsf{vcol}^{q_5}}$*, a Player-*$\mathsf{A}$ *positional strategy playing such a* GF*-strategy* σ_A *in* q_5 *parity dominates the valuation* $v^{1/2}_{Q_4,\mathsf{col}^{q_5}}$*. Hence, the pair* $(\mathsf{vcol}^{q_5}, \mathsf{Ev})$ *is faithful down to 4.*

Consider now the layer 3. First, the state q_6 *already has color 3, so it only remains to set its environment:* $\mathsf{Ev}(q_6) := E^3_{q_6,\mathsf{vcol}^{q_5}}$*. We then realize that* $\mathsf{NewCol}(q_4, \mathsf{vcol}^{q_5}) = 3$*. Indeed,* q_4 *is colored with 2 and may go with equal probability to a state colored with 0 and to a state colored with 3. The color of this state is therefore changed, thus obtaining a new virtual coloring function* $\mathsf{vcol}^{q_5,q_6,q_4}$*. We set its environment:* $\mathsf{Ev}(q_4) := E^3_{q_4,\mathsf{vcol}^{q_5}}$*. One can realize that the pair* $(\mathsf{vcol}^{q_5,q_6,q_4}, \mathsf{Ev})$ *witnesses the color 3 : a positional Player-*B *strategy generated by this environment would be so that (i) from* q_6*, it goes to* q_4 *with probability 1 (to avoid* k_4 *that is colored with 4) and (ii) from* q_5*, it goes to* q_6 *with positive probability (to see the color 3 with positive probability). Such a strategy has value 0 in the game* $\mathcal{L}^3_{\mathsf{vcol}^{q_5}} = \mathcal{L}^3_{\mathsf{vcol}}$ *from Figure 8.5, hence the pair* $(\mathsf{vcol}^{q_5,q_6,q_4}, \mathsf{Ev})$ *witnesses the color 3.*

We illustrate on this step why the environment needs to be set before changing the new color and not after. That is, we explain why it would not be correct to set $\mathsf{Ev}(q_4) := E^3_{q_4,\mathsf{col}^{q_5,q_6,q_4}}$ *instead of what we do above. In this environment, the state* q_4 *has color 3. Hence, looping with probability 1 on* q_4 *is an optimal* GF*-strategy for Player* B *w.r.t.* $(\mathsf{F}(q_4), \mathsf{Ev}(q_4))$*. Then, the corresponding pair of coloring and environment functions would not witness the color 3. Indeed, a Player* B *strategy that loops with probability 1 on* q_4 *is generated by this environment, and it has value* $1 \geq u$ *(because the real color of this state is 2, and not 3).*

This process is then repeated down to 0. In Figure 8.4, the depicted coloring function (with the appropriate environment function that is not shown in Figure 8.4) are in fact completely faithful (this is what outputs Algorithm 8.12 on the coloring function of Figure 8.3).

We give a proof sketch of Lemma 8.6, which explains the ideas for the first phase of the procedure for computing a first completely faithful pair, before Algorithm 8.13 is called – which we will discuss right after.

Proof sketch. We want to prove that the pair $(\mathsf{vcol}', \mathsf{Ev}')$ witnesses the color $n - 1$ (the other condition for faithfulness is ensured by the construction). We consider the case where $n - 1$ is even, the other case is similar (but one needs to take the point-of-view of Player B). Consider a Player-A positional strategy $\mathsf{s_A}$ generated by the environment function Ev' in the game $\mathcal{L}^{n-1}_{\mathsf{vcol}'}$. Let $Q_{n-1} := \mathsf{vcol}'^{-1}[n-1]$ and let $v := v^u_{n-1,\mathsf{vcol}'}$. For every $q \in Q_{n-1}$, let $Y_q := (\mathsf{F}(q), \mathsf{Ev}'(q))$ be the local environment at state q and let $\mathsf{Ev}'(q) = \langle c_q, e, p_q \rangle$. From the characterization of Lemma 8.2 (item (ii.1)), by carefully analyzing the links between the local games $\mathcal{G}_{Y_q}$ for all $q \in Q_{n-1}$ and the game $\mathcal{L}^{n-1}_{\mathsf{vcol}'}$, we deduce that the strategy $\mathsf{s_A}$ dominates the valuation v.

It remains to show that all BSCCs (that are not reduced to a stopping state and are) compatible with s_A are even-colored. Consider such a BSCC H and a Player-B deterministic positional strategy $\mathsf{s_B}$ which induces H. For every state $q \in H$, since no stopping state occurs in H, it must be that the probability to reach a stopping state is 0. That is, it amounts to have $\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q \setminus Q_u} \rangle](\sigma_\mathsf{A}, b) = 0$. For every state $q \in Q_{n-1}$, the coloring function vcol_q associated with environment $\mathsf{Ev}'(q)$ is such that $\mathsf{vcol}_q(q) \leq n-1$.[3] Hence, the color c_q is such that $c_q = \max(n-2, \mathsf{vcol}_q(q)) \leq n-1$. Now, assume that some state k_i is in H for some $i > n-1 \geq c_q$. In that case, as explained in Remark 8.3, the highest i such that k_i is in H must be even. Hence, H is even-colored. Assume now that no state k_i in H is such that $i > n-1$. In that case, if a state in H has color $n-1$ (like the state q_6 in Figure 8.4 in the case where $n-1 = 3$), then $n-1$ is the highest color in H and H is even-colored. Consider the first state q whose color is now $n-1$ (w.r.t. vcol') but whose previous color was not $n-1$. In that case, we have $c_q = \max(n-2, \mathsf{vcol}_q(q)) = n-2$ is odd. Furthermore, the state q has changed its color because $\mathsf{NewCol}(q, \mathsf{vcol}_q) = n-1$. With Remark 8.3, since $\mathsf{s_A}(q)$ is optimal w.r.t. Y_q, it follows that there is a positive probability to reach, in the game $\mathcal{G}_{Y_q}$ the state k_{n-1}. In the game $\mathcal{L}^{n-1}_{\mathsf{vcol}'}$, this corresponds to a positive probability to reach a state $q' \in H$ colored with $n-1$ w.r.t. vcol_q (recall Definition 8.11). Since q is the first state to have changed its color, we can deduce that q' already had color $n-1$ w.r.t. vcol. Furthermore, one can show that q' is colored with $n-1$ w.r.t. the real coloring function col. Overall, in the game $\mathcal{L}^{n-1}_{\mathsf{vcol}'}$, with the GF-strategy $\mathsf{s_A}(q)$, there is a positive probability to reach in one step a state q' colored with $n-1$. Iteratively, we obtain that, considering the k-th state whose color is now $n-1$ (i.e. w.r.t. vcol') but whose initial color was not $n-1$, there is a positive probability to reach (in at most k steps) a state colored with $n-1$. Hence, the highest color appearing in H is $n-1$, which is even. We obtain that $\mathsf{s_A}$ parity dominates the valuation v. $\square$

Overall, applying iteratively Algorithm 8.12 on all colors from e down to 0 starting with the initial coloring function induces a completely faithful pair $(\mathsf{vcol}, \mathsf{Ev})$. However, it may be the case that some states are mapped to an odd number, which does not allow to apply Lemma 8.5. The question is then:

[3]This is because all states $q \in Q_{n-1}$ satisfy $\mathsf{col}(q) \leq n-1$. This is one of the additional conditions for faithfulness that we did mention, but that is used in the Appendix in Definition 8.22.

from that completely faithful configuration, how can one make some progress towards a situation where Lemma 8.5 can be applied?

Example 8.7. *Consider the coloring function of Figure 8.4. As mentioned in Example 8.6, with an appropriate environment function (that is not shown in Figure 8.4), we can have a pair which is completely faithful. To gain some intuition on what should be done next, let us focus only on the states q_1, q_2, q_3. A simplified version is presented in Figure 8.7 (with a slight modification: instead of going to q_0, q_1 loops on itself): the initial (and true) colors of the states are in circles next to them and their color w.r.t. the current (virtual) coloring function (that is completely faithful with an appropriate environment function) is written in red. In this game, Player* B *plays alone, but it is obvious that Player* A *wins surely from q_2: indeed, either the game stays indefinitely in q_2, or it eventually reaches and settles in q_1.*

The current virtual color 1 assigned to both q_2 and q_3 does not properly reflect the fact that if the game reaches q_3, even though Player B *plays optimally according to the local game associated with q_2, it will end up looping in q_1, which will be losing for Player* B*. In a way, we would like to propagate the information that reaching q_1 is bad for Player* B*. Since 0 is the smallest color, there is no harm in increasing it to 2, the game from q_1 will be the same: it will be won by Player* A *by looping. Player* B *will now be able to know that going to q_1 is dangerous for him, which will be obtained by applying the previous iterative process.*

In a more general concurrent game, the next step of the process when we have a completely faithful configuration not satisfying the assumptions of Lemma 8.5 consists in changing all the states with the least (virtual) color n to the color $n+2$. However, note that there is a (very important) second step: the colors of all states (virtually) colored with $n+1$ should be reset to their initial colors. The reason why can be seen again in Figure 8.7. After the color of q_1 becomes 2, the color of q_3 will also become 2. However, if the color of the state q_2 is not reset, then it is not going to change since Player B *can choose to loop to q_2 and see the color 1 for ever (in game $\mathcal{G}^0_{q_2,\mathsf{vcol}}$). That is, from Player* B*'s perspective, looping indefinitely on q_2 is winning, which is not what happens in the real game (i.e. the coloring function does not faithfully describes what happens in the game). The changes made to the coloring function* vcol *from Figure 8.4 can be seen in Figure 8.8. Note that the process of increasing the colors of some states by 2 can only be done with the* least *color (otherwise faithfulness will not be preserved).*

The process described in Example 8.7 is implemented as Algorithm 8.13 in the Appendix, it ensures the lemma below.

Lemma 8.7 (Proof Page 332). *Let* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$, $\mathsf{Ev} : Q_u \to \mathsf{Env}(Q)$ *be a coloring and an environment functions. Let* $n := \min \mathsf{vcol}[Q]$. *Assume*

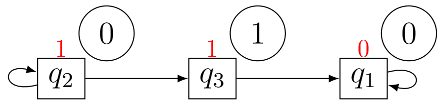

Figure 8.7: A (deterministic turn-based) game with only three states.

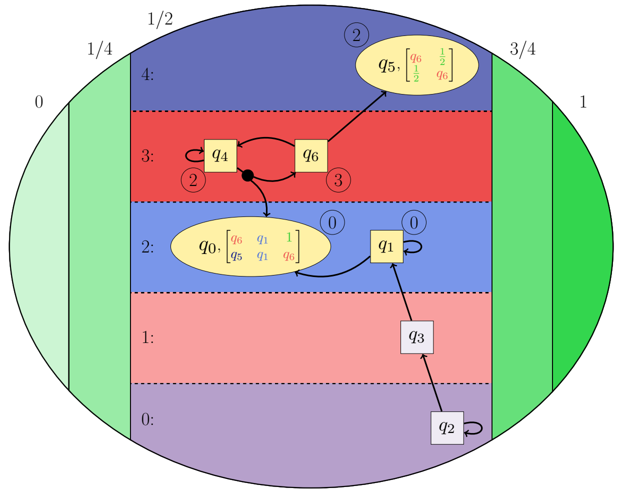

Figure 8.8: The same arena as in Figures 8.3,8.4 but with a different coloring function.

that $n \leq e - 2$ and the pair $(\mathsf{vcol}, \mathsf{Ev})$ *is completely faithful. If* $(\mathsf{vcol}', \mathsf{Ev}') \leftarrow \mathsf{IncLeast}(\mathsf{vcol}, \mathsf{Ev})$ *is the result of increasing the least-colored layer by 2 and resetting the environment of the last but least-colored layer (Algorithm 8.13), then* $(\mathsf{vcol}', \mathsf{Ev}')$ *is faithful down to* $n + 2$.

Proof sketch. Let $Q_n := \mathsf{vcol}^{-1}[n]$ and $Q_{n+2} := \mathsf{vcol}^{-1}[n+2]$. The algorithm has three steps: first, it increases the least color by 2; then it resets the environments of the $(n+1)$-colored states; finally it applies Algorithm 8.11 to these reset states. Let us argue that $(\mathsf{vcol}'', \mathsf{Ev})$ (obtained after the first step) witnesses the color $n + 2$.

Consider a Player-A positional strategy $\mathsf{s_A}$ generated by the environment Ev in the game $\mathcal{L}^{n+2}_{\mathsf{vcol}''}$. Let $v := v^u_{n+1,\mathsf{vcol}''}$. Similarly to the proof of Lemma 8.6, $\mathsf{s_A}$ dominates the valuation v. Consider a BSCC H compatible with $\mathsf{s_A}$. If $H \cap Q_{n+2} = \emptyset$, then H is even-colored. Indeed, $(\mathsf{vcol}, \mathsf{Ev})$ witnesses the color n. In addition, the probability to go to a state k_i^{n+2} that is $(n+1)$-colored in the game $\mathcal{L}^{n+2}_{\mathsf{vcol}''}$ is exactly the probability to go to a state k_i^n that is $(n+1)$-colored in the game $\mathcal{L}^n_{\mathsf{vcol}}$ (since n is the least color). Furthermore, H is also even-colored as soon as $H \cap Q_n = \emptyset$ since $(\mathsf{vcol}, \mathsf{Ev})$ witnesses the color $n+2$. Now, assume that none of these cases occur. Then, one can show that: either a state k_i is seen for some $i \geq n+2$, and H is even-colored; or, from some states in Q_{n+2}, there is a positive probability to exit Q_{n+2} and no state k_i is seen for $i \geq n+2$. Now, looking at what happens in game $\mathcal{L}^{n+2}_{\mathsf{vcol}}$, some states k_i are seen for $i \leq n+1$, and such states are colored with $n+1$. Hence, since $(\mathsf{vcol}, \mathsf{Ev})$ witnesses the color $n+2$, it must be that the highest color in H is $n+2$, which is even. Therefore it is also the case in the game $\mathcal{L}^{n+2}_{\mathsf{vcol}''}$. In all the cases, H is even-colored. □

As stated in Lemma 8.7, the update of colors described in Example 8.7 can be done only if, for a completely faithful pair, the least (virtual) color n appearing is at most $e-2$. If $n = e$, we are actually in the scope of Lemma 8.5 since in that case all states have (virtual) color e. However, there remains the case where we have $n = e-1$. In fact, this case cannot happen.

Lemma 8.8 (Proof Page 333). *Consider a coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$*, an environment function* $\mathsf{Ev} : Q_u \to \mathsf{Env}(Q)$*. Assume that* $(\mathsf{vcol}, \mathsf{Ev})$ *is completely faithful. Then, for* $C := \mathsf{vcol}[Q]$*, we have* $\min C \neq e-1$*.*

Proof sketch. Let $Q_{e-1} := \mathsf{vcol}^{-1}[e-1]$. Towards a contradiction, let $\mathsf{s_B}$ be a Player-B positional strategy generated by Ev in the game $\mathcal{L}^{e-1}_{\mathsf{vcol}}$. It parity dominates the valuation $v^{u'}_{e-1,\mathsf{vcol}}$ for some $u' < u$. Hence, all BSCCs compatible with $\mathsf{s_B}$ are odd-colored: they all stay in the layer Q_{e-1}. Indeed, since $e - 1 = \min C$, exiting Q_{e-1} while staying in Q_u means seeing $Q_e := \mathsf{vcol}^{-1}[e]$ with e even and the highest color in the game. Hence, either the game stays indefinitely in Q_{e-1} and Player B wins almost surely, or there is some positive probability to visit stopping states, and in that case their expected values is at most u'. Hence, in the game $\mathcal{G}^u$, the strategy $\mathsf{s_B}$ has values less than u from the states $Q_{e-1} \subseteq Q_u$, which is a contradiction. □

Finally, all these pieces are put together in Algorithm 8.14 in the Appendix, whose output is a completely faithful pair where all states are mapped to e. The only remaining step is to prove the termination of this algorithm. Let us consider the (virtual) coloring functions as vectors in $\mathbb{N}^{e+1}$ indicating the number of states mapped to each color. Then, one can realize that each step of Algorithm 8.14 increases this vector for a lexicographic order (i.e. we first compare the number of states mapped to e, then the number of states mapped to $e-1$, etc). Hence, Algorithm 8.14 does terminate in finitely many steps.

Lemma 8.9 (Proof Page 335). *Algorithm 8.14 computes a completely faithful pair of environment and coloring functions mapping each state to* e *in finitely many steps.*

We can now proceed to the (informal) proof of Theorem 8.3.

Proof sketch. Let us prove Theorem 8.3 for Player A. Consider some $u \in V_Q \setminus \{0\}$. By Lemma 8.9, there is a completely faithful pair of environment and coloring functions $(\mathsf{vcol}_u, \mathsf{Ev}^u_{\mathsf{A}})$ mapping each state in Q_u to e_u. Hence, by Lemma 8.5, all Player-A positional strategies generated by the environment function $\mathsf{Ev}_{\mathsf{A}}^u$ parity dominate the valuation $\chi_{\mathcal{G}}$ in the game $\mathcal{G}^u$. Since we assume that all game forms occurring in Q_u are positionally maximizable up to $e_u - \mathsf{col}(q)$ w.r.t. Player A, such positional strategies generated by $\mathsf{Ev}_{\mathsf{A}}^u$ do exist. Then, considering the environment function $\mathsf{Ev}_{\mathsf{A}} : Q \to \mathsf{Ev}(\mathsf{D})$ that merges all the environment functions $(\mathsf{Ev}_{\mathsf{A}}^u)_{u \in V_{\mathcal{G}} \setminus \{0\}}$ together (and that is defined arbitrarily on Q_0), it follows by Lemma 8.4, that all Player-A positional strategies generated by that environment function Ev are optimal. (And such strategies exist.) □

We finally state below an NSC-transfer.

Corollary 8.10. *Among standard finite game forms with finitely many outcomes, being positionally optimizable is an NSC-transfer for the existence of positional optimal strategies for both players in finite parity games.*

Proof. From a game form that is not positionally optimizable, one can build a simple parity game where one of the players has no positional optimal. This is by definition of positionally optimizable game forms (Definition 8.7) and by Definition 8.6. The other direction comes from Theorem 8.3. □

Note that, for simplicity, we have not stated an NSC-transfer as we did in Proposition 7.7 where the conditions on each state depends on the color of the state. It is plausible that such an NSC-transfer could be stated for the finite-state parity games that we have considered in this chapter. However, it would probably involve introducing the maximal color appearing in the parity games considered.

8.5 Discussion and future work

In this chapter, we have proved an NSC-transfer, among standard finite game forms, for the existence of positional optimal strategies in finite parity games for both players. A natural future work would be to prove the result for any one player, independently of the other player. As mentioned in the introduction of this chapter, this would require, a priori, to manipulate infinite-choice strategies for the player we are not handling, which is why we have not done it yet. One of the main difficulties would then be to glue infinite-choice local strategies (i.e. strategies in simple parity games induced by the different game forms $\mathsf{F}(q)$) together into a single global strategy, which is straightforward when all local strategies are positional. Furthermore, we could no longer use Proposition 8.2, which comes from Corollary 3.16 since it only applies to positional strategies. Instead, we would have to use the more general Corollary 3.14.

In a somewhat similar fashion than for Chapter 2 with Theorem 2.3, we believe that the benefit of this chapter does not only lie in the results stated in Theorem 8.3 and Corollary 8.10, but also in the method we took to prove Theorem 8.3. Due the the concurrent stochastic setting of this dissertation, the arguments are quite technical. It may be interesting to restrict the setting to finite turn-based stochastic games, which would (greatly) simplify the proof, or at least remove its most technical aspects. We believe that it could provide a new proof of the existence of positional optimal strategies in finite turn-based parity games for both players, which was originally proved in [27, 28], as it is straightforward to show that finite turn-based game forms are positionally optimizable.

8.6 Appendix

8.6.1 . Algorithms

The algorithms we mentioned in this chapter are gathered in this subsection. Let us quickly describe what each of them does:

- Consider a color $k \in [\![0, e]\!]$, a state $q \in Q_u$ and a virtual coloring function $\mathsf{vcol} : Q_u \to [\![0, e]\!]$. Then, the algorithm $\mathsf{CreateEnv}(k, q, \mathsf{vcol})$ builds an environment, at state q, for either of the players: A if k is even, B if it is odd.

- Consider a color $k \in [\![0, e]\!]$, a virtual coloring function $\mathsf{vcol} : Q_u \to [\![0, e]\!]$ and an environment function defined on $\mathsf{vcol}^{-1}[[\![k + 1, e]\!]]$. Then, the algorithm $\mathsf{UpdCurSta}(k, \mathsf{vcol}, \mathsf{Ev})$ sets the environment of all states virtually colored by k by calling algorithm $\mathsf{CreateEnv}(k, q, \mathsf{vcol})$.

- Consider a color $k \in [\![0, e]\!]$, a virtual coloring function $\mathsf{vcol} : Q_u \to [\![0, e]\!]$ and an environment function defined on $\mathsf{vcol}^{-1}[[\![k + 1, e]\!]]$. Then, the algorithm $\mathsf{UpdNewSta}(k, \mathsf{vcol}, \mathsf{Ev})$ changes the colors and sets the environment of all the states $q \in Q_u$ virtually colored by some $i < k$ for which the operator NewCol is equal to k.

- Consider a color $k \in [\![0, e]\!]$, a virtual coloring function $\mathsf{vcol} : Q_u \to [\![0, e]\!]$ and an environment function defined on $\mathsf{vcol}^{-1}[[\![k+1, e]\!]]$. Then, the algorithm $\mathsf{UpdateColEnv}(k, \mathsf{vcol}, \mathsf{Ev})$ calls successively algorithms $\mathsf{UpdCurSta}(k, \mathsf{vcol}, \mathsf{Ev})$ and $\mathsf{UpdNewSta}(k, \mathsf{vcol}, \mathsf{Ev})$.

- Consider a virtual coloring function $\mathsf{vcol} : Q_u \to [\![0, e]\!]$ and an environment function defined on Q_u. Then, the algorithm $\mathsf{IncLeast}(\mathsf{vcol}, \mathsf{Ev})$ increases the colors of states virtually colored by the least color c_{min} by 2. Then, it resets the color and environment of all states colored by $c_{\mathsf{min}} + 1$.

- Finally, consider a coloring function $\mathsf{col} : Q_u \to [\![0, e]\!]$. Then, the algorithm $\mathsf{ComputeEnv}(\mathsf{col})$ iteratively computes a pair of virtual coloring function and a environment function by successively calling algorithms UpdateColEnv and IncLeast until all states are colored by e.

8.6.2 . Proof of Proposition 8.1

Proof. The first observation we can make is that, if $\mathsf{s_A}$ dominates the valuation $v : Q \to [0,1]$, then it also dominates the valuation $v_{\mathsf{lt}} : \Omega_{\mathcal{C}}^{+} \to [0,1]$ (recall Definition 3.9) in the sense of Definition 3.7.

We can therefore apply Corollary 3.15: for all ECs H in the finite MDP $\Gamma_{\mathcal{C}}^{\mathsf{s_A}}$ induced by the strategy $\mathsf{s_A}$, there is a value $u(v,H) \in [0,1]$ such that $v[Q_H] = \{u(v,H)\}$. This proves the first part of the proposition.

As for the second part of the proposition, this is a direct consequence of Corollary 3.16. Indeed, consider any EC H with $u(v,H) > 0$ and state $q \in Q_H$. Consider any Player-B positional deterministic strategy $\mathsf{s_B}$. Then, all the BSCCs that can occur in Q_H, in the Markov chain induced by $\mathsf{s_B}$, are even-colored. That is, from q, the parity objective holds almost-surely. Since this holds against all Player-B positional deterministic strategies and since positional deterministic strategies are enough to play optimally in finite MDPs with parity objectives [27], this proves that $\chi_{\mathcal{C}_H^{\mathsf{s_A}}}(q) = 1$.

□

8.6.3 . Proof of Lemma 8.2

Proof. As in the proof of Proposition 8.1, the first equivalence comes Corollary 3.16 (the other direction) and the fact that positional deterministic strategies are enough to play optimally in finite MDPs with parity objectives [27]. Now assume that $u > 0$. Assume that the strategy $\mathsf{s}_Y^{\mathsf{A}}(\sigma_{\mathsf{A}})$ parity dominates the valuation v_Y^u in the game $\mathcal{G}_Y$. In particular, it dominates this valuation, i.e. item $(ii.1)$ of Lemma 8.2 is satisfied. Note that this effectively amounts to have σ_{A} optimal in the game in normal form $\langle \mathcal{F}, v_Y^u \circ p \rangle$, since $u = v_Y^u(q_{\mathsf{init}}) = \mathsf{val}[\langle \mathcal{F}^p, v_Y^u \rangle]$ by Lemma 3.9. Consider now some action $b \in \mathsf{Act_B}$ such that $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[0,1]} \rangle](\sigma_{\mathsf{A}}, b) = 0$ and a Player-B positional and deterministic strategy $\mathsf{s_B}$ such that $\mathsf{s_B}(q_0) := b$. Since $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[0,1]} \rangle](\sigma_{\mathsf{A}}, b) = 0$, no stopping state can be reached under $\mathsf{s}_Y^{\mathsf{A}}(\sigma_{\mathsf{A}})$ and $\mathsf{s_B}$. Consider the Markov chain $\mathcal{T}_{\mathcal{C}_Y}^{\mathsf{s}_Y^{\mathsf{A}}(\sigma_{\mathsf{A}}),\mathsf{s_B}}$. Besides stopping states, it is reduced to a BSCC H whose states are q_{init} and all states k_i reachable with that action b. That is, $H = \{q_{\mathsf{init}}\} \cup \{k_i \mid i \in [\![0,e]\!],\ \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[k_i]} \rangle](\sigma_{\mathsf{A}}, b) > 0\}$. Furthermore, q_{init} is colored with c. The value of any state in this BSCC H is equal to either 0 or 1 and it is equal to 1 if and only if H is even-colored (since every state in a BSCC is almost-surely seen infinitely often). Since $u > 0$, the value of H cannot be 0, thus H is even-colored and the maximum of the colors seen with that action is even. This exactly corresponds to item $(ii.2)$ of Lemma 8.2.

Assume now that the GF-strategy σ_{A} satisfies both items $(ii.1)$ and $(ii.2)$. Consider a positional deterministic Player-B strategy $\mathsf{s_B}$ in the game $\mathcal{G}_Y$. Let

```
if k is even then
    return E^k_{q,vcol};
end
if k is odd then
    m ← max(k − 2, 0);
    return E^m_{q,vcol};
end
```

Figure 8.9: CreateEnv(k, q, vcol)

```
for q ∈ Q_u do
    if vcol(q) = k then
        Ev(q) ← CreateEnv(k, q, vcol) ;
    end
end
return Ev;
```

Figure 8.10: UpdCurSta$(k, \text{vcol}, \text{Ev})$

```
change ← True;
while change do
    change ← False;
    for q ∈ Q_u do
        if NewCol(q, vcol) = k then
            Ev′(q) ← CreateEnv(k, q, vcol);
            vcol(q) ← k;
            change ← True; break;
        end
    end
end
return (vcol, Ev′);
```

Figure 8.11: UpdNewSta$(k, \text{vcol}, \text{Ev})$

```
Ev ← UpdCurSta(k, vcol, Ev);
UpdNewSta(k, vcol, Ev);
```

Figure 8.12: UpdateColEnv$(k, \text{vcol}, \text{Ev})$

```
c_min ← min vcol;
for q ∈ Q_u do
    if vcol(q) = c_min then
        vcol(q) ← c_min + 2;
    end
    if vcol(q) = c_min + 1 then
        vcol(q) ← col(q);
        Ev′(q) ← NoEnv;
    end
end
UpdNewSta(c_min + 2, vcol, Ev′);
```

Figure 8.13: IncLeast(vcol, Ev)

```
vcol ← col;
Ev ← EmptyEnv;
for k = e down to 0 do
    (vcol, Ev) ← UpdateColEnv(k, vcol, Ev);
end
while vcol[Q_u] ≠ {e} do
    (vcol, Ev) ← IncLeast(vcol, Ev);
    for k = e down to 0 do
        (vcol, Ev) ← UpdateColEnv(k, vcol, Ev);
    end
end
return (vcol, Ev)
```

Figure 8.14: ComputeEnv(col)

$b := \mathsf{s}_\mathsf{B}(q_0) \in \mathsf{Act}_B$. Consider a BSCC H in the induced Markov chain $\mathcal{T}_{\mathcal{C}_Y}^{\mathsf{s}_Y^\mathsf{A}(\sigma_\mathsf{A}),\mathsf{s}_\mathsf{B}}$ that is not reduced to a stopping state. In particular, this implies that $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[0,1]} \rangle](\sigma_\mathsf{A}, b) = 0$. Then, as previously, we have $H = \{q_{\mathsf{init}}\} \cup \{k_i \mid i \in [\![0,e]\!],\ \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[k_i]} \rangle](\sigma_\mathsf{A}, b) > 0\}$. The fact that σ_A satisfies item $(ii.2)$ ensures that the BSCC H is even-colored. It follows that the strategy $\mathsf{s}_Y^\mathsf{A}(\sigma_\mathsf{A})$ parity dominates the valuation v_Y^u.

□

8.6.4 . Proof of Lemma 8.4

Proof. Consider such a Player-A strategy $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^\mathcal{C}$. Let us show that it parity dominates the valuation $\chi_\mathcal{G}$. First, note that it dominates the valuation $\chi_\mathcal{G}$ since, for all $u \in V_Q \setminus \{0\}$, the strategy s_A^u dominates $\chi_\mathcal{G}$ in $\mathcal{G}^u$ (recall that the stopping states in $\mathcal{G}^u$ have the values of the original states in $\mathcal{G}$). Consider now a BSCC H compatible with s_A such that $\min \chi_\mathcal{G}[H] > 0$. By Proposition 8.1, there is a value $u_H \in (0,1]$ such that $\chi_\mathcal{G}[H] = \{u_H\}$. That is, $H \subseteq Q_{u_H}$. It follows that H is compatible with $\mathsf{s}_\mathsf{A}^{u_H}$. Since $\mathsf{s}_\mathsf{A}^{u_H}$ parity dominates the valuation $\chi_\mathcal{G}$ in $\mathcal{G}^{u_H}$, we can deduce that H is even-colored. Overall, the strategy s_A parity dominates the valuation $\chi_\mathcal{G}$. By Proposition 8.1, the strategy s_A guarantees it (i.e. it is optimal). □

8.6.5 . Proof of Theorem 8.3

In this section, we give all the technical details necessary to prove Theorem 8.3. The proof of this theorem is given in Page 311, provided that Lemmas 8.6, 8.7, 8.8 and 8.9 hold. We already stated and argued why these lemmas hold, however we have not given a detailed proof of them. This is what we do in this subsection. However, we first state and prove results that we use for the remainder of this chapter.

We first state and prove properties on the update of colors NewCol. This is done in Page 316. Then, we give the complete definition of faithfulness, see Page 322. We also state and prove three lemmas, that link the local and global behaviors, , see Page 323. We finally prove the lemmas mentioned above: we prove Lemma 8.6 in Page 328, Lemma 8.7 in Page 332, Lemma 8.8 in Page 333 and Lemma 8.9 in Page 335.

Properties ensured by the local operator NewCol

We introduce a useful notation which will allow us to rewrite Lemma 8.2 in our context.

Definition 8.18. *Consider a state $q \in Q$, a virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0,e]\!]$, *a color $n \in [\![0,e]\!]$ and two* GF*-strategies $(\sigma_\mathsf{A}, \sigma_\mathsf{B}) \in \Sigma_\mathsf{A}(\mathsf{F}(q)) \times \Sigma_\mathsf{B}(\mathsf{F}(q))$. We denote by $\mathsf{Col}(q, \mathsf{vcol}, \sigma_\mathsf{A}, \sigma_B)$ the set of colors reachable in one step with positive probability w.r.t. $(\sigma_\mathsf{A}, \sigma_\mathsf{B})$ in the local game $\mathcal{G}^n_{q,\mathsf{vcol}}$ (regard-*

less of the color $n \in [\![0,e]\!]$ considered). That is:

$$\mathsf{Col}(q, \mathsf{vcol}, \sigma_\mathsf{A}, \sigma_B) := \{i \in [\![0,e]\!] \mid \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}_{q,\mathsf{vcol}}[k_i]} \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) > 0\}$$

In particular, $\mathsf{Col}(q, \mathsf{vcol}, \sigma_\mathsf{A}, b) = \mathsf{Color}(\mathsf{F}(q), p_{q,\mathsf{vcol}}, \sigma_\mathsf{A}, b)$ *for all* $b \in \mathsf{Act}_\mathsf{B}$ *(notation from Lemma 8.2).*

Then, the set of colors $\mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_\mathsf{A}, \sigma_B)$ *is defined by:*

$$\mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_\mathsf{A}, \sigma_B) := \mathsf{Col}(q, \mathsf{vcol}, \sigma_\mathsf{A}, \sigma_B) \cup \{\max(\mathsf{vcol}(q), c_n)\}$$

with $c_n := n - 1$ *is* n *is even and* $c_n := n + 1$ *is* n *is odd.*

We obtain a corollary of Lemma 8.2 (which only consists in writing Lemma 8.2 in the context of a local game $\mathcal{G}^n_{q,\mathsf{vcol}}$).

Corollary 8.11. *Consider a state* $q \in Q$*, a virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0,e]\!]$*, a color* $n \in [\![0,e]\!]$ *and a local strategy* $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathsf{F}(q))$*. A Player-*A *GF-strategy* $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathcal{F})$ *is* optimal *w.r.t.* $Y := (\mathsf{F}(q), E^n_{q,\mathsf{vcol}})$ *if and only if, letting* $u := \chi_{\mathcal{G}_Y}(q_\mathsf{init})$*, either (i)* $u = 0$*, or (ii) the positional Player-*A *strategy* $\mathsf{s}^Y_\mathsf{A}(\sigma_\mathsf{A})$ *parity dominates the valuation* v^u_Y*.*

*Furthermore (ii) is equivalent to: (1) the Player-*A *positional strategy* $\mathsf{s}^Y_\mathsf{A}(\sigma_\mathsf{A})$ *dominates the valuation* v^u_Y *and (2) for all* $b \in \mathsf{Act}_\mathsf{B}$*, if* $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[0,1]} \rangle](\sigma_\mathsf{A}, b) = 0$*, then* $\max \mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_\mathsf{A}, \sigma_B)$ *is even.*

This is symmetrical for Player B*.*

Let us now state a proposition we will use to prove that the update of colors cannot decrease the colors of the states.

Proposition 8.12. *Let* $q \in Q_u$ *and some virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0,e]\!]$*, some color* $n \in [\![0,e]\!]$ *and a positive value* $z \in (0,1]$*. Let* $p := p_{q,\mathsf{vcol}}$ *and* $Y := (\mathsf{F}(q), E^n_{q,\mathsf{vcol}})$*. For all Player-*A *GF-strategies* $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathsf{F}(q))$*, the positional Player-*A *strategy* $\mathsf{s}^\mathsf{A}_Y(\sigma_\mathsf{A})$ *dominates the valuation* v^z_Y *if and only if for all* $b \in \mathsf{Act}_\mathsf{B}$*: if* $\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q \setminus Q_u} \rangle](\sigma_\mathsf{A}, b) > 0$*, then:*

$$\mathsf{out}[\langle \mathsf{F}(q), \mu^z \rangle](\sigma_\mathsf{A}, b) \geq 0$$

where $\mu^z : Q \to [0,1]$ *is such that, for all* $q \in Q$*:*

$$\mu^z(q) := \begin{cases} 0 & \text{if } q \in Q_u \\ x - z & \text{otherwise, if } q \in Q_x \subseteq Q \setminus Q_u \end{cases}$$

Proof. In the game $\mathcal{G}_Y$, the Player-A strategy $\mathsf{s}^\mathsf{A}_Y(\sigma_\mathsf{A})$ dominates the valuation v^z_Y — property we denote (1) — if and only if:

$$\begin{aligned} z = v^z_Y(q_\mathsf{init}) &\leq \mathsf{val}[\langle \mathsf{F}(q), v^z_Y \rangle](\mathsf{s}^\mathsf{A}_Y(\sigma_\mathsf{A})(q_\mathsf{init})) \\ &= \mathsf{val}[\langle \mathsf{F}(q), v^z_Y \rangle](\sigma_\mathsf{A}) \end{aligned}$$

In addition, $\mathsf{val}[\langle \mathsf{F}(q), v_Y^z \rangle](\sigma_\mathsf{A}) = \min_{b \in \mathsf{Act_B}} \mathsf{out}[\langle \mathsf{F}(q), v_Y^z \rangle](\sigma_\mathsf{A}, b)$. Furthermore, for all $b \in \mathsf{Act_B}$, we have that

$$\begin{aligned}\mathsf{out}[\langle \mathsf{F}(q), v_Y^z \rangle](\sigma_\mathsf{A}, b) = {} & \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[\{q_\mathsf{init}\} \cup K_e]} \rangle](\sigma_\mathsf{A}, b) \cdot z \\ & + \sum_{x \in V_{Q \setminus Q_u}} \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[x]} \rangle](\sigma_\mathsf{A}, b) \cdot x\end{aligned}$$

It follows that, if $\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q \setminus Q_u} \rangle](\sigma_\mathsf{A}, b) = 0$, we have $\mathsf{out}[\langle \mathsf{F}(q), v_Y^z \rangle](\sigma_\mathsf{A}, b) = z$. However, if $\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q \setminus Q_u} \rangle](\sigma_\mathsf{A}, b) > 0$, we have that

$$\begin{aligned}z \leq \mathsf{out}[\langle \mathsf{F}(q), v_Y^z \rangle](\sigma_\mathsf{A}, b) \Leftrightarrow {} & 0 \leq \sum_{x \in V_{Q \setminus Q_u}} \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[x]} \rangle](\sigma_\mathsf{A}, b) \cdot (x - z) \\ & 0 \leq \mathsf{out}[\langle \mathsf{F}(q), \sum_{x \in V_{Q \setminus Q_u}} \mathbb{1}_{p^{-1}[x]} \cdot (x - z) \rangle](\sigma_\mathsf{A}, b) \\ & 0 \leq \mathsf{out}[\langle \mathsf{F}(q), \mu^z \rangle](\sigma_\mathsf{A}, b)\end{aligned}$$

The result follows. □

The update of colors does not decrease the color. We define a successor operation compatible with the order $\prec_\mathsf{par}$.

Definition 8.19 (Parity successor). *For all $n \in \mathbb{N}$, we have $\mathsf{Succ}(n) := n - 2$ if $n \prec_\mathsf{par} 1$, $\mathsf{Succ}(1) := 0$ and $\mathsf{Succ}(n) := n + 2$ if $1 \prec_\mathsf{par} n$.*

Let us show that the local operator NewCol does not decrease (w.r.t. the usual order $<$ on natural numbers) the previous color of the state q given by a (virtual) coloring function vcol.

Proposition 8.13. *Consider a state $q \in Q_u$ and a coloring function $\mathsf{vcol} : Q_u \to [\![0, e]\!]$. We have $\mathsf{NewCol}(q, \mathsf{vcol}) \geq \mathsf{vcol}(q)$.*

Proof. We let $p := p_{q,\mathsf{vcol}}$ and, for all $n \in [\![0, e]\!]$, we let $Y_n := (\mathsf{F}(q), E^n_{q,\mathsf{vcol}})$. Note that $p_{[0,1]} \subseteq V_{Q \setminus Q_u}$.

There are two cases: either $\mathsf{vcol}(q) = e$ or $\mathsf{vcol}(q) < e$. First, assume that $\mathsf{vcol}(q) = e$. Let us show that in that case $\mathsf{NewCol}(q, \mathsf{vcol}) = e$. Assume towards a contradiction that $\chi_{\mathcal{G}_{Y_e}}(q_\mathsf{init}) = u' < u$ for some $u' \in [0, 1]$. Consider a Player-B GF-strategy σ_B that is optimal w.r.t. Y_e. For all $a \in \mathsf{Act_A}$, we have $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{Q \setminus Q_u} \rangle](a, \sigma_\mathsf{B}) > 0$. Indeed, otherwise, in the game $\mathcal{G}^e_{q,\mathsf{vcol}}$ where Player B plays the strategy defined by σ_B, Player A could loop indefinitely on q_init thus ensuring winning with probability 1 (since the color of the state q_init is e — as $\mathsf{vcol}(q) = e$ — which is both the highest color appearing in the game and even). We let $p_\mathsf{Exit} := \min_{a \in \mathsf{Act_A}} \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{Q \setminus Q_u} \rangle](a, \sigma_\mathsf{B}) > 0$.

Now, consider some $a \in \mathsf{Act_A}$. By Proposition 8.12 for Player B, we have: $\mathsf{out}[\langle \mathcal{F}, \mu^{u'} \rangle](a, \sigma_\mathsf{B}) \leq 0$ where $\mu^{u'} : Q \to [0, 1]$ comes from Proposition 8.12. Hence, we have:

$$\sum_{x \in X} \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{Q_x} \rangle](a, \sigma_\mathsf{B}) \cdot x \leq \sum_{x \in X} \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{Q_x} \rangle](a, \sigma_\mathsf{B}) \cdot u'$$

Now, Letting $v := \chi_{\mathcal{G}}$ the value vector in the game $\mathcal{G}$, we have:

$$\begin{aligned}
\sum_{q' \in Q} \mathbb{P}^{q,q'}(a, \sigma_{\mathsf{B}}) \cdot v(q') &= \sum_{q' \in Q} \mathsf{out}[\langle \mathcal{F}, q' \rangle](a, \sigma_{\mathsf{B}}) \cdot v(q') \\
&= \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{Q_u} \rangle](a, \sigma_{\mathsf{B}}) \cdot u + \sum_{x \in V_{Q \setminus Q_u}} \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{Q_x} \rangle](a, \sigma_{\mathsf{B}}) \cdot x \\
&\leq \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{Q_u} \rangle](a, \sigma_{\mathsf{B}}) \cdot u + \sum_{x \in V_{Q \setminus Q_u}} \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{Q_x} \rangle](a, \sigma_{\mathsf{B}}) \cdot u' \\
&= \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{Q_u} \rangle](a, \sigma_{\mathsf{B}}) \cdot u + \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{Q \setminus Q_u} \rangle](a, \sigma_{\mathsf{B}}) \cdot u' \\
&= u - \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{Q \setminus Q_u} \rangle](a, \sigma_{\mathsf{B}}) \cdot (u - u') \\
&\leq u - p_{\mathsf{Exit}} \cdot (u - u') < u = v(q)
\end{aligned}$$

Hence, letting $\delta := p_{\mathsf{Exit}} \cdot (u - u')/2 > 0$ and considering, in the original game $\mathcal{G}$, a Player-B strategy $\mathsf{s_B} \in \mathsf{S}_{\mathsf{B}}^{\mathcal{C}}$ such that $\mathsf{s_B}(q) := \sigma_{\mathsf{B}}$ and for all $q' \in Q$, we have $\mathsf{s_B}(q \cdot q')$ a δ-optimal Player-B strategy from q', it follows that, for all Player-A strategy $\mathsf{s_A}$ in the game $\mathcal{G}$, we have:

$$\begin{aligned}
\mathbb{P}^{\mathcal{C},q}_{\mathsf{s_A},\mathsf{s_B}}[W] &= \sum_{q' \in Q} \mathbb{P}^{q,q'}(\mathsf{s_A}(q), \sigma_{\mathsf{B}}) \cdot \mathbb{P}^{\mathcal{C},q'}_{\mathsf{s}^q_{\mathsf{A}},\mathsf{s}^q_{\mathsf{B}}}[W] \\
&\leq \sum_{q' \in Q} \mathbb{P}^{q,q'}(\mathsf{s_A}(q), \sigma_{\mathsf{B}}) \cdot (v(q') + \delta) \leq v(q) - 2\delta + \delta = v(q) - \delta
\end{aligned}$$

where $W := (\mathsf{col}^\omega)^{-1}[\mathsf{Parity_K}] \subseteq Q^\omega$. Thus, the value from q is less than $u = v(q)$. Hence the contradiction. In fact, $\chi_{\mathcal{G}^e_{q,\mathsf{vcol}}}(q_{\mathsf{init}}) \geq u$ and $\mathsf{NewCol}(q, \mathsf{vcol}) = e$.

Consider now the case where $\mathsf{vcol}(q) < e$. Assume towards a contradiction that $\mathsf{NewCol}(q, \mathsf{vcol}) < \mathsf{vcol}(q)$. Let $n := \mathsf{NewCol}(q, \mathsf{vcol})$.

Assume that n is even. By assumption, we have $\chi_{\mathcal{G}_{Y_n}}(q_{\mathsf{init}}) \geq u$. Consider a Player-A GF-strategy σ_{A} that is optimal w.r.t. Y_n. First, the positional Player-A strategy $\mathsf{s_A} := \mathsf{s}^A_{Y_n}(\sigma_{\mathsf{A}}) = \mathsf{s}^A_{Y_{n+2}}(\sigma_{\mathsf{A}})$ defined by σ_{A} dominates the valuation $v^u_{Y_n} = v^u_{Y_{n+2}}$ in the game $\mathcal{G}_{Y_n}$ and it also does in the game $\mathcal{G}_{Y_{n+2}}$. Consider now an action $b \in \mathsf{Act_B}$ and assume that $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[0,1]} \rangle](\sigma_{\mathsf{A}}, b) = 0$. By Corollary 8.11, it follows that, in the game $\mathcal{G}_{Y_n}$, we have $\max \mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_{\mathsf{A}}, b)$ even with $\mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_{\mathsf{A}}, b) = \mathsf{Col}(q, \mathsf{vcol}, \sigma_{\mathsf{A}}, b) \cup \{e_n\}$ with $e_n = \max(\mathsf{vcol}(q), c_n)$ and $c_n = n-1$ since n is even. Since $n < \mathsf{vcol}(q)$, we have $e_n = \max(\mathsf{vcol}(q), c_n) = \mathsf{vcol}(q)$. Furthermore, $e_{n+2} = \max(\mathsf{vcol}(q), c_{n+2})$ with $c_{n+2} = n + 1$ since n is even. Since $n < \mathsf{vcol}(q)$, we also have $e_{n+2} = \mathsf{vcol}(q) = e_n$. That is, $\mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_{\mathsf{A}}, b) = \mathsf{ColBSCC}(q, \mathsf{vcol}, n{+}2, \sigma_{\mathsf{A}}, b)$. Hence, $\mathsf{ColBSCC}(q, \mathsf{vcol}, n+2, \sigma_{\mathsf{A}}, b)$ is also even. By Corollary 8.11, the Player-A GF-strategy σ_{A} strongly dominates the valuation $v^u_{Y_{n+2}}$. Hence, by Proposition 8.1, the Player-A positional strategy $\mathsf{s}^{\mathsf{A}}_{Y_{n+2}}(\sigma_{\mathsf{A}})$ guarantees the valuation $v^u_{Y_{n+2}}$. Hence the contradiction since this implies $\mathsf{NewCol}(q, \mathsf{vcol}) \geq n + 2$.

When n is odd, the reasoning is symmetrical, by taking the point-of-view of Player B: we compare what happens in the games $\mathcal{G}_{Y_m}$ and $\mathcal{G}_{Y_m}$ for $m := \mathsf{Succ}(n)$. □

The update of colors is not affected by small changes of colors. Let us now tackle another property ensured by the local operator NewCol. Assume that the new color of a state q is $n \in [\![0, e]\!]$ w.r.t. a coloring function vcol. Now, consider another coloring function vcol' that coincide with vcol on colors at least n, and may differ for smaller colors. In that case, the new color of q will still be n w.r.t. the coloring function vcol'. This is (almost) what we prove here. First, we introduce the notion of equivalent and prevailing coloring functions.

Definition 8.20 (Equivalent and Prevailing coloring functions)**.** *Consider two coloring functions* $\mathsf{vcol}, \mathsf{vcol}' : Q_u \to [\![0, e]\!]$ *and some color* $n \in [\![0, e]\!]$. *The coloring functions* $\mathsf{vcol}, \mathsf{vcol}'$ *are* equivalent down to n *if, for all* $k \in [\![n, e]\!]$, *we have* $\mathsf{vcol}^{-1}[k] = \mathsf{vcol}'^{-1}[k]$.

Furthermore, vcol' *is said to be* $(n-1)$-prevailing compared to vcol *if* vcol *and* vcol' *are equivalent down to* n *and* $\mathsf{vcol}^{-1}[n-1] \subseteq \mathsf{vcol}'^{-1}[n-1]$.

Let us first state a lemma that we will use to prove the proposition (of interest for us) that we state below.

Lemma 8.14. *Consider a state* $q \in Q_u$ *and some color* $n \in [\![0, e]\!]$. *Let* $\mathsf{vcol}, \mathsf{vcol}' : Q_u \to [\![0, e]\!]$ *be two virtual coloring functions with* vcol' *n-prevailing compared to* vcol. *Consider a pair of* GF*-strategies* $(\sigma_\mathsf{A}, \sigma_\mathsf{B}) \in \Sigma_\mathsf{A}(\mathsf{F}(q)) \times \Sigma_\mathsf{B}(\mathsf{F}(q))$. *We have:*

$$\max \mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_\mathsf{A}, \sigma_\mathsf{B}) \geq n \Rightarrow$$
$$\max \mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_\mathsf{A}, \sigma_\mathsf{B}) = \max \mathsf{ColBSCC}(q, \mathsf{vcol}', n, \sigma_\mathsf{A}, \sigma_\mathsf{B})$$

Proof. Letting $c_n := n - 1$ if n is even and $c_n := n + 1$ if n is odd, we have (recall Definition 8.18):

$$\mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_\mathsf{A}, \sigma_\mathsf{B}) = \mathsf{Col}(q, \mathsf{vcol}, \sigma_\mathsf{A}, \sigma_\mathsf{B}) \cup \{\max(\mathsf{vcol}(q), c_n)\}$$

and

$$\mathsf{ColBSCC}(q, \mathsf{vcol}', n, \sigma_\mathsf{A}, \sigma_\mathsf{B}) = \mathsf{Col}(q, \mathsf{vcol}', \sigma_\mathsf{A}, \sigma_\mathsf{B}) \cup \{\max(\mathsf{vcol}'(q), c_n)\}$$

By assumption, we have that, for all $i \in [\![n+1, e]\!]$:

$$\mathsf{vcol}^{-1}[i] = \mathsf{vcol}'^{-1}[i]$$

Thus:

$$\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}_{q,\mathsf{vcol}}[k_i]} \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) = \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}_{q,\mathsf{vcol}}[k_i]} \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B})$$

Furthermore:

$$\mathsf{vcol}^{-1}[n] \subseteq \mathsf{vcol}'^{-1}[n]$$

Thus:

$$\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}_{q,\mathsf{vcol}}[k_n]} \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) > 0 \Rightarrow \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}_{q,\mathsf{vcol}'}[k_n]} \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) > 0$$

We therefore obtain that for any $k \in [\![n, e]\!]$:

(i) If $\max \mathsf{Col}(q, \mathsf{vcol}, \sigma_\mathsf{A}, \sigma_\mathsf{B}) \leq k$, then $\max \mathsf{Col}(q, \mathsf{vcol}', \sigma_\mathsf{A}, \sigma_\mathsf{B}) \leq k$;

(ii) If $\max \mathsf{Col}(q, \mathsf{vcol}, \sigma_\mathsf{A}, \sigma_\mathsf{B}) = k$, then $\max \mathsf{Col}(q, \mathsf{vcol}', \sigma_\mathsf{A}, \sigma_\mathsf{B}) = k$.

Furthermore:

(i') If $\mathsf{vcol}(q) \leq k$, then $\mathsf{vcol}'(q) \leq k$;

(ii') If $\mathsf{vcol}(q) = k$, then $\mathsf{vcol}'(q) = k$.

We have three cases, letting $n \leq j := \max \mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_\mathsf{A}, \sigma_\mathsf{B}) = \mathsf{Col}(q, \mathsf{vcol}, \sigma_\mathsf{A}, \sigma_\mathsf{B}) \cup \{\max(\mathsf{vcol}(q), c_n)\}$:

- If $j = c_n$, then $\max \mathsf{Col}(q, \mathsf{vcol}, \sigma_\mathsf{A}, \sigma_\mathsf{B}), \mathsf{vcol}(q) \leq j$. Hence, by (i) and (i'), we have $\max \mathsf{Col}(q, \mathsf{vcol}', \sigma_\mathsf{A}, \sigma_\mathsf{B}), \mathsf{vcol}'(q) \leq j$.
  That is, $\max \mathsf{ColBSCC}(q, \mathsf{vcol}', n, \sigma_\mathsf{A}, \sigma_\mathsf{B}) = j$.
- If $j = \mathsf{vcol}(q)$ then $\max \mathsf{Col}(q, \mathsf{vcol}, \sigma_\mathsf{A}, \sigma_\mathsf{B}), c_n \leq j$. Hence, by (ii'), we have $j = \mathsf{vcol}'(q)$ and, by (i), we have $\max \mathsf{Col}(q, \mathsf{vcol}', \sigma_\mathsf{A}, \sigma_\mathsf{B}) \leq j$.
  That is, $\max \mathsf{ColBSCC}(q, \mathsf{vcol}', n, \sigma_\mathsf{A}, \sigma_\mathsf{B}) = j$.
- If $j = \max \mathsf{Col}(q, \mathsf{vcol}, \sigma_\mathsf{A}, \sigma_\mathsf{B})$ then $\mathsf{vcol}(q), c_n \leq j$. Hence, by (ii), we have $j = \max \mathsf{Col}(q, \mathsf{vcol}', \sigma_\mathsf{A}, \sigma_\mathsf{B})$ and, by (i'), we have $\mathsf{vcol}'(q) \leq j$. That is, $\max \mathsf{ColBSCC}(q, \mathsf{vcol}', n, \sigma_\mathsf{A}, \sigma_\mathsf{B}) = j$.

The lemma follows. □

Proposition 8.15. *Consider a state $q \in Q_u$ and a coloring function $\mathsf{vcol} : Q_u \to [\![0, e]\!]$. Let $n := \mathsf{NewCol}(q, \mathsf{vcol}) \in [\![0, e]\!]$. Assume that another coloring function $\mathsf{vcol}' : Q_u \to [\![0, e]\!]$ is n-prevailing compared to vcol. In that case, $\mathsf{NewCol}(q, \mathsf{vcol}') = n$.*

Proof. Let us consider two such coloring functions vcol and vcol'. We have $n = \mathsf{NewCol}(q, \mathsf{vcol}) \in [\![0, e]\!]$. Consider a Player-$\mathsf{A}$ GF-strategy σ_A that is optimal w.r.t. Y for $Y := (\mathsf{F}(d), E^n_{q,\mathsf{vcol}})$. Let $Y' := (\mathsf{F}(d), E^n_{q,\mathsf{vcol}'})$, $p := p_{q,\mathsf{vcol}}$ and $p' := p_{q,\mathsf{vcol}'}$. Let us show that $n \preceq_{\mathsf{par}} \mathsf{NewCol}(q, \mathsf{vcol}')$. This straightforwardly holds if $n = e - 1$. Assume now that $n \neq e - 1$. The Player-A strategy $\mathsf{s}^\mathsf{A}_Y(\sigma_\mathsf{A})$ dominates the valuations $v^u_Y = v^u_{Y'}$.

Consider an action $b \in \mathsf{Act_B}$ and assume that $\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[0,1]} \rangle](\sigma_\mathsf{A}, b) = 0$. By Corollary 8.11, we have $j := \max \mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_\mathsf{A}, b)$ is even with $\max \mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_{\mathsf{A}}, b) \geq \max(\mathsf{vcol}(q), c_n)$ for $c_n = n - 1$ if n is even and $c_n = n + 1$ otherwise. Since $n - 1$ is odd, it follows that, in any case, $\max \mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_{\mathsf{A}}, b) \geq n$. Then, by Lemma 8.14, we have $\max \mathsf{ColBSCC}(q, \mathsf{vcol}', n, \sigma_{\mathsf{A}}, b) = \max \mathsf{ColBSCC}(q, \mathsf{vcol}, n, \sigma_{\mathsf{A}}, b)$, which is even. As this holds for all $b \in \mathsf{Act}_{\mathsf{B}}$, the GF-strategy σ_{A} strongly dominates the valuation $v^u_{Y'}$, hence, by Proposition 8.1, the value of the state q_{init} in the game $\mathcal{G}_{Y'}$ is at least u: $\chi_{\mathcal{G}_{Y'}}[q_{\mathsf{init}}] \geq u$. Hence, $n \preceq_{\mathsf{par}} \mathsf{NewCol}(q, \mathsf{vcol}')$.

Let us now show that $\mathsf{NewCol}(q, \mathsf{vcol}') \preceq_{\mathsf{par}} n$, which straightforwardly holds if $n = e$. Hence, assume that $n \neq e$ and let $m := \mathsf{Succ}(n) \in [\![0, e]\!]$. The proof is very similar than in the previous case. Let $Z := (\mathsf{F}(d), E^m_{q,\mathsf{vcol}})$. Let also $Z' := (\mathsf{F}(d), E^m_{q,\mathsf{vcol}'})$. The value of the state q_{init} in the game $\mathcal{G}_Z$ is at most u' for some $u' < u$: $\chi_{\mathcal{G}_Z}[q_{\mathsf{init}}] \leq u' < u$. Consider a Player-$\mathsf{B}$ GF-strategy σ_{B} that is optimal w.r.t. Z. The Player-B strategy $\mathsf{s}^{\mathsf{B}}_Z(\sigma_{\mathsf{B}})$ dominates the valuation $v^{u'}_Z = v^{u'}_{Z'}$. Consider an action $a \in \mathsf{Act}_{\mathsf{A}}$ and assume that $\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[0,1]} \rangle](a, \sigma_{\mathsf{B}}) = 0$. We have $\max \mathsf{ColBSCC}(q, \mathsf{vcol}, m, a, \sigma_{\mathsf{B}})$ odd with $j := \max \mathsf{ColBSCC}(q, \mathsf{vcol}, m, a, \sigma_{\mathsf{B}}) \geq \max(\mathsf{vcol}(q), c_m)$ for $c_m = m - 1$ if m is even and $c_m = m + 1$ otherwise. Let us show that $j \geq n$.

- If m is odd (and $n = m + 2$), we have $c_m = m + 1 = n - 1$ which is even. Hence, $j \geq n$.
- If $m = 0$ (and $n = 1$), since j is odd, it must be that $j \geq n$.
- If $m \geq 2$ is even (and therefore $n = m - 2$), we have $c_m = m - 1 = n + 1$. Hence, we have $j \geq n + 1$.

We can therefore apply Lemma 8.14 to obtain that $\max \mathsf{ColBSCC}(q, \mathsf{vcol}', m, a, \sigma_{\mathsf{B}}) = \max \mathsf{ColBSCC}(q, \mathsf{vcol}, m, a, \sigma_{\mathsf{B}})$, which is odd. As this holds for all $a \in \mathsf{Act}_{\mathsf{A}}$, we have that the GF-strategy σ_{B} strongly dominating the valuation $v^{u'}_{Z'}$, hence, by Proposition 8.1, the value of the state q_{init} in the game $\mathcal{G}_{Z'}$ is at most u': $\chi_{\mathcal{G}_{Z'}}[q_{\mathsf{init}}] \leq u'$. Hence, $\mathsf{NewCol}(q, \mathsf{vcol}') \prec_{\mathsf{par}} m$.

In any case, we have $n = \mathsf{NewCol}(q, \mathsf{vcol}')$. □

Complete definition of faithfulness

We want to formally give the definition of faithfulness that we will use in the proof. First, we define the notion of coherent coloring and environment functions. Informally, for each state $q \in Q_u$: the colors given by the coloring function correspond to the environment provided at each state (the environments being defined from coloring functions as in Definition 8.13).

Definition 8.21 (Coherent coloring and environment functions). *Consider a virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$*, some color* $n \in [\![0, e]\!]$ *and let* $Q_n := \mathsf{vcol}^{-1}[n]$. *Consider an environment function* $\mathsf{Ev} : Q_n \to \mathsf{Env}(Q)$ *and, for all* $q \in Q_n$*, let* $\mathsf{vcol}_q : Q_u \to [\![0, e]\!]$ *be the coloring function associated with* *the environment* $\mathsf{Ev}(q)$. *Let* $q \in Q_n$. *We say that* $(\mathsf{vcol}, \mathsf{Ev})$ *is* coherent *at state* q *if, letting* $n_q := \mathsf{NewCol}(q, \mathsf{vcol}_q)$*:*

- $\mathsf{col}(q), \mathsf{vcol}_q(q), n_q \leq n$ *and the coloring function* vcol *is* n*-prevailing compared to the coloring function* vcol_q*;*
- $n \equiv n_q \mod 2$ *and* $\mathsf{Ev}(q) = \mathsf{CreateEnv}(n_q, q, \mathsf{vcol}_q)$ *where* $\mathsf{CreateEnv}$ *corresponds to Algorithm 8.9.*

If this holds for all $q \in Q_n$, $(\mathsf{vcol}, \mathsf{Ev})$ *is* coherent *at color* n.

Definition 8.22 (Faithful pair of coloring and environment functions)**.** *Consider a virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$*, some* $n \in [\![0, e+1]\!]$ *and a partial environment function* $\mathsf{Ev} : Q_u \to \mathsf{Env}(Q)$ *defined on* $\mathsf{vcol}^{-1}[[\![n, e]\!]]$. *We say that* $(\mathsf{vcol}, \mathsf{Ev})$ *is* faithful down to n *if:*

1f. for all $k \in [\![n, e]\!]$*, the pair* $(\mathsf{vcol}, \mathsf{Ev})$ *witnesses the color* k*;*

2f. for all $k \in [\![n, e]\!]$*, the pair* $(\mathsf{vcol}, \mathsf{Ev})$ *is coherent at color* k*;*

3f. for all $q \in Q_u$*, if* $\mathsf{vcol}(q) < n$*, then we have* $\mathsf{col}(q) = \mathsf{vcol}(q)$ *and* $\mathsf{NewCol}(q, \mathsf{vcol}) < n$*;*

When $n = 0$*, we say that the pair* $(\mathsf{vcol}, \mathsf{Ev})$ *is* completely faithful.

Three central lemmas

In the following, we state three lemmas that we will use in remainder of this chapter. The first lemma relates probability distributions in a local game and in a global game. This is particularly useful as it allows to use the assumptions made on the local strategies to obtain various properties on global games where such local strategies are used.

The second lemma states that any Player-A positional strategy generated by an environment such that the new color (w.r.t. that environment) is even dominates a specific valuation. This is analogous for Player B.

Finally, the third lemma states that any Player-A local strategy defined by local strategies that are optimal w.r.t. to an even color ensure that in a BSCC H compatible with such a strategy, if some high enough color occurs in H, then H is even-colored. This is analogous for Player B.

Lemma 8.16. *Consider a virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$*, some color* $n \in [\![0, e]\!]$ *and a partial environment function* $\mathsf{Ev} : Q_u \to \mathsf{Env}(Q)$ *defined over* Q_n *for* $Q_n := \mathsf{vcol}^{-1}[\{n\}]$. *Let* $q \in Q_n$ *and* $\mathsf{vcol}_q : Q_u \to [\![0, e]\!]$ *be the coloring function associated with environment* $\mathsf{Ev}(q)$. *We also let* $p := p_{q,\mathsf{vcol}_q}$. *Then, for all Player-*A *and Player-*B *strategies* $\mathsf{s_A}$ *and* $\mathsf{s_B}$ *in the arena* $\mathcal{C}^n_{\mathsf{vcol}}$*, we have:*

$$\forall x \in V_{Q \setminus Q_u},\ \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[x] = \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[x]} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \tag{8.1}$$

and

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[Q_n \cup K^n] = \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[\{q_{\mathsf{init}}\} \cup K_e]} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \tag{8.2}$$

Furthermore, if vcol *is equivalent down to* $n+1$ *to* vcol_q*:*

$$\forall i \in [\![n+1, e]\!],\ \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[k^n_i] = \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[k_i]} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \tag{8.3}$$

and if, in addition, vcol *is* n*-prevailing compared to* vcol_q*:*

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[\mathsf{vcol}_q^{-1}[n] \setminus \{q\}] = \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[k_n]} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \tag{8.4}$$

Proof. By Definition 8.11, we have:

- for all $q' \in Q_n$, $p^{-1}_{n,\mathsf{vcol}}[q'] = \{q'\}$;
- for all $i \in [\![0, e]\!]$: $p^{-1}_{n,\mathsf{vcol}}[k^n_i] = \mathsf{vcol}^{-1}[i] \setminus Q_n$;
- for all $x \in V_{Q \setminus Q_u}$, $p^{-1}_{n,\mathsf{vcol}}[x] = Q_x$.

Furthermore, from Definition 8.13, we have:

- $p^{-1}[q_{\mathsf{init}}] = \{q\}$;
- for all $i \in [\![0, e]\!]$: $p^{-1}[k_i] = \mathsf{vcol}^{-1}[i] \setminus \{q\}$;
- for all $x \in V_{Q \setminus Q_u}$, $p^{-1}[x] = Q_x = p^{-1}_{n,\mathsf{vcol}}[x]$.

By Definition 8.12, for all $x \in V_{Q \setminus Q_u}$. We have:

$$\begin{aligned}
\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}(x) &= \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}_{n,\mathsf{vcol}}[x]} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \\
&= \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[x]} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q))
\end{aligned}$$

This proves Equation (8.1).

Furthermore, using this Equation for the second equality, we have:

$$\begin{aligned}
\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[Q_n \cup K^n] &= 1 - \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[V_{Q \setminus Q_u}] \\
&= 1 - \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[V_{Q \setminus Q_u}]} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \\
&= \mathsf{out}[\langle \mathsf{F}(q), 1 - \mathbb{1}_{p^{-1}[V_{Q \setminus Q_u}]} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \\
&= \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[\{q_{\mathsf{init}}\} \cup K_e]} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q))
\end{aligned}$$

This proves Equation (8.2).

Assume now that vcol is equivalent down to $n+1$ to vcol_q. Let $i \in [\![n+1, e]\!]$. We have $\mathsf{vcol}^{-1}[i] = \mathsf{vcol}_q^{-1}[i]$ (recall Definition 8.20). Furthermore, for all $q' \in Q_n$, we have $\mathsf{vcol}(q') = n \neq i$. Hence, $\mathsf{vcol}^{-1}[i] \setminus Q_n = \mathsf{vcol}^{-1}[i] = \mathsf{vcol}^{-1}[i] \setminus \{q\}$. Hence we have, again by Definition 8.12, for all $i \in [\![n+1, e]\!]$:

$$\begin{aligned}
\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}(k^n_i) &= \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}_{n,\mathsf{vcol}}[k^n_i]} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \\
&= \mathsf{out}[\langle \mathsf{F}(q), \mathsf{vcol}^{-1}[i] \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \\
&= \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[k_i]} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q))
\end{aligned}$$

This gives Equation (8.3). Assume in addition that vcol is n-prevailing compared to vcol_q (recall Definition 8.20). That is, $\mathsf{vcol}_q^{-1}[n] \subseteq \mathsf{vcol}^{-1}[n] = Q_n$. Hence, $\mathsf{vcol}_q^{-1}[n] \setminus \{q\} \subseteq Q_n$. It follows that:

$$\begin{aligned}
\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[\mathsf{vcol}_q^{-1}[n] \setminus \{q\}] &= \sum_{q' \in \mathsf{vcol}_q^{-1}[n]\setminus\{q\}} \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}_{n,\mathsf{vcol}}[q']}\rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \\
&= \sum_{q' \in \mathsf{vcol}_q^{-1}[n]\setminus\{q\}} \mathsf{out}[\langle \mathsf{F}(q), q'\rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \\
&= \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{\mathsf{vcol}_q^{-1}[n]\setminus\{q\}}\rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \\
&= \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[k_n]}\rangle](\mathsf{s_A}(q), \mathsf{s_B}(q))
\end{aligned}$$

We obtain Equation (8.4). □

Lemma 8.17. *Consider a virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$*, some color* $n \in [\![0, e]\!]$ *and let* $Q_n := \mathsf{vcol}^{-1}[n]$*. Consider an environment function* $\mathsf{Ev} : Q_n \to \mathsf{Env}(Q)$*. Assume that* $(\mathsf{vcol}, \mathsf{Ev})$ *is coherent at color* n*. In the arena* $\mathcal{C}^n_{\mathsf{vcol}}$*:*

- *if* n *is even, then all positional Player-*A *strategies generated by the environment* Ev *dominate the valuation* $v^u_{n,\mathsf{vcol}}$ *in the arena* $\mathcal{C}^n_{\mathsf{vcol}}$ *(see Definition 8.16);*
- *if* n *is odd, then there is some* $y < u$ *such that, for all* $z \geq y$*, all positional Player-*B *strategies generated by the environment* Ev *dominate the valuation* $v^z_{n,\mathsf{vcol}}$ *in the arena* $\mathcal{C}^n_{\mathsf{vcol}}$ *(see Definition 8.16).*

Proof. For all states $q \in Q_n$, we let $\mathsf{vcol}_q : Q_u \to [\![0, e]\!]$ be the coloring function associated with the environment $\mathsf{Ev}(q)$ and we let $p := p_{q,\mathsf{vcol}_q}$. Let $q \in Q_n$. For all Player-A and B strategies $\mathsf{s_A}, \mathsf{s_B}$ in the arena $\mathcal{C}^n_{\mathsf{vcol}}$, for all $z \in [0, 1]$, we have:

$$\mathsf{out}[\langle \mathsf{F}(q), v^z_{n,\mathsf{vcol}}\rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) = z \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[Q_n \cup K] + \sum_{x \in V_{Q\setminus Q_u}} x \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}(x)$$

Assume that n is even and consider a Player-A positional strategy $\mathsf{s_A}$ generated by the environment function Ev. Letting $Y_q := (\mathsf{F}(q), \mathsf{Ev}(q))$, in the game $\mathcal{G}_{Y_q}$, the positional Player-A strategy $\mathsf{s}^{\mathsf{A}}_{Y_q}(\mathsf{s_A}(q))$ dominates the valuation $v^{u'}_{Y_q}$ for $u' := \chi_{\mathcal{G}_{Y_q}}[q_{\mathsf{init}}] \geq u$ (recall Lemma 8.2) since $(\mathsf{vcol}, \mathsf{Ev})$ is coherent at color n. Hence:

$$\begin{aligned}
u' &\leq \mathsf{out}[\langle \mathsf{F}(q), v^{u'}_{Y_q}\rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \\
&= \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[\{q_{\mathsf{init}}\}\cup K_e]}\rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \cdot u' \\
&+ \sum_{x \in V_{Q\setminus Q_u}} \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{p^{-1}[x]}\rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) \cdot x \\
&= u' \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[Q_n \cup \mathsf{K}^n] + \sum_{x \in V_{Q\setminus Q_u}} x \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[x]
\end{aligned}$$

Then, we have (since $u' \geq u$):

$$\begin{aligned}
u = u' + (u - u') &\leq u' \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[Q_n \cup K^n] + \sum_{x \in V_{Q \setminus Q_u}} x \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[x] + (u - u') \\
&\leq u' \cdot \mathbb{P}^{\mathcal{C}^n_{\mathsf{vcol}},q}_{\mathsf{s_A},\mathsf{s_B}}[Q_n \cup K^n] + \sum_{x \in V_{Q \setminus Q_u}} x \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[x] + (u - u') \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[Q_n \cup K^n] \\
&= u \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[Q_n \cup K^n] + \sum_{x \in V_{Q \setminus Q_u}} x \cdot \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[x] \\
&= \mathsf{out}[\langle \mathsf{F}(q), v^u_{n,\mathsf{vcol}} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q))
\end{aligned}$$

Since this holds for all Player-B positional strategies strategies $\mathsf{s_B}$ and for all $q \in Q_n$, it follows that the Player-A strategy $\mathsf{s_A}$ dominates the valuation $v^u_{n,\mathsf{vcol}}$ in the arena $\mathcal{C}^n_{\mathsf{vcol}}$.

In the case where n is odd, the proof is analogous, from Player-B's point-of-view. ☐

Lemma 8.18. *Consider a virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$*, some* $n \in [\![0, e]\!]$ *and* $Q_n := \mathsf{vcol}^{-1}[\{n\}]$*. Consider an environment function* $\mathsf{Ev} : Q_n \to \mathsf{Env}(Q)$ *and assume that* $(\mathsf{vcol}, \mathsf{Ev})$ *is coherent at the color* n*. Then, denoting* $K^{\geq n} := \{k^n_i \mid i \in [\![n, e]\!]\}$*:*

- *if* n *is even, then in the arena* $\mathcal{C}^n_{\mathsf{vcol}}$*, all positional Player-A strategies generated by the environment* Ev *ensure that for all BSCCs* H *compatible with* $\mathsf{s_A}$*, if* $K^{\geq n}$ *occurs in* H*, then* H *is even-colored;*

- *if* n *is odd, then in the arena* $\mathcal{C}^n_{\mathsf{vcol}}$*, all positional Player-B strategies generated by the environment* Ev *ensure that for all BSCCs* H *compatible with* $\mathsf{s_B}$*, if* $K^{\geq n}$ *occurs in* H*, then* H *is odd-colored.*

Proof. Assume that n is even and consider a Player-A positional strategy $\mathsf{s_A}$ generated by the environment Ev in the game $\mathcal{L}^n_{\mathsf{vcol}}$ (recall Definition 8.12). Consider a Player-B positional deterministic strategy $\mathsf{s_B}$ and consider a BSCC H compatible with $\mathsf{s_A}$ and $\mathsf{s_B}$.

For all $q \in Q_n$, let $\mathsf{vcol}_q : Q_u \to [\![0, e]\!]$ be the coloring function associated with the environment $\mathsf{Ev}(q)$. Let $q \in Q_n$ and $n_q := \mathsf{NewCol}(q, \mathsf{vcol}_q)$. Since $(\mathsf{vcol}, \mathsf{Ev})$ is coherent at n, n_q is even and less than or equal to n. Recall that $\mathcal{G}^{n_q}_{q,\mathsf{vcol}} = \mathcal{G}_{Y_q}$, for $Y_q := (\mathsf{F}(q), E^{n_q}_{q,\mathsf{vcol}_q})$ (see Definition 8.13). Recall also that $E^{n_q}_{q,\mathsf{vcol}_q} = \langle \max(c_{n_q}, \mathsf{vcol}_q(q)), e, p_{\{q\},\mathsf{vcol}_q} \rangle$ (also see Definition 8.13) with $c_{n_q} = n_q - 1$ as n_q is even. Let $p^q := p_{q,\mathsf{vcol}_q}$. Note that $p^q_{[0,1]} \subseteq V_{Q \setminus Q_u}$ (recall Definition 8.13). By Equation (8.1) from Lemma 8.16, if q is in H, then $\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q \setminus Q_u} \rangle](\mathsf{s_A}(q), \mathsf{s_B}(q)) = 0$. In that case, since the Player-A GF-strategy $\mathsf{s_A}(q)$ is optimal w.r.t. Y_q (recall Lemma 8.2), we have the following.:

$$\max(\mathsf{Color}(\mathsf{F}(q), p^q, \mathsf{s_A}(q), \mathsf{s_B}(q)), n_q - 1, \mathsf{vcol}_q(q)) \text{ is even} \tag{8.5}$$

Since $(\mathsf{vcol}, \mathsf{Ev})$ is coherent (see Definition 8.21) at n, we have $\mathsf{col}(q) \leq n$ for all $q \in Q_n$. Hence, our goal is to show that:

$$K^{\geq n} \text{ occurs in } H \Rightarrow \max M \text{ is even}$$

with $M := \{i \in [\![n, e]\!] \mid k_i^n \text{ occurs in } H\}$.

To prove this, we use the following characterization: for all colors $i \in [\![n+1, e]\!]$, by Equation (8.3) in Lemma 8.16 since vcol and vcol_q are equivalent down to $n+1$, we have the following equivalence:

$$k_i^n \text{ occurs in } H \Leftrightarrow \exists q \in H,\ i \in \mathsf{Color}(\mathsf{F}(q), p^q, \mathsf{s_A}(q), \mathsf{s_B}(q)) \tag{8.6}$$

Furthermore, for all $q \in H$, by assumption (the pair $(\mathsf{vcol}, \mathsf{Ev})$ being coherent), $\mathsf{vcol}_q(q) \leq n$. Hence, since $n_q \leq n$, Equation (8.5) gives that, for all $q \in H$:

$$\max \mathsf{Color}(\mathsf{F}(q), p^q, \mathsf{s_A}(q), \mathsf{s_B}(q)) \geq n{+}1 \Rightarrow \max \mathsf{Color}(\mathsf{F}(q), p^q, \mathsf{s_A}(q), \mathsf{s_B}(q)) \text{ is even} \tag{8.7}$$

Furthermore, if $\max M = n$, then n is the highest color appearing in H (since vcol and Ev are coherent at n) and H is then even-colored. Otherwise:

$$\begin{aligned}
K^{\geq n} \text{ occurs in } H \Leftrightarrow\ & M \neq \emptyset \\
\Rightarrow\ & M \neq \emptyset \wedge k_m^n \text{ occurs in } H \text{ for } m := \max M \\
\Rightarrow\ & M \neq \emptyset \wedge \exists q_m \in H,\ m \in \mathsf{Color}(\mathsf{F}(q_m), p, \mathsf{s_A}(q_m), \mathsf{s_B}(q_m)) \\
& \text{by Equation (8.6)} \\
\Rightarrow\ & M \neq \emptyset \wedge \exists q_m \in H,\ \max \mathsf{Color}(\mathsf{F}(q_m), p^{q_m}, \mathsf{s_A}(q_m), \mathsf{s_B}(q_m)) = m \\
& \text{since } m = \max M \text{ and by Equation (8.6)} \\
\Rightarrow\ & M \neq \emptyset \wedge m = \max M \text{ is even} \\
& \text{since } m \geq n+1 \text{ and by Equation (8.7)}
\end{aligned}$$

We obtain the desired result.

The case of n odd is analogous by reversing the instances of 'odd' and 'even'. However, the arguments for obtaining the analogue of Equation (8.5) is slightly different. Indeed, letting $Y_q := (fsf(q), E^{\mathsf{Succ}(n_q)}_{\{q\}, \mathsf{vcol}_q})$, the Player-B GF-strategy $\mathsf{s_B}(q)$ is optimal w.r.t. Y_q (recall Lemma 8.2 and the definition of $\mathsf{CreateEnv}$ (i.e. Algorithm 8.9)). Recall that $E^{\mathsf{Succ}(n_q)}_{\{q\}, \mathsf{vcol}_q} = \langle \max(c_{\mathsf{Succ}(n_q)}, \mathsf{vcol}_q(q)), e, \mathsf{vcol}_q \rangle$ where $c_{\mathsf{Succ}(n_q)} = -1$ if $n_q = 1$ and $c_{\mathsf{Succ}(n_q)} = n_q - 1$ otherwise (i.e. if $n_q \geq 3$). Overall, from Lemma 8.2, we do obtain:

$$\max(\mathsf{Color}(\mathsf{F}(q), p, \mathsf{s_A}(q), \mathsf{s_B}(q)), n_q - 1, \mathsf{vcol}_q(q)) \text{ is odd} \tag{8.8}$$

with $n_q - 1$ even. □

Proof of Lemma 8.6

Let us define a notion which is slightly weaker then being faithful: being non-deceiving.

Definition 8.23 (Non-deceiving environment and coloring functions)**.** *Consider a virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$*, a partial environment function* $\mathsf{Ev} : Q_u \to \mathsf{Env}(Q)$ *and some color* $n \in [\![0, e+1]\!]$*. We say that the pair* $(\mathsf{vcol}, \mathsf{Ev})$ *is* non-deceiving down to n *if:*

1n-d. for all $k \in [\![n, e]\!]$*, the pair* $(\mathsf{vcol}, \mathsf{Ev})$ *witnesses the color* k*;*

2n-d. for all $k \in [\![n, e]\!]$*, the pair* $(\mathsf{vcol}, \mathsf{Ev})$ *is coherent at color* k*;*

3n-d. for all $q \in Q_u$*, if* $\mathsf{vcol}(q) < n$*, then we have* $\mathsf{col}(q) = \mathsf{vcol}(q)$ *and* $\mathsf{NewCol}(q, \mathsf{vcol}) \leq n$*.*

The difference with being faithful lies in the fact that some state $q \in Q_u$ *with* $\mathsf{vcol}(q) < n$ *could be such that* $\mathsf{NewCol}(q, \mathsf{vcol}) = n$ *(which is not possible with a faithful pair). When* $n = 0$*, being non-deceiving down to* n *is equivalent to being faithful down to* n *(i.e. to being completely faithful).*

Then, Algorithm 8.12 is composed of Algorithm 8.10 and Algorithm 8.11. In fact, Algorithm 8.10 transforms a faithful pair into a non-deceiving one (one level below). And Algorithm 8.11 transforms a non-deceiving pair into a faithful one (at the same level). We state one lemma per algorithm formally stating the specifications of these algorithms.

Lemma 8.19. *Consider a virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$*, some color* $n \in [\![1, e+1]\!]$ *and a partial environment function* $\mathsf{Ev} : Q_u \to \mathsf{Env}(Q)$ *defined on* $\mathsf{vcol}^{-1}[[\![n, e]\!]]$*. Assume that* $(\mathsf{vcol}, \mathsf{Ev})$ *is faithful down to* n*. Let* $\mathsf{Ev}' \leftarrow \mathsf{UpdCurSta}(n-1, \mathsf{vcol}, \mathsf{Ev})$*. Then,* $(\mathsf{vcol}, \mathsf{Ev}')$ *is non-deceiving down to* $n-1$*.*

Proof. For all $q \in Q_u$, such that $\mathsf{vcol}(q) \geq n-1$, we denote by vcol_q the coloring function corresponding to the environment $\mathsf{Ev}'(q)$.

2n-d. Let us show that $(\mathsf{vcol}, \mathsf{Ev}')$ is coherent at $n-1$. We let $Q_{n-1} := \mathsf{vcol}^{-1}[n-1]$ and $q \in Q_{n-1}$. Since $(\mathsf{vcol}, \mathsf{Ev})$ is faithful down to n, we have $\mathsf{vcol}(q) = \mathsf{col}(q) = n-1$. Furthermore, $\mathsf{vcol}_q = \mathsf{vcol}$. Hence, $\mathsf{vcol}_q = (q) = \mathsf{col}(q) = n-1$. In addition, since the pair $(\mathsf{vcol}, \mathsf{Ev})$ is faithful down to n, we have $\mathsf{NewCol}(q, \mathsf{vcol}) < n$. By Proposition 8.13, we have $n-1 = \mathsf{vcol}(q) \leq \mathsf{NewCol}(q, \mathsf{vcol}) < n$. Hence, $\mathsf{vcol}(q) = \mathsf{NewCol}(q, \mathsf{vcol}) = n-1$. By definition of Algorithm 8.9, we have $\mathsf{Ev}'(q) = \mathsf{CreateEnv}(n-1, q, \mathsf{vcol})$. As this holds for all $q \in Q_{n-1}$, $(\mathsf{vcol}, \mathsf{Ev})$ is coherent at $n-1$.

3n-d. This condition straightforwardly holds since the coloring function has not changed and the pair $(\mathsf{vcol}, \mathsf{Ev})$ is faithful down to n.

1n-d. This condition straightforwardly holds for $k \geq n$ since the pair $(\mathsf{vcol}, \mathsf{Ev})$ is faithful down to n. Note that since $(\mathsf{vcol}, \mathsf{Ev}')$ is coherent at $n-1$, both Lemma 8.17 and Lemma 8.18 can be applied to $(\mathsf{vcol}, \mathsf{Ev}')$ at $n-1$.

Assume that $n-1$ is even. Consider a Player-A positional strategy $\mathsf{s_A}$ generated by the environment Ev in the game $\mathcal{L}^{n-1}_{\mathsf{vcol}}$. By Lemma 8.17, this strategy dominates the valuation $v^u_{n-1,\mathsf{vcol}}$ in the arena $\mathcal{C}^{n-1}_{\mathsf{vcol}}$. Consider a Player-B positional deterministic strategy $\mathsf{s_B}$ and consider a BSCC H in $\mathcal{C}^{n-1}_{\mathsf{vcol}}$ compatible with $\mathsf{s_A}$ and $\mathsf{s_B}$. By Lemma 8.18, if $K^{\geq n-1}$ occurs in H, then H is even-colored. Assume now that $K^{\geq n-1}$ does not occur in H. In that case, there are some states in H in $Q_{n-1} \subseteq \mathsf{col}^{-1}[n-1]$ and possibly some states in $\{k_i^{n-1} \mid i \in [\![0, n-2]\!]\}$ occur in H with $\mathsf{col}^{n-1}(k_i^{n-1}) = n-2$ for all $i \in [\![0, n-2]\!]$. Hence, the highest color appearing in H is $n-1$ and therefore H is even-colored. Thus, $\mathsf{s_A}$ parity dominates the valuation $v^u_{n-1,\mathsf{vcol}}$.

The case where $n-1$ is odd is analogous, from Player-B point-of-view.

□

Lemma 8.20. *Consider a virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0, e]\!]$*, some color* $n \in [\![0, e]\!]$*, and a partial environment function* $\mathsf{Ev} : Q_u \to \mathsf{Env}(Q)$ *defined on* $[\![n, e]\!]$*. Assume that the pair* $(\mathsf{vcol}, \mathsf{Ev})$ *is non-deceiving down to* n*. Let* $(\mathsf{vcol}', \mathsf{Ev}') \leftarrow \mathsf{UpdNewSta}(n, \mathsf{vcol}, \mathsf{Ev})$*. Then,* $(\mathsf{vcol}', \mathsf{Ev}')$ *is faithful down to* n*.*

Proof. For all $q \in Q_u$ such that $\mathsf{vcol}(q) \geq n$, we denote by vcol_q the coloring function associated with the environment $\mathsf{Ev}'(q)$. Let $Q_n := (\mathsf{vcol}')^{-1}[n]$.

2f. For all $k \in [\![n+1, e]\!]$, the pair $(\mathsf{vcol}', \mathsf{Ev}')$ is coherent at k since the pair $(\mathsf{vcol}, \mathsf{Ev})$ is non-deceiving down to n. Now, let $q \in Q_n$. If $\mathsf{vcol}(q) = n$, then straightforwardly, $(\mathsf{vcol}', \mathsf{Ev}')$ is coherent at state q since $(\mathsf{vcol}, \mathsf{Ev})$ is non-deceiving down to n, hence vcol is n-prevailing compared to vcol_q and vcol' is also n-prevailing compared to vcol. Consider now a state $q \in Q_n$ such that $\mathsf{vcol}(q) < n$. By definition of Algorithm 8.11, we have $\mathsf{vcol}(q) = \mathsf{vcol}_q(q) \leq \mathsf{NewCol}(q, \mathsf{vcol}_q) = n$ (by Proposition 8.13). That is, $\mathsf{vcol}_q(q) \leq n-1$. Since $(\mathsf{vcol}, \mathsf{Ev})$ is non-deceiving down to n, we have $\mathsf{col}(q) = \mathsf{vcol}(q) = \mathsf{vcol}_q(q) \leq n-1$. Furthermore, by definition of Algorithm 8.11, the colors of the states in $\mathsf{vcol}^{-1}[[\![n, e]\!]]$ is not changed and more and more states are colored with n. Hence, vcol' is n-prevailing compared to vcol_q. In addition, $n = \mathsf{NewCol}(q, \mathsf{vcol}_q)$ and $\mathsf{Ev}(q) = \mathsf{CreateEnv}(n, q, \mathsf{vcol}_q)$. It follows that the pair $(\mathsf{vcol}', \mathsf{Ev}')$ is coherent at state q, and this holds for all $q \in Q_n$.

3f. By definition of Algorithm 8.11, for all $k < n$ and state $q \in Q_u$ such that $\mathsf{vcol}'(q) = k$, we have $\mathsf{NewCol}(q, \mathsf{vcol}') \neq n$ and $\mathsf{vcol}'(q) = \mathsf{vcol}(q) = \mathsf{col}(q)$

(since $(\mathsf{vcol}, \mathsf{Ev})$ is non-deceiving down to n). Furthermore, assume towards a contradiction that a state $q \in Q_u$ is such that $\mathsf{vcol}'(q) = k = \mathsf{vcol}(q)$ and $\mathsf{NewCol}(q, \mathsf{vcol}') = m > n$. In that case, by Proposition 8.15, since vcol and vcol' are equivalent down to m, we would have $\mathsf{NewCol}(q, \mathsf{vcol}) = m > n$, which is not possible since $(\mathsf{vcol}, \mathsf{Ev})$ is non-deceiving down to n. In fact, $\mathsf{NewCol}(q, \mathsf{vcol}') < n$.

1f. Let $k \geq n+1$. We have $\mathsf{vcol}^{-1}[k] = (\mathsf{vcol}')^{-1}[k]$ and for all $q \in \mathsf{vcol}^{-1}[k]$, we have $\mathsf{Ev}(q) = \mathsf{Ev}'(q)$. Hence, the first condition for faithfulness holds for k since $(\mathsf{vcol}, \mathsf{Ev})$ is non-deceiving down to n. We consider now the case where $k = n$. Assume that n is even. Consider a Player-A positional strategy $\mathsf{s_A}$ generated by the environment function Ev in the game $\mathcal{L}^n_{\mathsf{vcol}'}$. Let us show that this strategy parity dominates the valuation $v^u_{n,\mathsf{vcol}'}$ in the game $\mathcal{L}^n_{\mathsf{vcol}'}$. By Lemma 8.17, this strategy dominates the valuation $v^u_{n,\mathsf{vcol}'}$. Now, fix a Player-B positional deterministic strategy $\mathsf{s_B}$ and consider a BSCC H compatible with $\mathsf{s_A}$ and $\mathsf{s_B}$ in $\mathcal{L}^n_{\mathsf{vcol}'}$. If $K^{\geq n}$ occurs in H, then by Lemma 8.18, the BSCC H is even-colored. Assume now that $K^{\geq n}$ does not occur in H. Let $X^0_n := \mathsf{vcol}^{-1}[n]$ and let $j := |Q_n| - |X^0_n|$. For all $1 \leq i \leq j$, we denote by q_i the i-th element of Q_n whose color was changed by Algorithm 8.11 and $X^i_n := X^0_n \cup \cup_{1 \leq k \leq i}\{q_i\}$. In particular, $X^j_n = Q_n$ and:

$$\forall i \in [\![1, j]\!],\ \mathsf{vcol}^{-1}_{q_i}[n] = X^{i-1}_n \tag{8.9}$$

First, assume that $H \subseteq X^0_n$. In that case, this is a BSCC compatible with the strategy $\mathsf{s_A}$ in the game $\mathcal{L}^n_{\mathsf{vcol}}$ in which this strategy, by assumption, parity dominates the valuation $v^u_{n,\mathsf{vcol}}$. Hence, this BSCC is even-colored. Second, assume that $H \not\subseteq X^0_n$ and let $C_n := Q_n \cap \mathsf{col}^{-1}[n]$. For all $i \in [\![0, j]\!]$, let $H_i := H \cap X^i_n$. Let us show by induction on $i \in [\![0, j]\!]$ the following property $\mathcal{P}(i)$: for all $q \in H_i \setminus C_n$, there is a positive probability to visit C_n from q while only visiting states in H_i, that is:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}'},q}[H_i^* \cdot C_n] > 0$$

Assume towards a contradiction that $\mathcal{P}(0)$ does not hold. That is, there is some $q \in H_0 \setminus C_n$ such that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}'},q}[H_0^* \cdot C_n] = 0$. Let $Z \subseteq H_0$ be the set of states in H_0 such that $Z := \{q\} \cup \{q' \in H_0 \mid \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}'},q}[H_0^* \cdot q'] > 0\}$. Note that $Z \cap C_n = \emptyset$. Our goal is now to exhibit a BSCC in $\mathcal{L}^n_{\mathsf{vcol}}$ that is compatible with $\mathsf{s_A}$ and that is odd-colored. This will be a contradiction with the fact that $\mathsf{s_A}$ parity dominates the valuation $v^u_{n,\mathsf{vcol}}$ in the game $\mathcal{L}^n_{\mathsf{vcol}}$. Consider a Player-$\mathsf{B}$ strategy s'_B such that, for all $q' \in X^0_n$, $\mathsf{s}'_\mathsf{B}(q') := \mathsf{s_B}(q')$ and for all $i \in [\![0, n-1]\!]$, we have $\mathsf{s_B}(k^n_i)$ playing deterministically to reach q. By construction, we do have a BSCC H' compatible with $\mathsf{s_A}$ and s'_B in $\mathcal{L}^n_{\mathsf{vcol}}$ that is included in $Z \cup \{k^n_i \mid i \in [\![0, n-1]\!]\}$. Furthermore, since the BSCC H in the game

$\mathcal{L}^n_{\mathsf{vcol}'}$ is not equal to H_0 (because this would imply $H \subseteq X_n^0$), it must be that at least one state in $\{k_i^n \mid i \in [\![0, n-1]\!]\}$ is in H'. Then, since each state in $\{k_i^n \mid i \in [\![0, n-1]\!]\}$ is colored with $n-1$ and no state in Z is in C_n, it follows that the highest color in H is $n-1$, which is odd. That is, the BSCC H' his odd-colored, which is a contradiction. In fact, $\mathcal{P}(0)$ holds.

Assume now that $\mathcal{P}(i)$ holds for some $0 \leq i \leq j-1$. Consider some state $q \in H_{i+1} \setminus C_n$. If $q \in H_i$, then we can apply $\mathcal{P}(i)$. Assume now that $q \notin H_i$, that is $q \notin X_n^i$. Hence, $q = q_{i+1}$. Since $q \notin C_n$, we have $\mathsf{col}(q) \neq n$, and therefore $\mathsf{col}(q) \leq n - 1$[4]. Furthermore, letting $Y_q := (\mathsf{F}(q), E^n_{\{q\},\mathsf{vcol}_q})$, the Player-A GF-strategy $\mathsf{s_A}(q)$ is optimal w.r.t. Y_q (recall Lemma 8.2). Hence, letting $p^q := p_{q,\mathsf{vcol}_q}$, we have:

$$\max(\mathsf{Color}(\mathsf{F}(q), p^q, \mathsf{s_A}(q), \mathsf{s_B}(q)), n-1) \text{ is even} \tag{8.10}$$

with $n-1$ odd. Recall that $\mathsf{Color}(\mathsf{F}(q), p^q, \mathsf{s_A}(q), \mathsf{s_B}(q)) := \{i \in [\![0, e]\!] \mid \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{(p^q)^{-1}[k_i]}\rangle](\mathsf{s_A}(q), b) > 0\}$. In addition, since we assume that $K^{\geq n}$ does not occur in H, we have $\max \mathsf{Color}(\mathsf{F}(q), p^q, \mathsf{s_A}(q), \mathsf{s_B}(q)) \leq n$ by Equation (8.3) from Lemma 8.16. We can conclude that we have $\max \mathsf{Color}(\mathsf{F}(q), p^q, \mathsf{s_A}(q), \mathsf{s_B}(q)) = n$. Hence, by definition, we have that $\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{(p^q)^{-1}[k_n]}\rangle](\mathsf{s_A}(q), b) > 0$. It follows that, by Equation (8.4) in Lemma 8.16 and since vcol' is n-prevailing compared to vcol_q, we have $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}'},q}[\mathsf{vcol}_q^{-1}[n] \setminus \{q\}] > 0$. Since $\mathsf{vcol}_q^{-1}[n] = X_n^i$ (by Equation 8.9), this is equivalent to: $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}'},q}[X_n^i] > 0$. That is, from q, there is a positive probability to reach a state $q' \in X_n^i$ for which, by $\mathcal{P}(i)$, we have: $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}'},q'}[H_i^* \cdot C_n] > 0$. It follows that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}'},q}[H_{i+1}^* \cdot C_n] > 0$ and $\mathcal{P}(i+1)$ holds. We can conclude that from all states in H there is a positive probability to reach C_n, that is $H \cap C_n \neq \emptyset$. Hence, n is the highest color in H and it is even. Thus, H is even-colored.

The case where n is odd is symmetrical, from Player-B's point-of-view.

□

The proof of Lemma 8.6 is now straightforward.

Proof. Indeed, by Lemma 8.19, we have $(\mathsf{vcol}, \mathsf{Ev}')$ non-deceiving down to $n-1$. Then, by Lemma 8.20, the pair of coloring and environment functions that Algorithm UpdNewSta (i.e. Algorithm 8.11) outputs is faithful down to $n-1$. □

[4] We have shown that $(\mathsf{vcol}, \mathsf{Ev})$ is coherent at n, hence all states in Q_n are such that $\mathsf{col}(q) \leq n$.

Proof of Lemma 8.7

Proof. Recall that $C := \mathsf{vcol}[Q]$, $k := \min C$ and assume that $k \leq e - 2$.

In fact, we consider $(\mathsf{vcol}', \mathsf{Ev}')$ to be the new environment and coloring functions obtained with Algorithm 8.13 before calling Algorithm 8.11. Let us show that $(\mathsf{vcol}', \mathsf{Ev}')$ is non-deceiving down to $n := k + 2$. First, note that vcol' is n-prevailing compared to vcol. Let $Q_n := (\mathsf{vcol}')^{-1}[n]$. For all $q \in Q_n$, we let vcol_q be the virtual coloring function associated with the environment $\mathsf{Ev}'(q)$.

2n-d. Let $q \in Q_n$. We have either $\mathsf{vcol}(q) = n - 2$ or $\mathsf{vcol}(q) = n$. In both cases, since $(\mathsf{vcol}, \mathsf{Ev})$ is completely faithful and therefore coherent at colors $n - 2$ and n, it holds that $\mathsf{col}(q) \leq n$, $\mathsf{vcol}_q(q) \leq n$, $n_q \leq n$ (with $n_q := \mathsf{NewCol}(q, \mathsf{vcol})$). Furthermore, vcol' is n-prevailing compared to vcol, which is n-prevailing compared to vcol_q. Finally, $n \equiv n-2 \mod 2$ and therefore this condition for being non-deceiving holds.

3n-d. This condition for being non-deceiving also straightforwardly holds. Indeed, assume towards a contradiction that a state $q \in Q_u$ such that $\mathsf{col}(q) = k = \mathsf{vcol}'(q) < n$ is such that $\mathsf{NewCol}(q, \mathsf{vcol}') = m > n$. We have vcol that is equivalent down to m to vcol'. This implies by Lemma 8.15, that we have $\mathsf{NewCol}(q, \mathsf{vcol}) = m > n$, which is not possible since $(\mathsf{vcol}, \mathsf{Ev})$ is assumed completely faithful.

1n-d. Consider now the first condition for being non-deceiving. It holds for all $k > n$, since $\mathsf{vcol}^{-1}[k] = (\mathsf{vcol}')^{-1}[k]$, and in both games $\mathcal{L}^k_{\mathsf{vcol}}$ and $\mathcal{L}^k_{\mathsf{vcol}'}$, the states in $\mathsf{vcol}^{-1}[[\![0, k-1]\!]] = \mathsf{vcol}'^{-1}[[\![0, k-1]\!]]$ are colored with $k-1$.

Assume that n is even. Consider a Player-A positional strategy $\mathsf{s_A}$ generated by the environment function Ev in the game $\mathcal{L}^n_{\mathsf{vcol}'}$. Let us show that this strategy parity dominates the valuation $v^u_{n,\mathsf{vcol}'}$ in the game $\mathcal{L}^n_{\mathsf{vcol}'}$. By Lemma 8.17, this strategy dominates the valuation $v^u_{n,\mathsf{vcol}'}$. Now, fix a Player-B positional deterministic strategy $\mathsf{s_B}$ and consider a BSCC H compatible with $\mathsf{s_A}$ and $\mathsf{s_B}$ in $\mathcal{L}^n_{\mathsf{vcol}'}$. If $K^{\geq n}$ occurs in H, by Lemma 8.18, H is even-colored. Assume now that $K^{\geq n}$ does not occur in H. Let $S_n := (\mathsf{vcol})^{-1}[n]$ (before running the algorithm, the set colored by n) and $T_n := Q_n \setminus S_n = \mathsf{vcol}^{-1}[n-2]$ (before running the algorithm, the set colored by $n-2$). Assume that $H \cap S_n = \emptyset$. In that case, this is a BSCC compatible with the strategy $\mathsf{s_A}$ in the game $\mathcal{L}^{n-2}_{\mathsf{vcol}}$ in which this strategy, by assumption, parity dominates the valuation $v^u_{n-2,\mathsf{vcol}}$. Furthermore, consider the state k^n_{n-1} in $\mathcal{L}^n_{\mathsf{vcol}}$ and the state k^{n-2}_{n-1} in $\mathcal{L}^{n-2}_{\mathsf{vcol}}$. These states are colored by $n - 1$ and exactly the same edges in both games $\mathcal{L}^n_{\mathsf{vcol}}$ and $\mathcal{L}^{n-2}_{\mathsf{vcol}}$ lead to these states. (These correspond to the states $q \in Q_u$ such that $\mathsf{vcol}(q) = n - 1$.) Hence, since the

BSCC H is even-colored in $\mathcal{L}^{n-2}_{\mathsf{vcol}}$, it also is in $\mathcal{L}^{n}_{\mathsf{vcol}}$[5]. Assume now that $H \cap T_n = \emptyset$. In that case, this is a BSCC compatible with the strategy $\mathsf{s_A}$ in the game $\mathcal{L}^{n}_{\mathsf{vcol}}$ in which this strategy, by assumption, parity dominates the valuation $v^u_{n,\mathsf{vcol}}$. Hence, this BSCC is even-colored. Finally, assume that $H \cap T_n \neq \emptyset$ and $H \cap S_n \neq \emptyset$. Let $C_n := Q_n \cap \mathsf{col}^{-1}[n]$. Let us show that for all $q \in (H \cap S_n) \setminus C_n$, there is a positive probability to visit C_n from q while only visiting states in $H \cap S_n$, that is:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}'},q}[(H \cap S_n)^* \cdot C_n] > 0$$

Assume towards a contradiction that this does not hold. That is, there is some $q \in (H \cap S_n) \setminus C_n$ such that $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}'},q}[(H \cap S_n)^* \cdot C_n] = 0$. Let $Z \subseteq H \cap S_n$ be the set of states in $H \cap S_n$ such that $Z := \{q\} \cup \{q' \in H \cap S_n \mid \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}},q}[(H \cap S_n)^* \cdot q'] > 0\}$. Note that $Z \cap C_n = \emptyset$. Our goal is now to exhibit a BSCC in $\mathcal{L}^{n}_{\mathsf{vcol}}$ that is compatible with $\mathsf{s_A}$ and odd-colored. This is a contradiction with the fact that $\mathsf{s_A}$ parity dominates the valuation $v^u_{n,\mathsf{vcol}}$ in the game $\mathcal{L}^{n}_{\mathsf{vcol}}$. Consider a Player-$\mathsf{B}$ strategy s'_B such that, for all $q \in S_n$, $\mathsf{s}'_\mathsf{B}(q) := \mathsf{s_B}(q)$ and for all $i \in [\![0, n-1]\!]$, we have $\mathsf{s_B}(k^n_i)$ playing in a deterministic way to reach q. By construction, we do have a BSCC H' compatible with $\mathsf{s_A}$ and s'_B in $\mathcal{L}^{n}_{\mathsf{vcol}}$ that is included in $Z \cup \{k^n_i \mid i \in [\![0, n-1]\!]\}$. Furthermore, since the BSCC H in the game $\mathcal{L}^{n}_{\mathsf{vcol}'}$ is not equal to $H \cap S_n$, it must be that at least one state in $\{k^n_i \mid i \in [\![0, n-1]\!]\}$ is in H'. Then, since each state in $\{k^n_i \mid i \in [\![0, n-1]\!]\}$ is colored with $n-1$ and no state in Z is in C_n, it follows that the highest color in H' is $n-1$, which is odd. That is, the BSCC H' is odd-colored. This is a contradiction. In fact, from all states $q \in (H \cap S_n) \setminus C_n$, we have $\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}^n_{\mathsf{vcol}'},q}[(H \cap S_n)^* \cdot C_n] > 0$. Furthermore, the colors of all states in H are at most n (recall what we shown above in 2n-d). Hence, the highest color in H is n, which is even (the color of the states in C_n). Therefore the BSCC H is even-colored. That is, the Player-A strategy $\mathsf{s_A}$ parity dominates the valuation $v^u_{n,\mathsf{vcol}'}$ in the game $\mathcal{L}^{n}_{\mathsf{vcol}'}$.

The case where n is odd is identical.

Hence, the pair $(\mathsf{vcol}', \mathsf{Ev}')$ is non-deceiving down to n. The result then follows from Lemma 8.20. □

Proof of Lemma 8.8

We can proceed to the proof of Lemma 8.8.

[5]Note that this is why we increase the least color, and not an intermediate color. Indeed, if there were a state colored by the coloring function vcol by $n-3$, then the state k^{n-2}_{n-3} would be colored by $n-3$ in $\mathcal{L}^{n-2}_{\mathsf{vcol}}$, whereas its counterpart k^n_{n-3} would be colored by $n-1$ in $\mathcal{L}^{n}_{\mathsf{vcol}}$. Hence, we could not deduce anymore that since the BSCC H is even-colored in $\mathcal{L}^{n-2}_{\mathsf{vcol}}$, it also is in $\mathcal{L}^{n}_{\mathsf{vcol}}$.

Proof. Assume towards a contradiction that $\min C = e-1$. Let $Q_{e-1} := \mathsf{vcol}^{-1}[e-1]$. Consider a Player-$\mathsf{B}$ strategy $\mathsf{s_B}$ generated by the environment function Ev. Because the pair $(\mathsf{vcol}, \mathsf{Ev})$ is completely faithful and $e-1$ is odd, it follows that $\mathsf{s_B}$ parity dominates the valuation $v^{u'}_{e-1,\mathsf{vcol}}$ in the game $\mathcal{L}^{e-1}_{\mathsf{vcol}}$ for some $u' \leq u$. Consider the game $\mathcal{G}^{Q_{e-1}}$ from Lemma 1.5.10 (we do not mention any player since the game has a value). This lemma also gives that all states have the same values in $\mathcal{G}$ and $\mathcal{G}^{Q_{e-1}}$. Let us show that there is some $x \in [0,1]$ such that $x < u$ and such that the strategy $\mathsf{s_B}$ parity dominates the valuation $v^x_{Q_{e-1}}$. First, note that for all $y \in [0,1]$, if $\mathsf{s_B}$ dominates the valuation $v^y_{Q_{e-1}}$, then it also parity dominates it (since $\mathsf{s_B}$ parity dominates the valuation $v^{u'}_{e-1,\mathsf{vcol}}$ in the game $\mathcal{L}^{e-1}_{\mathsf{vcol}}$). Now, consider some state $q \in Q_{e-1}$. Let vcol_q be the coloring function associated with the environment $\mathsf{Ev}(q)$. Denoting $Y_q := (\mathsf{F}(q), \mathsf{Ev}(q))$ and $p^q := p_{q,\mathsf{vcol}_q}$, we have that the Player-B GF-strategy $\mathsf{s_B}(q)$ is optimal w.r.t. Y_q. Hence, by Lemma 8.2 for Player B, we have that for all Player-A actions $a \in \mathsf{Act_A}$, if $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{(p^q)^{-1}[p^q_{[0,1]}]}\rangle](a, \mathsf{s_B}(q)) = 0$ then $\mathsf{Color}(\mathsf{F}(q), p^q, a, \mathsf{s_B}(q)) \cup \{e-2\}$ is odd with $\mathsf{Color}(\mathsf{F}(q), p^q, a, \mathsf{s_B}(q)) := \{i \in [\![0, e]\!] \mid \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{(p^q)^{-1}[k_i]}\rangle](a, \mathsf{s_B}(q)) > 0\}$. Since e is even, it must be that, for all states $q \in Q_{e-1}$ and $a \in \mathsf{Act}^q_{\mathsf{A}}$:

$$\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{(p^q)^{-1}[p^q_{[0,1]}]}\rangle](a, \mathsf{s_B}(q)) = 0 \Rightarrow \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{(p^q)^{-1}[k_e]}\rangle](a, \mathsf{s_B}(q)) = 0 \tag{8.11}$$

For $q \in Q_{e-1}$, let $\mathsf{NZ}(q) := \{a \in \mathsf{Act}^q_{\mathsf{A}} \mid \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{(p^q)^{-1}[p^q_{[0,1]}]}\rangle](a, \mathsf{s_B}(q)) > 0\}$. Then, we let:

$$p_m := \min_{q \in Q_{e-1}} \min_{a \in \mathsf{NZ}(q)} \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{(p^q)^{-1}[p^q_{[0,1]}]}\rangle](a, \mathsf{s_B}(q)) > 0 \tag{8.12}$$

and

$$x := u + p_m \cdot (u' - u) < u \tag{8.13}$$

Let us show that the strategy $\mathsf{s_B}$ dominates the valuation $v^x_{Q_{e-1}}$ in the game $\mathcal{G}^{Q_{e-1}}$. Let $q \in Q_{e-1}$. We have $\mathsf{val}[\langle \mathsf{F}(q), v\rangle](\mathsf{s_B}(q)) = \max_{a \in \mathsf{Act}^q_{\mathsf{A}}} \mathsf{out}[\langle \mathsf{F}(q), v\rangle](a, \mathsf{s_B}(q))$. Consider some $a \in \mathsf{Act}^q_{\mathsf{A}}$ and a Player-A strategy $\mathsf{s_A}$ with $\mathsf{s_A}(q) := a$. Denoting $P_a[T] := \mathbb{P}^{\mathcal{C}^{e-1}_{\mathsf{vcol}},q}_{\mathsf{s_A},\mathsf{s_B}}[T]$ for $T \subseteq Q_{e-1} \cup \{k^{e-1}_e\} \cup V_{Q \setminus Q_u}$, we have:

$$\begin{aligned}
\mathsf{out}[\langle \mathsf{F}(q), v^x_{Q_{e-1}}\rangle](a, \mathsf{s_B}(q)) &= x \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q_{e-1}}\rangle](a, \mathsf{s_B}(q)) \\
&\quad + u \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q_e}\rangle](a, \mathsf{s_B}(q)) \\
&\quad + \sum_{y \in V_{Q \setminus Q_u}} y \cdot \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{Q_y}\rangle](a, \mathsf{s_B}(q)) \\
&= x \cdot P_a[Q_{e-1}] + u \cdot P_a[\{k^{e-1}_e\}] + \sum_{y \in V_{Q \setminus Q_u}} y \cdot P_a[y]
\end{aligned}$$

In addition, since the strategy $\mathsf{s_B}$ dominates the valuation $v^{u'}_{e-1,\mathsf{vcol}}$ in the game $\mathcal{L}^{e-1}_{\mathsf{vcol}}$, we have:

$$u' \geq \mathsf{out}[\langle \mathsf{F}(q), v^{u'}_{e-1,\mathsf{vcol}} \rangle](a, \mathsf{s_B}(q)) \\ = u' \cdot P_a[Q_{e-1} \cup \{k_e^{e-1}\}] + \sum_{y \in V_{Q \setminus Q_u}} y \cdot P_a[y]$$

That is:

$$\sum_{y \in V_{Q \setminus Q_u}} y \cdot P_a[y] \leq u' \cdot P_a[V_{Q \setminus Q_u}]$$

By Equation (8.1) in Lemma 8.16, for all $y \in V_{Q \setminus Q_u}$, we have:

$$P_a[y] = \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{(p^q)^{-1}[y]} \rangle](a, \mathsf{s_B}(q))$$

Furthermore, by Equation (8.3) in Lemma 8.16 which can be applied since $(\mathsf{vcol}, \mathsf{Ev})$ is faithful down to $e-1$, it also is coherent at $e-1$, we have:

$$P_a(\{k^e_{e-1}\}) = \mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{(p^q)^{-1}[k_e]} \rangle](a, \mathsf{s_B}(q))$$

Now, assume that $\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{(p^q)^{-1}[V_{Q \setminus Q_u}]} \rangle](a, \mathsf{s_B}(q)) = 0$. That is, for all $y \in V_{Q \setminus Q_u}$, we have $P_a(y) = 0$. By Equation (8.11) it follows that $P_a[\{k^e_{e-1}\}] = 0$. Hence: $v^x_{Q_{e-1}}(q) = x = \mathsf{out}[\langle \mathsf{F}(q), v^x_{Q_{e-1}} \rangle](a, \mathsf{s_B}(q))$.

Assume now that $\mathsf{out}[\langle \mathsf{F}(q), \mathbb{1}_{(p^q)^{-1}[V_{Q \setminus Q_u}]} \rangle](a, \mathsf{s_B}(q)) > 0$ (i.e. $a \in \mathsf{NZ}(q)$). By Equation (8.12), it implies that $P_a[V_{Q \setminus Q_u}] \geq p_m$. It follows that, by Equation (8.13):

$$\begin{aligned} \mathsf{out}[\langle \mathsf{F}(q), v^x_{Q_{e-1}} \rangle](a, \mathsf{s_B}(q)) &= x \cdot P_a[Q_{e-1}] + u \cdot P_a[\{k^e_{e-1}\}] + \sum_{y \in V_{Q \setminus Q_u}} y \cdot P_a[y] \\ &\leq u \cdot P_a[Q_{e-1}] + u \cdot P_a[\{k^e_{e-1}\}] + u' \cdot P_a[V_{Q \setminus Q_u}] \\ &= u \cdot (1 - P_a[V_{Q \setminus Q_u}]) + u' \cdot P_a[V_{Q \setminus Q_u}] \\ &= u + P_a[V_{Q \setminus Q_u}] \cdot (u' - u) \leq u + p_m \cdot (u' - u) \\ &= x = v^x_{Q_{e-1}}(q) \end{aligned}$$

As this holds for all $x \in Q_{e-1}$, it follows that the Player-B strategy $\mathsf{s_B}$ dominates the valuation $v^x_{Q_{e-1}}$, and therefore parity dominates it. Hence, for all $q \in Q_{e-1}$, we have $\chi_{\mathcal{G}^{Q_{e-1}}}[q] \leq x < u$. That is a contradiction with Lemma 1.5.10. □

8.6.6 . Proof of Lemma 8.9

First, let us formally define a way to compare coloring functions.

Definition 8.24. *We define the (transitive) relation $\prec_{[\![0,e]\!]}$ on coloring functions which corresponds to the lexicographic order. Specifically, for all virtual coloring functions $\mathsf{vcol}_1, \mathsf{vcol}_2 : Q_u \to [\![0,e]\!]$, we have $\mathsf{vcol}_1 \prec_{[\![0,e]\!]} \mathsf{vcol}_2$ if and only if there is some $k \in [\![0,e]\!]$ such that:*

- *for all* $i \in [\![k+1, e]\!]$*, we have* $|\mathsf{vcol}_1^{-1}[i]| = |\mathsf{vcol}_2^{-1}[i]|$*;*

- $|\mathsf{vcol}_1^{-1}[k]| < |\mathsf{vcol}_1^{-1}[k]|$.

We write $\mathsf{vcol}_1 =_{[\![0,e]\!]} \mathsf{vcol}_2$ *when such a property is ensured for* $k = -1$*. We also write* $\mathsf{vcol}_1 \preceq_{[\![0,e]\!]} \mathsf{vcol}_2$ *for* $\mathsf{vcol}_1 =_{[\![0,e]\!]} \mathsf{vcol}_2$ *or* $\mathsf{vcol}_1 \prec_{[\![0,e]\!]} \mathsf{vcol}_2$*.*

Straightforwardly, this relation ensures the following proposition.

Proposition 8.21. *Let* $n \in \mathbb{N}$*. There is no infinite sequence* $(c_k)_{k \in \mathbb{N}} \in ([\![0,e]\!]^n)^{\mathbb{N}}$ *such that* $c_k \prec_{[\![0,e]\!]} c_{k+1}$ *for all* $k \in \mathbb{N}$*.*

Proof. This is because the relation $\prec_{[\![0,e]\!]}$ corresponds to a lexicographic order on vectors taking their values in $[\![0,e]\!]$. □

We can then consider all steps of Algorithm 8.14 one by one.

Lemma 8.22. *Consider a virtual coloring function* $\mathsf{vcol} : Q_u \to [\![0,e]\!]$ *and some* $k \in [\![0,e]\!]$*. For all partial environment functions* $\mathsf{Ev} : Q_u \to \mathsf{Env}(Q)$ *defined on* $[\![k+1,e]\!]$*:*

- *for* $(\mathsf{vcol_{upd}}, \mathsf{Ev_{upd}}) \leftarrow \mathsf{UpdateColEnv}(k, \mathsf{vcol}, \mathsf{Ev})$*, we have* $\mathsf{vcol} \preceq_{[\![0,e]\!]} \mathsf{vcol_{upd}}$*;*

For all environment functions $\mathsf{Ev} : Q_u \to \mathsf{Env}(Q)$*:*

- *for* $(\mathsf{vcol_{Inc}}, \mathsf{Ev_{Inc}}) \leftarrow \mathsf{IncLeast}(\mathsf{vcol}, \mathsf{Ev})$*, we have* $\mathsf{vcol} \prec_{[\![0,e]\!]} \mathsf{vcol_{Inc}}$*;*

Proof. The procedure UpdateColEnv consists in two part. First, the call Algorithm 8.10. This does not change the coloring function. Then, there is the call to Algorithm 8.11. By Proposition 8.13, if $\mathsf{NewCol}(q, \mathsf{vcol}) = k$, then $\mathsf{vcol}(q) \leq k$. Hence, the change of colors of the states does not decrease the colors of the states. It follows that the resulting coloring function $\mathsf{vcol_{upd}}$ is such that $\mathsf{vcol} \preceq_{[\![0,e]\!]} \mathsf{vcol_{upd}}$.

Consider now the procedure IncLeast. By definition, before the call of this algorithm, some states are colored with $c_{\min}$. Hence, the coloring function resulting of the call to Algorithm 8.13 has more states colored with $c_{\min} + 2$, whereas states colored with a color higher than $c_{\min} + 2$ are left unchanged. Hence: $\mathsf{vcol} \prec_{[\![0,e]\!]} \mathsf{vcol_{Inc}}$. □

We can now proceed to the proof of Lemma 8.9.

Proof. By Lemma 8.22, each call to the procedure IncLeast strictly increase (w.r.t. $\prec_{[\![0,e]\!]}$) the coloring function. Furthermore, each call to the procedure UpdateColEnv does not decrease (w.r.t. $\prec_{[\![0,e]\!]}$) the coloring function. Hence, by Proposition 8.21, the call to the procedure IncLeast is done only finitely many times. □

9 - Study of standard finite game forms

In this part of the dissertation, we have defined several classes of game forms, such as the class of determined (Definition 6.1), finitely/uniquely maximizable (Definitions 6.3, 6.4), or positionally optimizable (Definition 8.7) game forms. Each of these classes is associated with transfers (though, not necessarily NSC-transfers). See Chapters 6, 7 and 8 for more details.

In this final chapter, we focus on standard finite game forms. The classes of game forms we study are the above mentioned classes of game forms restricted to standard finite game forms. We will state two kinds of results. First, we state decidability/complexity results on class membership problems: given a game form $\mathcal{F}$, if $\mathcal{F}$ belongs to one of the classes mentioned above. The standard finite assumption facilitates the reasoning about decidability/complexity issues. When considering this issue, for the same reason, we also only consider deterministic game forms. Second, we state expressiveness results, that is we compare the different classes of game forms defined in this part, i.e. we establish which class subsumes which.

More specifically, in Section 9.1, we study the complexity of deciding if a standard finite deterministic game form is determined. It is straightforward that this decision problem is in co-NP. We show that it is equivalent, under polynomial time reduction, to the decision problem MonotoneDual (the dualization of monotone CNF formulas), see Proposition 9.4. It is an open problem whether this decision problem MonotoneDual is in P or is co-NP-complete.

Then, in Section 9.2, we encode the decision problems w.r.t. some classes mentioned above with formulas of the first order theory of the reals. These formulas consist in (existential or universal) quantifications over (real) variables followed by a combination, using the classical logical operators such as and, or, etc., of inequalities between multi-variate polynomials. This is formally defined in Definition 9.5. Interestingly for us, the first order theory of the reals is decidable. Using this fact, we are then able to show that it is decidable (resp. semi-decidable) if a game form is uniquely (resp. finitely) maximizable w.r.t. Player A, see Corollary 9.7. Furthermore, we show that the assertion that a game form is positionally optimizable (resp. up to some $n \in \mathbb{N}$) can be encoded in the first order theory of the reals, see Proposition 9.9, thus showing that the corresponding decision problem is decidable.

Finally, in Section 9.3, Theorem 9.10 gives the complete picture of how the sets of determined, finitely/uniquely maximizable w.r.t. Player A and positionally optimizable game forms compare. The results are summarized in Figure 9.3. The main difficulty lies in establishing a strict hierarchy: for all $n \in \mathbb{N}$, the set of game forms positionally optimizable up to $n + 1$ is strictly included in the set of game forms positionally optimizable up to n. This is done in Proposition 9.12.

9.1 Determined game forms

In this section, we study the complexity of deciding if a standard finite deterministic game form is determined (recall Definition 6.1). This section comes from [38]. We formally define below the decision problem we will study in this section, along with the definition of the size of standard finite deterministic game forms.

Definition 9.1. *The decision problem* DetGF *is as follows:*

- *Input: a standard finite deterministic game form $\mathcal{F}$;*
- *Output: yes if and only if the game form $\mathcal{F}$ is determined.*

Such a game form $\mathcal{F} = \langle \mathsf{Act_A}, \mathsf{Act_B}, \mathsf{O}, \varrho \rangle_\mathsf{s}$ is represented as a bi-dimensional table where O is inferred from the contents of the cells of the table (i.e. $\mathsf{O} = \varrho(\mathsf{Act_A}, \mathsf{Act_B})$). We assume that $\mathsf{Act_A} = [\![1, |\mathsf{Act_A}|]\!]$ and similarly for $\mathsf{Act_B}$ (where $|\mathsf{Act_A}|$ refers to the cardinal of $\mathsf{Act_A}$).

The size $|\mathcal{F}|$ of such a game form $\mathcal{F}$ is equal to $|\mathcal{F}| := |\mathsf{Act_A}| \times |\mathsf{Act_B}|$.

It is straightforward that this decision problem is in co-NP. Indeed, if a game form is not determined, it suffices to guess a valuation $v : \mathsf{O} \to \{0, 1\}$ and check — in polynomial time — that there is no row full of 1 nor any column full of 0. In fact, in [71] (where determinacy is refered to as tightness), the authors mentioned that DetGF could be solved in quasi-polynomial time via a reduction to the dualization of monotone CNF formulas (called MonotoneDual), which can be solved in quasi-polynomial time [75]. Note that it is an open problem whether MonotoneDual is in P or is coNP-complete [76].

The goal of this section is to show that DetGF is equivalent, under polynomial time reduction, to MonotoneDual, thus showing that answering if DetGF is in P or coNP-complete directly answers the same question for MonotoneDual. To do so, we will actually show that DetGF is equivalent, under polynomial time reduction, to the decision problem co-IMSAT, which is equivalent, under polynomial time reduction, to MonotoneDual [77, Corollary 1, Theorem 2].

Let us define this decision problem IMSAT. In the classical decision problem SAT, a logical formula given in conjunctive normal form is given as input, the output being yes if and only if this formula is satisfiable. The problem IMSAT refers to a version of SAT where the formula taken as input is monotone — i.e. in each clause, either all variables are positive (not negated) or all variables are negative (negated) — and intersecting — each pair of positive and negative clauses has at least one variable in common. Positive and negative clauses are formally defined below.

Definition 9.2 (Positive and negative clauses). *Consider a non-empty set of variables $X \neq \emptyset$. A* positive *clause C^+ on X represented by a subset $X_{C^+} \subseteq X$ of X is equal to:*

$$C^+ := \bigvee_{x \in X_{C^+}} x$$

The set of all positive clauses on the set X of variables is denoted $\mathsf{PosClause}(X)$*. Similarly, a* negative *clause C^- on X represented by a subset $X_{C^-} \subseteq X$ of X is equal to:*

$$C^- := \bigvee_{x \in X_{C^-}} \neg x$$

The set of all negative clauses on the set X of variables is denoted $\mathsf{NegClause}(X)$*.*

We define below the notion of intersecting monotone formula.

Definition 9.3 (Intersecting Monotone formula). *Consider a non-empty set of variables $X \neq \emptyset$. An* intersecting monotone *formula φ (*IM*-formula for short) on the set of variables X is a CNF formula:*

$$\varphi := \bigwedge_{1 \leq i \leq n} C_i^+ \wedge \bigwedge_{1 \leq j \leq k} C_j^-$$

where, for all $i \in [\![1, n]\!]$, we have $C_i^+ \in \mathsf{PosClause}(X)$ a positive clause and for all $j \in [\![1, k]\!]$, we have $C_j^- \in \mathsf{NegClause}(X)$ a negative clause, with $k, n \geq 1$. Furthermore, for all $i \in [\![1, n]\!]$ and $j \in [\![1, k]\!]$, we have $X_{C_i^+} \cap X_{C_j^-} \neq \emptyset$.

This induces the decision problem co-IMSAT.

Definition 9.4. *The decision problem* co-IMSAT *is as follows:*

- *Input: a non-empty set of variables X and an intersecting monotone formula φ and X;*
- *Output: yes if and only if no valuation of the variables in X satisfies φ.*

As mentioned above, co-IMSAT is equivalent to MonotoneDual.

Theorem 9.1. *The decision problem* MonotoneDual *is equivalent, under polynomial time reductions, to the decision problem* co-IMSAT*.*

The remainder of this section is devoted to the proof that DetGF is equivalent, under polynomial time reduction, to co-IMSAT, which in turns, shows that DetGF is equivalent, under polynomial time reduction, to MonotoneDual. In the proof of the following two lemmas, $\{0, 1\}$-valuations of the outcomes will be seen as $\{\mathsf{False}, \mathsf{True}\}$-valuations where False corresponds to 0 and True corresponds to 1.

Lemma 9.2. *The decision problem* co-IMSAT *is at least as hard, under polynomial time reduction, as* DetGF*.*

$$\mathcal{F}_1 = \begin{bmatrix} x & x & z \\ x & y & y \\ z & y & z \end{bmatrix}$$

Figure 9.1: A determined game form (already depicted in Figure 6.1).

Proof. We exhibit a polynomial time reduction from DetGF to $\mathsf{co\text{-}IMSAT}$. Specifically, consider a standard game form $\mathcal{F} = \langle \mathsf{St_A}, \mathsf{St_B}, \mathsf{O}, \varrho \rangle_{\mathsf{s}}$ and assume that $\mathsf{Act_A} = [\![1, k]\!]$ and $\mathsf{Act_B} = [\![1, n]\!]$ for some $n, k \geq 1$. We consider the set of variables $X = \mathsf{O}$ and the IM-formula $\varphi_{\mathcal{F}}$, defined by:

$$\varphi_{\mathcal{F}} := \bigwedge_{1 \leq i \leq n} C_i^+ \wedge \bigwedge_{1 \leq j \leq k} C_j^-$$

where, for all $i \in \mathsf{Act_B} = [\![1, n]\!]$, we have $C_i^+ \in \mathsf{PosClause}(X)$ and $X_{C_i^+} := \varrho(\mathsf{Act_A}, i) \subseteq \mathsf{O} = X$ and for all $j \in [\![1, k]\!]$, we have $C_j^- \in \mathsf{NegClause}(X)$ and $X_{C_j^-} := \varrho(j, \mathsf{Act_B}) \subseteq \mathsf{O} = X$. That is, the positive clauses encode the columns whereas the negative clauses encode the rows.

Note that the formula φ is indeed an IM-formula as it is monotone by definition and, for all $i \in [\![1, n]\!]$ and $j \in [\![1, k]\!]$, we have $\varrho(j, i) \in X_{C_i^+} \cap X_{C_j^-} \neq \emptyset$. It is also clear that this transformation is computable in polynomial time. As an example, the formula $\varphi_{\mathcal{F}_1}$ corresponding to the game form $\mathcal{F}_1$ of Figure 9.1 is equal to:

$$\varphi_{\mathcal{F}} = (x \vee z) \wedge (x \vee y) \wedge (y \vee z) \wedge (\neg x \vee \neg z) \wedge (\neg x \vee \neg y) \wedge (\neg y \vee \neg z)$$

In that case, the formula $\varphi_{\mathcal{F}_1}$ is not satisfiable and the game form $\mathcal{F}_1$ is determined.

We now have to show that the formula $\varphi_{\mathcal{F}}$ is not satisfiable if and only if the game form $\mathcal{F}$ is determined. First, assume that $\varphi_{\mathcal{F}}$ is not satisfiable and let $v : \mathsf{O} \to \{0, 1\}$ be a $\{0, 1\}$-valuation of the outcomes O. Note that it can be seen as a $\{\mathsf{False}, \mathsf{True}\}$-valuation v of the variables in $X = \mathsf{O}$. Since $\varphi_{\mathcal{F}}$ is not satisfiable, there exists a clause C that is not satisfied by the valuation v. Assume for instance that it is a positive clause $C = C_i^+$ for some $i \in [\![1, n]\!] = \mathsf{Act_B}$. Then, for all $x \in X_{C_i^+}$, we have $v(x) = 0$. That is, $\{0\} = v[X_{C_i^+}] = v[\varrho(\mathsf{Act_A}, i)]$. Similarly, if C is a negative clause $C = C_j^-$ for some $j \in [\![1, k]\!] = \mathsf{Act_A}$, then, for all $x \in X_{C_j^-}$, we have $v(x) = 1$, i.e. $\{1\} = v[X_{C_j^-}] = v[\varrho(j, \mathsf{Act_B})]$. As this holds for all $\{0, 1\}$-valuations, the game form $\mathcal{F}$ is determined.

It is analogous to show that if $\mathcal{F}$ is determined then $\varphi_{\mathcal{F}}$ is not satisfiable: for any $\{\mathsf{False}, \mathsf{True}\}$-valuation of the outcomes $v : \mathsf{O} \to \{\mathsf{False}, \mathsf{True}\}$, which

$$\mathcal{F}_\varphi = \begin{bmatrix} b & a & a & x_1 & x_1 \\ d & a & e & x_1 & x_1 \\ b & c & x & x_0 & x_1 \\ x_0 & x_0 & x_1 & x_1 & x_1 \\ x_0 & x_0 & x_0 & x_0 & x_0 \end{bmatrix}$$

Figure 9.2: The game form that is the translation of the formula φ.

can be seen as a $\{0, 1\}$-valuation, if there is a row full of 1 in $\langle \mathcal{F}, v \rangle$, then the corresponding negative clause is unsatisfied by v, and similarly if there is a column full of 0 in $\langle \mathcal{F}, v \rangle$, then the corresponding positive clause is unsatisfied by v. This proves the desired result. □

Let us now show the other direction, i.e. that the decision problem DetGF is at least as hard as the decision problem co-IMSAT. This is stated in the lemma below.

Lemma 9.3. *The decision problem* DetGF *is at least as hard under polynomial time reduction as* co-IMSAT.

Proof. We define a polynomial time reduction from co-IMSAT to DetGF. Specifically, consider an IM-formula φ on a non-empty set of variables X such that:

$$\varphi := \bigwedge_{1 \leq i \leq n} C_i^+ \wedge \bigwedge_{1 \leq j \leq k} C_j^-$$

with $k, n \geq 1$. We want to build a game form $\mathcal{F}_\varphi$ that is determined if and only if φ is not satisfiable. The idea of the reduction is close to what we have done in the proof of the previous lemma: the positive clauses are encoded in the columns and the negative clauses are encoded in the rows, the intersection of a row and a column being well defined since the formula φ is intersecting. However, some technical difficulties arise from the fact that if, for instance, a negative clause has more variables than the number of positive clauses, then there will be some outcomes in $\mathcal{F}_\varphi$ in a row corresponding to a negative clause but whose column does not correspond to any positive clause. Hence, this column should be made so that it cannot be full of 0s (for the relevant valuations) as to not affect the determinacy of $\mathcal{F}_\varphi$. Let us illustrate this on an example. Consider the following IM-formula:

$$\varphi := (b \vee d) \wedge (a \vee c) \wedge (\neg a \vee \neg b) \wedge (\neg a \vee \neg d \vee \neg e)$$

Figure 9.2 depicts the game form $\mathcal{F}_\varphi$ we build from φ (the variables x, x_0, x_1 are fresh, the variable x is only added to fill the game form). The negative clauses are encoded in the red rows and the positive clauses are encoded in the blue columns. The variables x_0 and x_1 are added so that only the red rows and the blue lines are of interest for the determinacy of the game form $\mathcal{F}_\varphi$. More precisely, it is straightforward that for any $\{0,1\}$-valuation $v : \mathsf{O} \to \{0,1\}$ with either $v(x_0) = 1$ or $v(x_1) = 0$, there is either a row full of 1 or a column full of 0. Furthermore, for any $\{0,1\}$-valuation $v : \mathsf{O} \to \{0,1\}$ such that $v(x_0) = 0$ and $v(x_1) = 1$, there is a row full of 1 or a column full of 0 if and only if a red row is full of 1 or a blue column is full of 0.

Formally, let $m_k := \max_{j \in [\![1,k]\!]} |X_{C_j^-}|$ and $m_n := \max_{i \in [\![1,n]\!]} |X_{C_i^+}|$ be the maximum number of variables occurring, respectively, in a positive and a negative clause. We define the game form $\mathcal{F} := \langle \mathsf{Act_A}, \mathsf{Act_B}, \mathsf{O}, \varrho \rangle$ by:

- $\mathsf{Act_A} := [\![1, k + m_k + 2]\!]$;
- $\mathsf{Act_B} := [\![1, n + m_n + 2]\!]$;
- $\mathsf{O} := X \cup \{x, x_0, x_1\}$ for three fresh outcomes $x, x_0, x_1 \notin X$;
- Let $j \leq k$. Let us define the row corresponding to j. For all $i \leq n$, we set $\varrho(j,i) \in \mathsf{O}$ such that $\varrho(j,i) \in C_j^- \cap C_i^+$. Then, $\varrho(j,i)$ for $n+1 \leq i \leq n + m_n$ is defined such that $\varrho(j, [\![1, n+m_n]\!]) = X_{C_j^-}$ (note that this is possible since $|X_{C_j^-}| \leq m_n$). Furthermore, $\varrho(j, n + m_n + 1) := x_1$ and $\varrho(j, n + m_n + 2) := x_1$.

  Consider the column corresponding to $i \in [\![1,n]\!]$, its value for $j \leq k$ is already defined. Then, $\varrho(j,i)$ for $k+1 \leq j \leq k+m_k$ is defined such that $\varrho([\![1, k+m_k]\!], i) = X_{C_i^+}$ (note that this is possible since $|X_{C_i^+}| \leq m_k$). Furthermore, $\varrho(i, k + m_k + 1) := x_0$ and $\varrho(i, k + m_k + 2) := x_0$.

  Furthermore, for all $n+1 \leq i \leq n+m_n$ and $k+1 \leq j \leq k+m_k$, we set $\varrho(j,i) := x$. In addition, $\varrho(k+m_k+1, [\![n+1, n+m_n+1]\!]) := \{x_1\}$, $\varrho([\![k+1, k+m_k]\!], n+m_n+1) := \{x_0\}$. Finally, $\varrho(k+m_k+2, \mathsf{Act_B}) := \{x_0\}$ and $\varrho(\mathsf{Act_A} \setminus \{k+m_k+2\}, n+m_n+2) := \{x_1\}$.

First, note that this reduction can indeed be computed in polynomial time. Consider now a $\{0,1\}$-valuation $v : \mathsf{O} \to \{0,1\}$ and the game in normal form $\langle \mathcal{F}, v \rangle$. We have the following:

- Assume that $v(x_0) = 1$ or $v(x_1) = 0$. Then, we have $v[\varrho(k + m_k + 2, \mathsf{Act_B})] = \{v(x_0)\}$ and $v[\varrho(\mathsf{Act_A}, n + m_n + 2)] := \{v(x_0), v(x_1)\}$. That is, if $v(x_0) = 1$ there is a row full of 1, and if $v(x_0) = v(x_1) = 0$, then there is a column full of 0.
- Assume now that $v(x_0) = 0$ and $v(x_1) = 1$. Let $j \in \mathsf{Act_A} \setminus [\![1,k]\!]$. If $j \leq k + m_k$, then $\varrho(j, n+m_n+1) = x_0$, hence the row j is not full of 1 (w.r.t. the valuation v). Similarly, if $j = k + m_k + 1$ or $j = k + m_k + 2$, then $\varrho(j, 1) = x_0$ and the row j is also not full of 1. Furthermore, for all

$i \in \mathsf{Act_B} \setminus [\![1, n]\!]$, we have $\varrho(k + m_k + 1, i) = x_1$. Hence, the column i is not full of 1.

Furthermore, for all $a \in [\![1, k]\!]$, $v[\varrho(a, \mathsf{Act_B} \setminus [\![1, n+m_n]\!])] = \{v(x_1)\} = \{1\}$ and for all $b \in [\![1, n]\!]$, $v[\varrho(\mathsf{Act_A} \setminus [\![1, k + m_k]\!], b)] = \{v(x_0)\} = \{0\}$. Therefore, there is a row full of 1 or a column full of 0 in $\langle \mathcal{F}, v \rangle$ if and only if there is $a \in [\![1, k]\!] \subseteq \mathsf{Act_A}$ such that $v[\varrho(a, [\![1, n + m_n]\!])] = v[X_{C_a^-}] = \{1\}$ or $b \in [\![1, n]\!] \subseteq \mathsf{Act_B}$ such that $v[\varrho([\![1, k + m_k]\!], b)] = v[X_{C_b^+}] = \{0\}$.

Proving that $\varphi_{\mathcal{F}}$ is not satisfiable if and only if $\mathcal{F}_\varphi$ is determined is now direct. Indeed, if φ is not satisfiable, then for all valuations $v : \mathsf{O} \to [0, 1]$, there is either: some $a \in [\![1, k]\!]$ such that the negative clause C_a^- is not satisfied, i.e. $v[X_{C_a^-}] = \{1\}$, and in that case the row a is full of 1 in $\langle \mathcal{F}, v \rangle$; or some $b \in [\![1, n]\!]$ such that the positive clause C_b^+ is not satisfied, i.e. $v[X_{C_b^+}] = \{0\}$, and in that case the column b is full of 0 in $\langle \mathcal{F}, v \rangle$. The other direction is very similar. □

We obtain the proposition below.

Proposition 9.4. *The decision problem* DetGF *is equivalent under polynomial time reductions to the decision problem* MonotoneDual.

Proof. This comes from Lemmas 9.2 and 9.3 and Theorem 9.1. □

9.2 First-order theory of the reals

In this section, we use the first order theory of the reals to show that it is possible to decide if a game form is (finitely or uniquely) maximizable or if it is positionally maximizable. We restrict ourselves to standard finite deterministic game forms, though we use the deterministic assumptions only to simplify the proofs.

9.2.1 . Definition

Let us first formally define the first order theory of the reals formulas we will consider.

Definition 9.5 (First-order theory of the reals)**.** *In the first order theory of the reals (*FO-$\mathbb{R}$ *for short), we consider formulas* Φ *of the shape:*

$$\Phi = Q_1\ x_1 \in \mathbb{R}, \ldots, Q_n\ x_n \in \mathbb{R}, \varphi(x_1, \ldots, x_n)$$

where $n \in \mathbb{N}$ *and*

- *for all* $i \in [\![1, n]\!]$, $Q_i \in \{\exists, \forall\}$. *That is, it is either an existential or a universal quantifier;*

- *$\varphi(x_1, \ldots, x_n)$ is a classical formula without quantifiers, with $\wedge$ (and), $\vee$ (or) and $\neg$ (negation) as logical connectors. The atomic propositions considered in $\varphi(x_1, \ldots, x_n)$ are of the shape:*

$$P(x_1, \ldots, x_n) \bowtie 0$$

where $\bowtie \in \{=, \neq, \geq, >, \leq, <\}$ and $P : \mathbb{R}^n \to \mathbb{R}$ is a real multi-variate polynomial with integer coefficients.

When all quantifiers Q_i for $i \in [\![1, n]\!]$ are existential quantifiers, the formula Φ belongs to the existential theory of the reals ($\exists$-$\mathbb{R}$ for short).

The semantics behind the first order theory of the reals is quite intuitive, though it would take some space to formally define. Instead, let us illustrate it on a couple of examples.

Example 9.1. *Consider the* FO-$\mathbb{R}$ *formula below.*

$$\forall x \in \mathbb{R},\ \exists y \in \mathbb{R},\ x + y \geq 0$$

This FO-$\mathbb{R}$ *formula is true since it does hold that, for all $x \in \mathbb{R}$, taking $y := -x \in \mathbb{R}$, we have $x + y = 0$. However, the* FO-$\mathbb{R}$ *formula*

$$\exists x \in \mathbb{R},\ \forall y \in \mathbb{R},\ x + y \geq 0$$

is false in FO-$\mathbb{R}$ *since there is no $x \in \mathbb{R}$, such that, for all $y \in \mathbb{R}$, we have $x + y \geq 0$.*

Let us now consider the decision problem associated with FO-$\mathbb{R}$ formulas.

Definition 9.6. *We denote by $\mathsf{True}_{\mathsf{FO}-\mathbb{R}}$ the set of all first-order theory of the reals formulas that evaluate to true. The decision problem $\mathsf{Is}_{\mathsf{FO}-\mathbb{R}}$ is as follows:*

- *Input: an* FO-$\mathbb{R}$ *formula Φ;*

- *Output: yes if and only if the formula Φ is in $\mathsf{True}_{\mathsf{FO}-\mathbb{R}}$.*

Theorem 9.5. *The decision problem $\mathsf{Is}_{\mathsf{Fo}-\mathbb{R}}$ is decidable.*

This result was first shown in [78]. It was then improved in [79, Theorem 1.1] it as it was shown that $\mathsf{Is}_{\mathsf{FO}-\mathbb{R}}$ could be decided in doubly-exponential time via a quantifier elimination procedure.

As a side remark, in [80], it was shown that the decision problem $\mathsf{Is}_{\mathsf{Fo}-\mathbb{R}}$ for those formulas belonging to the existential theory of the reals can be decided in polynomial space.

9.2.2 . Finitely and uniquely maximizable game forms

Let us first focus on game forms finitely and uniquely maximizable w.r.t. Player A (recall Definitions 6.3 and 6.4). Given some $n \in \mathbb{N}$, we encode in FO-$\mathbb{R}$ the fact that a standard finite deterministic game form is maximizable w.r.t. Player A by a set of cardinal n. We formally define below this notion.

Definition 9.7 (n-maximizable game forms). *Consider a game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ on a set of outcomes O. For all $n \in \mathbb{N}$, we say that the game form $\mathcal{F}$ is n-*maximizable w.r.t. Player A *if there is a set of* GF*-strategies $S_A \subseteq \Sigma_{\mathsf{A}}$ that maximizes the game form $\mathcal{F}$ such that $|S_{\mathsf{A}}| \leq n$. We denote by n-$\mathsf{Max}_{\mathsf{A}}$ the set of game forms n-maximizable w.r.t. Player* A.

By definition, a game form is uniquely maximizable if and only if it is in 1-$\mathsf{Max}_{\mathsf{A}}$.

The fact that standard finite deterministic game form is in 1-$\mathsf{Max}_{\mathsf{A}}$ can be encoded in the first-order theory of the reals.

Proposition 9.6. *Consider a standard finite deterministic game form $\mathcal{F}$. For all $l \in \mathbb{N}$, the fact that $\mathcal{F} \in l$-$\mathsf{Max}_{\mathsf{A}}$ can be encoded, with a formula whose size is polynomial in l and $|\mathcal{F}|$ (as defined in Definition 9.1), in* FO-$\mathbb{R}$.

Proof. Consider some $l \in \mathbb{N}$. We encode the fact that a standard finite deterministic game form $\mathcal{F} = \langle \mathsf{Act}_{\mathsf{A}}, \mathsf{Act}_{\mathsf{B}}, \mathsf{O}, \varrho \rangle$ is in l-$\mathsf{Max}_{\mathsf{A}}$. Without loss of generality, since the game form $\mathcal{F}$ is deterministic, we assume that $\mathsf{O} = \varrho(\mathsf{Act}_{\mathsf{A}}, \mathsf{Act}_{\mathsf{B}})$. We assume that $\mathsf{Act}_{\mathsf{A}} = [\![1, n]\!]$, $\mathsf{Act}_{\mathsf{B}} = [\![1, m]\!]$ and we consider the formula $\Phi_{\mathcal{F}}^{l-\mathsf{Max}_{\mathsf{A}}}$ below, that we will explain line by line in the following.

$$\begin{aligned}
\Phi_{\mathcal{F}}^{l-\mathsf{Max}_{\mathsf{A}}} := \exists(\sigma_k = \sigma_k^1, \ldots, \sigma_k^n)_{1 \leq k \leq l}, \bigwedge_{1 \leq k \leq l} \mathsf{IsStrategy}_{\mathsf{A}}(\sigma_k) \land \\
\forall v = (v_o)_{o \in \mathsf{O}}, \mathsf{RealBetweenZeroOne}(v) \land \\
\exists u, (0 \leq u \leq 1) \land \\
\exists \sigma_{\mathsf{B}} = \sigma_{\mathsf{B}}^1, \ldots, \sigma_{\mathsf{B}}^m, \mathsf{IsStrategy}_{\mathsf{B}}(\sigma_{\mathsf{B}}) \land \mathsf{Val}_{\mathsf{B}}(\sigma_{\mathsf{B}}, v, u) \land \\
\bigvee_{1 \leq k \leq l} \mathsf{Val}_{\mathsf{A}}(\sigma_k, v, u)
\end{aligned}$$

Before detailing all the predicates occurring in the formula $\Phi_{\mathcal{F}}^{l-\mathsf{Max}_{\mathsf{A}}}$, note that the formula $\Phi_{\mathcal{F}}^{l-\mathsf{Max}_{\mathsf{A}}}$ does not fit exactly the formalism of Definition 9.5: all the quantifiers are not at the beginning of the formula. However, the semantics of the formula does not change if these quantifiers are moved at the beginning. We present the formula $\Phi_{\mathcal{F}}^{l-\mathsf{Max}_{\mathsf{A}}}$ in that way for readability.

The first line encodes the existence of a set of l Player-A GF-strategies among which we will later find optimal GF-strategies for Player A. For all $1 \leq k \leq l$, the predicate $\mathsf{IsStrategy}(\sigma_k)$ checks that σ_k is indeed a Player-A

GF-strategy in $\mathcal{F}$. That is:

$$\mathsf{IsStrategy}_{\mathsf{A}}(\sigma_k) := \bigwedge_{1 \leq i \leq n} ((0 \leq \sigma_k^i \leq 1) \wedge (\sum_{i=1}^{n} \sigma_k^i = 1))$$

Then, we quantify over all the possible valuations v of the outcomes. This is checked by the predicate $\mathsf{RealBetweenZeroOne}(v)$:

$$\mathsf{RealBetweenZeroOne}(v) := \bigwedge_{o \in \mathsf{O}} (0 \leq v_o \leq 1)$$

The value of the game in normal form $\langle \mathcal{F}, v \rangle$ is then equal to u (which is in $[0, 1]$), which will be checked by the following predicates. First, we exhibit a Player-B GF-strategy $\sigma_{\mathsf{B}} \in \Sigma_{\mathsf{B}}(\mathcal{F})$ whose value is at most u. The predicate $\mathsf{IsStrategy}_{\mathsf{B}}(\sigma_B)$ checks that σ_{B} is indeed a Player-B GF-strategy:

$$\mathsf{IsStrategy}_{\mathsf{B}}(\sigma_{\mathsf{B}}) := \bigwedge_{1 \leq j \leq m} ((0 \leq \sigma_{\mathsf{B}}^j \leq 1) \wedge (\sum_{j=1}^{m} \sigma_{\mathsf{B}}^j = 1))$$

The predicate $\mathsf{Val}_{\mathsf{B}}(\sigma_{\mathsf{B}}, v, u)$ checks that the value of this Player-B GF-strategy σ_{B} in the game in normal form $\langle \mathcal{F}, v \rangle$ is at most u:

$$\mathsf{Val}_{\mathsf{B}}(\sigma_{\mathsf{B}}, v, u) := \bigwedge_{1 \leq i \leq n} (\sum_{j=1}^{m} \sigma_{\mathsf{B}}^j \cdot v_{\varrho(i,j)} \leq u)$$

Finally, we check that there is a Player-A GF-strategy σ_k for some $1 \leq k \leq l$ whose value is at least u in the game in normal form $\langle \mathcal{F}, v \rangle$. This is checked by the predicate $\mathsf{Val}_{\mathsf{A}}(\sigma_k, v, u)$:

$$\mathsf{Val}_{\mathsf{A}}(\sigma_k, v, u) := \bigwedge_{1 \leq j \leq m} (\sum_{i=1}^{n} \sigma_k^i \cdot v_{\varrho(i,j)} \geq u)$$

Therefore, we have that $\mathcal{F}$ is l-$\mathsf{Max}_{\mathsf{A}}$ if and only if $\Phi_{\mathcal{F}}^{l-\mathsf{Max}_{\mathsf{A}}} \in \mathsf{True}_{\mathsf{FO}-\mathbb{R}}$. In addition, the size of the formula $\Phi_{\mathcal{F}}^{l-\mathsf{Max}_{\mathsf{A}}}$ is polynomial in the size of $\mathcal{F}$ and l. □

We can therefore conclude that deciding if a game form is uniquely (resp. finitely) maximizable is decidable (resp. semi-decidable).

Corollary 9.7. *It is decidable if a standard finite deterministic game form is uniquely maximizable w.r.t. Player* A *and it is semi-decidable if a standard finite deterministic game form is finitely maximizable w.r.t. Player* A*.*

Proof. This is direct consequence of Proposition 9.6 for game forms uniquely maximizable w.r.t. Player A . Consider now the case of game forms finitely maximizable w.r.t. Player A. Recall that a decision problem L is semi-decidable if there is an algorithm terminating on all positive instances of L (but that does not necessarily terminate on all instances) such that, whenever it terminates, it accepts all positives instances of L and rejects all negative instances. Let us design such an algorithm. Given a game form $\mathcal{F}$, one can check, for all $n \in \mathbb{N}$, if $\mathcal{F}$ is n-maximizable, in which case one accepts $\mathcal{F}$. The correctness of the algorithm is a direct consequence of the fact that all game forms finitely maximizable w.r.t. Player A are n-maximizable for some $n \in \mathbb{N}$. □

Open Question 9.1. *It is not known if it can be decided that a standard finite deterministic game form is finitely maximizable w.r.t. Player* A.

The reason why we cannot encode this problem in the first order theory of the reals as we did above is the following. We do not know any bound in the size of the finite set maximizing a game form, assuming there is one that does. Hence, encoding a set of arbitrary, yet finite, size seems to require infinitely many variables.

9.2.3 . Relevant environments and positionally optimizable game forms

Let us now focus on positionally optimizable game forms. Recall that these are the game forms we defined in in the previous chapter. The formal definition is given in Definition 8.7. The definition of environments (i.e. Definition 8.4) will be extensively used in this subsection (and the next). Before encoding the corresponding decision problem in the first order theory of the reals, let us first state that, to prove that a game form is positionally optimizable, one does not need to consider all possible environments. In fact, it is sufficient to consider relevant environments, defined below.

Definition 9.8 (Relevant environments). *For a set of outcomes* O*, an environment* $E = \langle c, e, p \rangle \in \mathsf{Env}(\mathsf{O})$ *is* relevant *if* $c \in \{0, 1\}$, $p^{-1}[\{c-1\}] = p^{-1}[\{q_{\mathsf{init}}\}] = \emptyset$ *and, for all* $i \in [\![c, e]\!]$*, there is* $o \in \mathsf{O}$ *such that* $p(o) = k_i$*. The* size *of a relevant environment* E *is equal to* $\mathsf{Sz}(E) := e - c$.

As mentioned above, we may only consider relevant environment to decide if a game form is positionally maximizable.

Proposition 9.8 (Proof 9.5.1). *Consider a set of outcomes* O *and a game form* $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$*. Consider a Player* $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$*. For all* $n \in \mathbb{N}$*, the game form* $\mathcal{F}$ *is positionally maximizable w.r.t. Player* C *up to* n *if and only if, for all relevant environments* E *with* $\mathsf{Sz}(E) \leq n-1$*, there is an optimal* GF*-strategy in* $\mathcal{F}$ *for Player* C *w.r.t.* $(\mathcal{F}, E)$.

The benefit of considering only relevant environment is that, given any $l \in \mathbb{N}$, deciding if a game form is positionally optimizable up to l can be done by considering environments where all the outcomes are mapped to the indices at most l. Then, the fact that a standard finite deterministic game form is positionally optimizable can be encoded in the first-order theory of the reals.

Proposition 9.9. *Consider a standard finite deterministic game form $\mathcal{F}$. The fact that $\mathcal{F} \in \mathsf{ParO}$ (resp. $\mathcal{F} \in \mathsf{ParO}(n)$, for some $n \in \mathbb{N}$) can be encoded, with formula of size polynomial in $|\mathcal{F}|$ (resp. and n), in* FO-$\mathbb{R}$.

Proof. Consider some $l \in \mathbb{N}$. We encode the fact that a standard finite deterministic game form $\mathcal{F} = \langle \mathsf{Act_A}, \mathsf{Act_B}, \mathsf{O}, \varrho \rangle$ is positionally optimizable up to l. Without loss of generality, since the game form $\mathcal{F}$ is deterministic, we assume that $\mathsf{O} = \varrho(\mathsf{Act_A}, \mathsf{Act_B})$. To do so, we use the characterization of Proposition 9.8 and we express in FO-$\mathbb{R}$ the fact that, for all relevant environments $E = \langle c, e, p \rangle$ of size at most $l-1$, both players have an optimal GF-strategy. We use the characterization given by Lemma 8.2. We assume that $\mathsf{Act_A} = [\![1, n]\!]$, $\mathsf{Act_B} = [\![1, m]\!]$ and we consider the formula $\Phi_{\mathcal{F}}^{\mathsf{ParO}(l)}$ below, that we will explain line by line in the following.

$$\Phi_{\mathcal{F}}^{\mathsf{ParO}(l)} := \forall c,\ \mathsf{ZeroOrOne}(c) \land$$
$$\forall \alpha = (\alpha_o)_{o \in \mathsf{O}},\ \mathsf{ZeroOrOne}(\alpha) \land$$
$$\forall v = (v_o)_{o \in \mathsf{O}},\ \mathsf{RealBetweenZeroOne}(v) \land$$
$$\forall k = (k_o)_{o \in \mathsf{O}},\ \mathsf{IntBetweenZeroL}(k, c) \land$$
$$\exists u,\ (0 \leq u \leq 1) \land$$
$$\exists S_\mathsf{A} = S_\mathsf{A}^1, \ldots, S_\mathsf{A}^n,$$
$$\exists \sigma_\mathsf{A} = \sigma_\mathsf{A}^1, \ldots, \sigma_\mathsf{A}^n,\ \mathsf{IsIndicator}(S_\mathsf{A}) \land \mathsf{IsStrategy}_\mathsf{A}(\sigma_\mathsf{A}) \land \mathsf{IsSupp}(\sigma_\mathsf{A}, S_\mathsf{A}) \land \mathsf{Val}_\mathsf{A}(\sigma_\mathsf{A}, \alpha, v, u) \land$$
$$\bigwedge_{1 \leq j \leq m} \mathsf{MaxIntegerEven}(S_\mathsf{A}, k, j, c) \land$$
$$\exists S_\mathsf{B} = S_\mathsf{B}^1, \ldots, S_\mathsf{B}^m,$$
$$\exists \sigma_\mathsf{B} = \sigma_\mathsf{B}^1, \ldots, \sigma_\mathsf{B}^n,\ \mathsf{IsIndicator}(S_\mathsf{B}) \land \mathsf{IsStrategy}_\mathsf{B}(\sigma_\mathsf{B}) \land \mathsf{IsSupp}(\sigma_\mathsf{B}, S_\mathsf{B}) \land \mathsf{Val}_\mathsf{B}(\sigma_\mathsf{B}, \alpha, v, u) \land$$
$$\bigwedge_{1 \leq i \leq n} \mathsf{MaxIntegerOdd}(S_\mathsf{B}, k, i, c)$$

As in the proof of Proposition 9.6, the formula $\Phi_{\mathcal{F}}^{\mathsf{ParO}(l)}$ does not fit exactly the formalism of Definition 9.5, since all the quantifiers are not at the beginning of the formula. However, as the proof of Proposition 9.6, the semantics is changed if these quantifiers are moved at the beginning of the formula. We presented $\Phi_{\mathcal{F}}^{\mathsf{ParO}(l)}$ in that way for readability.

The first four lines encode the relevant environment E. Specifically:

$$\mathsf{ZeroOrOne}(c) := (c = 0) \lor (c = 0)$$

Since the environment E is relevant, c is equal to either 0 or 1. Furthermore:

$$\mathsf{ZeroOrOne}(\alpha) := \bigwedge_{o \in \mathsf{O}} ((\alpha_o = 0) \lor (\alpha_o = 1))$$

where, for all $o \in \mathsf{O}$, $\alpha_o = 1$ means that $p(o) \in [0,1]$ and $\alpha_o = 0$ means that $p(o) \in K_l$. Furthermore:

$$\mathsf{RealBetweenZeroOne}(v) := \bigwedge_{o \in \mathsf{O}} (0 \leq v_o \leq 1)$$

where, for all $o \in \mathsf{O}$, $v_o \in [0,1]$ corresponds to the value of o w.r.t. p (assuming $p(o) \in [0,1]$). Furthermore:

$$\mathsf{IntBetweenZeroL}(k,c) := \bigwedge_{o \in \mathsf{O}} (\bigvee_{0 \leq p \leq l} k_o = p) \wedge ((c = 0) \Rightarrow (k_o \leq l-1))$$

where, for all $o \in \mathsf{O}$, $k_o \in [\![0,l]\!]$ is the index of o given by p (assuming that $p(o) \in K_l$). Furthermore, since the size of E is at most $l-1$, if $c = 0$, then the maximum of the colors should be at most $l-1$.

Then, u is the value of the game $\mathcal{G}_{(\mathcal{F},E)}$ from the state q_{init}. This is ensured by the remainder of the formula: it exhibits a GF-strategy per player whose corresponding positional strategy parity dominates the valuation v_Y^u for $Y := (\mathcal{F}, E)$. The predicates are used for both players, we give them only for Player A, the case of Player B being analogous.

$$\mathsf{IsIndicator}(S_{\mathsf{A}}) := \bigwedge_{1 \leq i \leq n} ((S_{\mathsf{A}}^i = 0) \vee (S_{\mathsf{A}}^i = 1))$$

where S_{A} encodes the support of the GF-strategy σ_{A}. The fact that it is indeed a Player-A GF-strategy is ensured by the predicates below:

$$\mathsf{IsStrategy}_{\mathsf{A}}(\sigma_{\mathsf{A}}) := \bigwedge_{1 \leq i \leq n} ((0 \leq \sigma_{\mathsf{A}}^i \leq 1) \wedge (\sum_{i=1}^{n} \sigma_{\mathsf{A}}^i = 1))$$

In addition:

$$\mathsf{IsSupp}(\sigma_{\mathsf{A}}, S_{\mathsf{A}}) := \bigwedge_{1 \leq i \leq n} ((\sigma_{\mathsf{A}}^i > 0) \Leftrightarrow (S_{\mathsf{A}}^i = 1))$$

This predicate ensures that S_{A} does indeed correspond to the support of the GF-strategy σ_{A}. Furthermore:

$$\mathsf{Val}_{\mathsf{A}}(\sigma_{\mathsf{A}}, \alpha, v, u) := \bigwedge_{1 \leq j \leq k} (\sum_{i=1}^{n} \sigma_{\mathsf{A}}^i \cdot (\alpha_{\varrho(i,j)} \cdot v_{\varrho(i,j)} + (1 - \alpha_{\varrho(i,j)}) \cdot u) \geq u)$$

This predicates encodes the fact that the Player-A GF-strategy σ_{A} dominates the valuation v_Y^u. Indeed, note that for all $o \in \mathsf{O}$, if $\alpha_o = 1$, then $p(o) \in [0,1]$, and therefore $v_Y^u(o) = p(o) = v_o$. Similarly, if $\alpha_o = 0$, then $p(o) \in K_l$, and therefore $v_Y^u(o) = u$. Finally, when l is even, denoting $l = 2 \cdot x$ with $x \in \mathbb{N}$:

$$\begin{aligned}\mathsf{MaxIntegerEven}(S_{\mathsf{A}}, k, j, c) := & \bigvee_{1 \leq i \leq n} (S_{\mathsf{A}}^i = 1 \wedge \alpha_{\varrho(i,j)} = 1) \vee \\ & \bigvee_{0 \leq e \leq x} (2e \geq c) \wedge ((\bigvee_{1 \leq i \leq n} (S_{\mathsf{A}}^i = 1) \wedge (k_{\varrho(i,j)} = 2e)) \wedge \\ & (\bigwedge_{1 \leq i \leq n} (S_{\mathsf{A}}^i = 1) \Rightarrow (k_{\varrho(i,j)} \leq 2 \cdot e)))\end{aligned}$$

It is similar if l is odd, up to small changes. This predicate encodes item $(ii.2)$ of Lemma 8.2: for all columns $j \in \mathsf{Act_B}$, either there is positive probability to see an outcome mapped to a value in $[0,1]$ (i.e. $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[0,1]} \rangle](\sigma_{\mathsf{A}}, j) > 0$) or the maximum of the colors in the support of the strategy σ_{A} is even (i.e. $\max(\mathsf{Color}(\mathcal{F}, p, \sigma_{\mathsf{A}}, j) \cup \{c\})$ is even).

By Lemma 8.2 and Proposition 9.8, we have that $\mathcal{F} \in \mathsf{ParO}(l)$ if and only if $\Phi_{\mathcal{F}}^{\mathsf{ParO}(l)} \in \mathsf{True}_{\mathsf{FO}-\mathbb{R}}$. Furthermore, the size of $\Phi_{\mathcal{F}}^{\mathsf{ParO}(l)}$ is polynomial in the size of $\mathcal{F}$ (and in l).

Finally, since given an environment, an outcome is mapped to at most one index, and since in relevant environments of size $n \in \mathbb{N}$, the outcomes are mapped to at least $n-1$ different indices, we have $\mathcal{F} \in \mathsf{ParO} \Leftrightarrow \mathcal{F} \in \mathsf{ParO}(n+2)$ for $n := |\mathsf{O}|$. The proposition follows. □

9.3 Comparing classes of game forms

In this section, we compare the strengths of the game form properties we have defined in this part. For these properties to be well defined, we need to restrict ourselves to standard finite game forms (note however that we do not restrict ourselves to deterministic game forms). In addition, when applicable, the properties we consider on game forms hold for both players. This is the case for finitely/uniquely maximizable and positionally optimizable game forms. That way, how all these properties compare is not too complicated to describe. More specifically, the goal of this section is to prove the results summarized in Figure 9.3, and stated formally in Theorem 9.10 below.

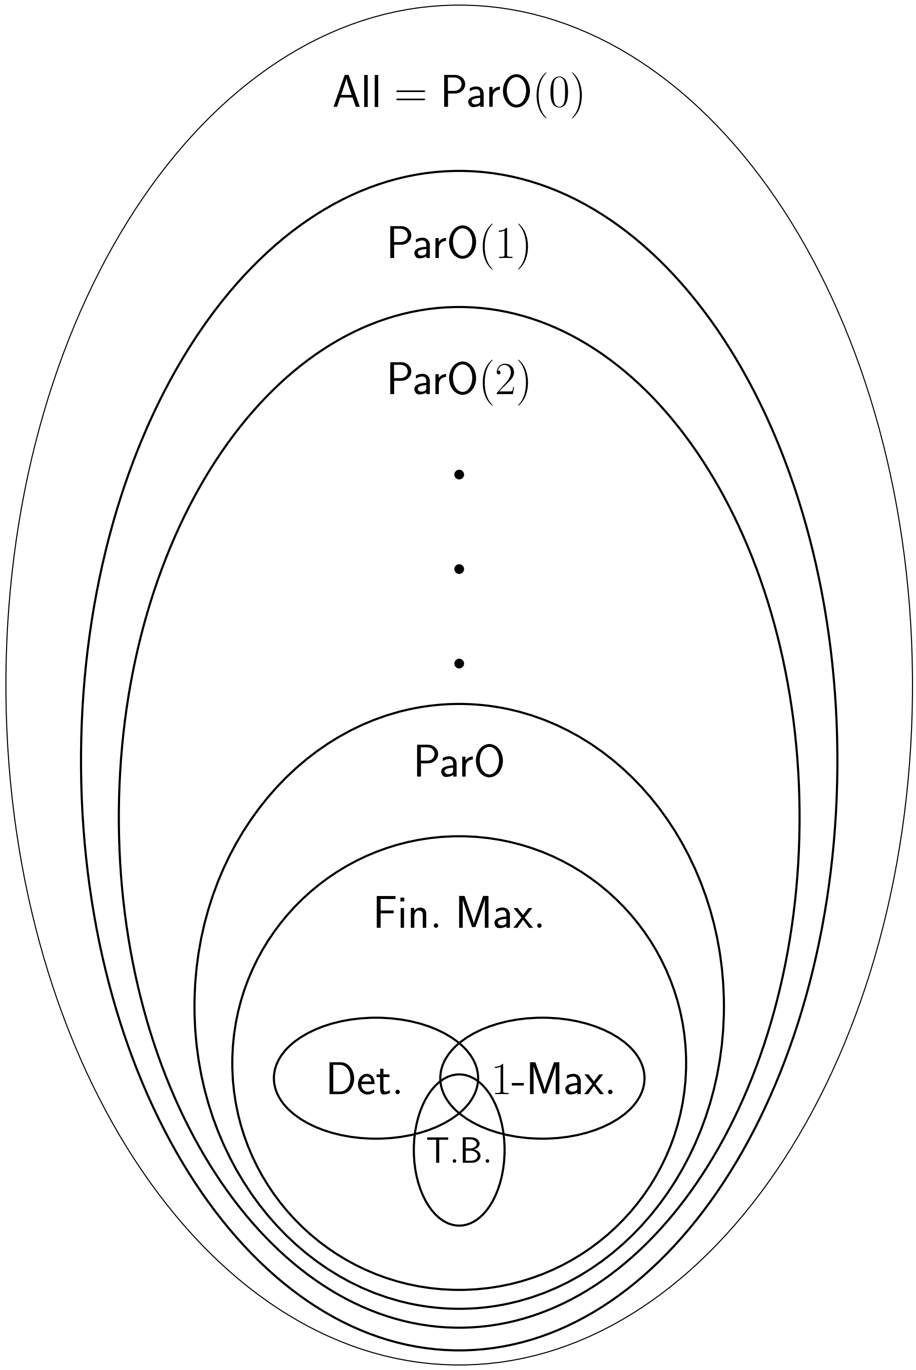

Figure 9.3: This summarizes how properties on standard finite game forms compare in strength, that is formally stated in Theorem 9.10. If an ellipse (or an intersection of ellipses) labeled with X is contained in another ellipse labeled with X, this means that the set of Y game forms is strictly included in the set of X game forms.

Theorem 9.10. *Among the set of all standard finite game forms, denoting* All *the set of such game forms,* $\mathsf{Det.}$ *the set of deterministic game forms that are determined,* $\mathsf{Fin.\ Max.}$ *(resp.* $\mathsf{1\text{-}Max.}$*) the set of game forms finitely (resp. uniquely) maximizable for both players, and* $\mathsf{T.B.}$ *the set of turn-based game forms, from top to bottom in Figure 9.3, we have:*

1. $\mathsf{All} = \mathsf{ParO}(0)$*;*
2. *for all* $n \in \mathbb{N}$*,* $\mathsf{ParO}(n) \supsetneq \mathsf{ParO}(n+1)$*;*
3. *for all* $n \in \mathbb{N}$*,* $\mathsf{ParO}(n) \supsetneq \cap_{k \in \mathbb{N}} \mathsf{ParO}(k) = \mathsf{ParO}$*;*
4. $\mathsf{ParO} \supsetneq \mathsf{Fin.Max.}$*;*
5. $\mathsf{Fin.Max.} \supsetneq \mathsf{Det.} \cup \mathsf{1\text{-}Max.} \cup \mathsf{T.B.}$*;*
6. $\forall X \in \{\mathsf{Det.}, \mathsf{1\text{-}Max.}\}$*,* $X \supsetneq \mathsf{Det.} \cap \mathsf{1\text{-}Max.}$*;*
7. $\forall X \in \{\mathsf{Det.}, \mathsf{T.B.}\}$*,* $X \supsetneq \mathsf{Det.} \cap \mathsf{T.B.}$*;*
8. $\forall X \in \{\mathsf{T.B.}, \mathsf{1\text{-}Max.}\}$*,* $X \supsetneq \mathsf{T.B.} \cap \mathsf{1\text{-}Max.}$*;*
9. $\forall X \in \{\mathsf{Det.} \cap \mathsf{1\text{-}Max.}, \mathsf{Det.} \cap \mathsf{T.B.}, \mathsf{T.B.} \cap \mathsf{1\text{-}Max.}\}$*,* $X \supsetneq \mathsf{Det.} \cap \mathsf{T.B.} \cap \mathsf{1\text{-}Max.}$.

Among all these items, there are only two that are not straightforward to prove: items 2 and 4. We state and prove propositions corresponding to these items and then formally prove Theorem 9.10. Let us first start with the simplest of these items: item 4.

Proposition 9.11. *The set of standard finite game forms positionally optimizable (i.e.* ParO*) strictly contains the set of game forms finitely maximizable w.r.t. both players (i.e.* $\mathsf{Fini.\ Max.}$*). That is,* $\mathsf{ParO} \supsetneq \mathsf{Fini.\ Max.}$.

Proof. The inclusion $\mathsf{ParO} \supseteq \mathsf{Fin.\ Max.}$ is a direct corollary of Corollary 6.12 applied to both players: in a finite concurrent parity game where all local interactions are finitely maximizable w.r.t. both players, both players have an optimal positional strategy.

Let us now show that this inclusion is strict, that is let us exhibit a standard finite game form that is positionally optimizable but not finitely maximizable for any player. Consider the standard finite game form $\mathcal{F} \in \mathsf{Form}(\mathsf{O})$ depicted in Figure 9.4 where $\mathsf{O} := \{x, y, z\}$. Both players have two available actions, we let $\mathsf{Act}_\mathsf{A} := \{a_\mathsf{t}, a_\mathsf{b}\}$ where a_t (resp. a_b) is the Player-A action corresponding to the top (resp. bottom) row in $\mathcal{F}$. Similarly, we let $\mathsf{Act}_\mathsf{B} := \{b_\mathsf{l}, b_\mathsf{r}\}$ where b_l (resp. b_r) is the Player-B action corresponding to the left (resp. right) column in $\mathcal{F}$. Let us show that this game form is positionally optimizable.

Consider any environment $E = \langle c, e, p \rangle \in \mathsf{Env}(\mathsf{O})$ and the corresponding parity game $\mathcal{G}_{(\mathcal{F},E)}$ (recall Definition 8.5). We let $u := \chi_{\mathcal{G}_{(\mathcal{F},E)}}(q_{\mathsf{init}})$. For all

$$\mathcal{F} = \begin{bmatrix} \frac{2x+y+z}{4} & \frac{3y+z}{4} \\ & \\ \frac{3y+z}{4} & \frac{y+3z}{4} \end{bmatrix}$$

Figure 9.4: A game form $\mathcal{F}$ that is positionally optimizable but not finitely maximizable.

$$\langle \mathcal{F}, v \rangle = \begin{bmatrix} \frac{2\alpha+1}{4} & \frac{3}{4} \\ & \\ \frac{3}{4} & \frac{1}{4} \end{bmatrix}$$

Figure 9.5: The game in normal form $\langle \mathcal{F}, v_\alpha \rangle$ obtained from the game form $\mathcal{F}$ of Figure 9.4.

$t \in \mathsf{O}$, if $p(t) \notin [0,1]$, we let $n_t \in [\![0,e]\!]$ denote the integer that t is mapped to w.r.t. p, i.e. the integer ensuring $p(t) = k_{n_t}$. There are two cases:

- Assume that $p(y), p(z) \notin [0,1]$. Then, $u \in \{0,1\}$. Indeed, by playing actions $a_{\mathsf{b}} \in \mathsf{Act}_{\mathsf{A}}$ and $b_{\mathsf{r}} \in \mathsf{Act}_{\mathsf{B}}$ at state q_{init} in the game $\mathcal{G}_{(\mathcal{F},E)}$, both players can ensure that: 1) surely, no stopping state is reached and, 2) almost surely, the highest color seen infinitely often is $\max(c, n_y, n_z)$. Hence, if $\max(c, n_y, n_z)$ is even, we have that $u = 1$, the Player-A GF-strategy $a_{\mathsf{b}} \in \mathcal{D}(\mathsf{Act}_{\mathsf{A}})$ is optimal w.r.t. $(\mathcal{F}, E)$ and any Player-B GF-strategy is optimal w.r.t. $(\mathcal{F}, E)$. This is symmetrical if $\max(c, n_y, n_z)$ is odd.

- Assume that $p(y) \in [0,1]$ or $p(z) \in [0,1]$. In that case, one can realize that, for all Player-A GF-strategies $\sigma_{\mathsf{A}} \in \mathcal{D}(\mathsf{Act}_{\mathsf{A}})$ and Player-B action $b \in \mathsf{Act}_{\mathsf{B}}$, we have: $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[0,1]} \rangle](\sigma_{\mathsf{A}}, b) \geq \frac{1}{4}$. In other words, in the parity game $\mathcal{G}_{(\mathcal{F},E)}$, regardless of what the players do, there is probability at least $\frac{1}{4}$ to exit to a stopping state. Therefore, by Lemma 8.2, any Player-A GF-strategy σ_{A} optimal in the game in normal form $\langle \mathcal{F}, v^u_{(\mathcal{F},E)} \circ p \rangle$ is optimal w.r.t. $(\mathcal{F}, E)$, and such GF-strategies exist since $\mathcal{F}$ is standard finite. This is similar for Player B.

In any case, both players have an optimal GF-strategy. Thus, the game form $\mathcal{F}$ is positionally optimizable.

Let us now show that $\mathcal{F}$ is not finitely maximizable w.r.t. any player. For all $\alpha \in [0,1]$, we let $v_\alpha : \mathsf{O} \to [0,1]$ be a valuation of the outcomes such that $v(y) := 1$, $v(z) := 0$ and $v(x) := \alpha$. The game in normal form $\langle \mathcal{F}, v_\alpha \rangle$ that we obtain is depicted in Figure 9.5. Then, one can check that the value $u_\alpha := \mathsf{val}[\langle \mathcal{F}, v_\alpha \rangle] \in [0,1]$ of the game in normal form $\langle \mathcal{F}, v_\alpha \rangle$ is equal to:

$$u_\alpha = \frac{4-\alpha}{8-4\alpha}$$

Furthermore, the only Player-A GF-strategy $\sigma^\alpha_{\mathsf{A}}$ that is optimal in the game in normal form $\langle \mathcal{F}, v \rangle$ is such that:

$$\sigma^\alpha_{\mathsf{A}}(a_{\mathsf{t}}) := \frac{1}{2-\alpha}$$

and similarly the only Player-B GF-strategy $\sigma_{\mathsf{B}}^{\alpha}$ that is optimal in the game in normal form $\langle \mathcal{F}, v \rangle$ is such that:

$$\sigma_{\mathsf{B}}^{\alpha}(b_{\mathsf{l}}) := \frac{1}{2-\alpha}$$

Therefore, for both players, playing optimally in all games in normal form that can be obtained from $\mathcal{F}$ requires playing infinitely many different GF-strategies. □

We can now consider item 2 of Theorem 9.10. It is formally stated as Proposition 9.12 below.

Proposition 9.12. *For all $n \in \mathbb{N}$, the set of game forms positionally optimizable up to n (i.e. $\mathsf{ParO}(n)$) strictly contains the set of game forms optimizable up to $n+1$ (i.e. $\mathsf{ParO}(n+1)$). That is, we have $\mathsf{ParO}(n) \supsetneq \mathsf{ParO}(n+1)$.*

The inclusion $\mathsf{ParO}(n) \supseteq \mathsf{ParO}(n+1)$ comes directly from the definition of positionally optimizable game forms (i.e. Definition 8.7). Then, for all $n \geq 1$, we exhibit a standard finite game form $\mathcal{F}_n$ that is in $\mathsf{ParO}(n-1)$ but not in $\mathsf{ParO}(n)$. This is done in Definition 9.9 below, where we describe a game form that is $\mathsf{ParO}(n)$ but not in $\mathsf{ParO}(n-1)$ in the case where n is even.

Definition 9.9. *Consider some even $n \geq 2$. We consider the set of outcomes $\mathsf{O}_n := \{x_0, x_1, \ldots, x_{n-1}, y, z\}$ and we consider the standard finite game form $\mathcal{F}_n = \langle \mathsf{Act}_{\mathsf{A}}^n, \mathsf{Act}_{\mathsf{B}}^n, \mathsf{O}_n, \varrho_n \rangle_{\mathsf{s}}$ depicted in Figure 9.6. Let us describe the set of actions of the players. We set $\mathsf{Act}_{\mathsf{A}}^n := \{a_{\mathsf{t}}, a_{\mathsf{b}}, a_1, \ldots, a_{n-1}, a_{1,0}, \ldots, a_{n-1,n-2}, a_{\mathsf{Ex}}\}$ and $\mathsf{Act}_{\mathsf{B}}^n := \{b_{\mathsf{l}}, b_{\mathsf{r}}, b_0, \ldots, b_{n-2}, b_{2,1}, \ldots, b_{n-2,n-3}, b_{\mathsf{Ex}}\}$. The actions $a_{\mathsf{t}}, a_{\mathsf{b}}$ (resp. $b_{\mathsf{l}}, b_{\mathsf{r}}$) refer to the two topmost rows (resp. leftmost columns): a_{t} (resp. b_{l}) leads to $\frac{x_0+\ldots+x_{n-1}}{n}, \frac{3y+z}{4}$ whereas a_{b} (resp. b_{r}) leads to $\frac{3y+z}{4}, \frac{y+3z}{4}$. Then, for all odd $i \leq n-1$ (resp. even $j \leq n-2$) a_i and $a_{i,i-1}$ (resp. b_j and $b_{j,j-1}$ — only if $j \geq 2$) correspond to the rows leading to x_i and $\frac{x_i+x_{i-1}}{2}$ (resp. columns leading to x_j and $\frac{x_j+x_{j-1}}{2}$) respectively. Finally, action a_{Ex} (resp. b_{Ex}) correspond to the bottommost row (resp. rightmost column).*

The game forms described in Definition 9.9 satisfies the lemma below.

Lemma 9.13 (Proof 9.5.2)**.** *Consider some even $n \geq 2$. The game form of $\mathcal{F}_n$ Definition 9.9 is:*

- *positionally maximizable w.r.t. Player* B*;*
- *positionally maximizable w.r.t. Player* A *up to $n-1$ but not up to n.*

That is: $\mathcal{F}_n \in \mathsf{ParO}(n) \setminus \mathsf{ParO}(n-1)$.

The proof of this lemma is very tedious since it requires to consider a lot of different possibilities in terms of which outcome is mapped to which value or color. Let us first give below an informal proof of this lemma in the case where $n = 2$. The game form $\mathcal{F}_2$ is depicted in Figure 9.7.

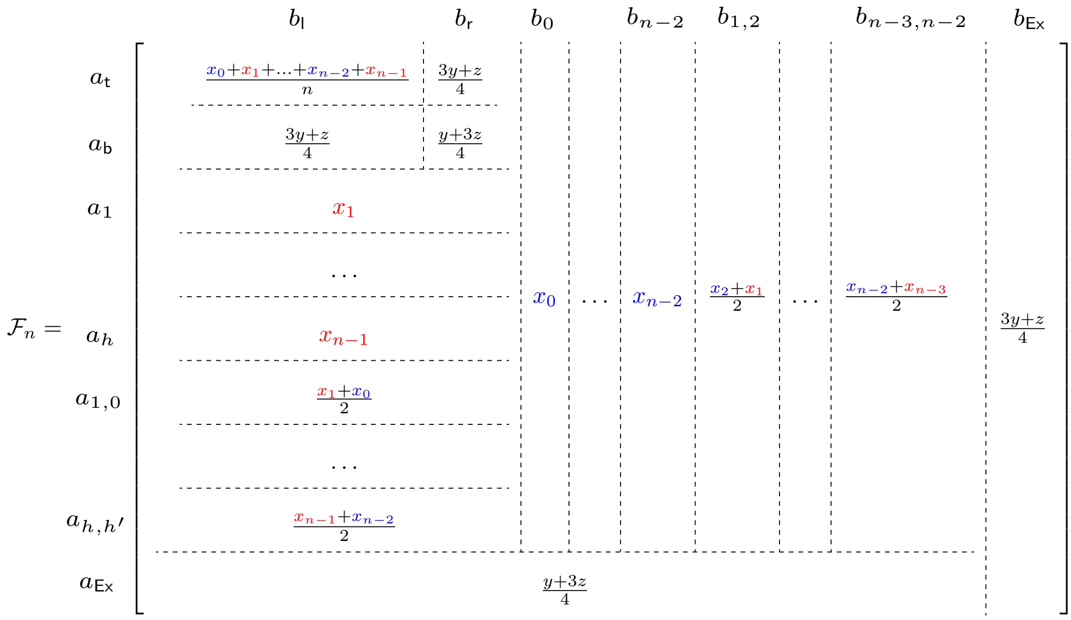

Figure 9.6: The game form $\mathcal{F}_n$. Due to a lack of space, h refers to $n-1$ and h' to $n-2$.

Proof sketch. Consider a relevant environment $E = \langle c, e, p \rangle \in \mathsf{Env}(\mathsf{O})$ and the corresponding parity game $\mathcal{G}_{(\mathcal{F},E)}$ (recall Definition 8.5). We let $u := \chi_{\mathcal{G}_{(\mathcal{F}_2,E)}}(q_\mathsf{init})$. For all $t \in \mathsf{O}$, if $p(t) \notin [0,1]$, we let $n_t \in [\![0,e]\!]$ denote the integer that t is mapped to w.r.t. p, i.e. the integer ensuring $p(t) = k_{n_t}$. We want to show that, in any case, Player B has a GF-strategy that is optimal w.r.t. $(\mathcal{F}, E)$ and that this also holds for Player A as long as $c = e$ (i.e. if $\mathsf{Sz}(E) = 0 = n-2$). However, there is a relevant environment E with $\mathsf{Sz}(E) = 1$ such that Player A has no GF-strategy optimal w.r.t. $(\mathcal{F}, E)$. Let us explain what happens for Player A. There are several cases, we detail some of them.

- Assume that $p(y), p(z) \notin [0,1]$. As in the proof of Proposition 9.11, this implies $u \in \{0,1\}$ and playing action a_Ex for Player A and action b_Ex for Player B is optimal w.r.t. $(\mathcal{F}, E)$.

- Assume now that $p(y) \in [0,1]$ and $p(z) \notin [0,1]$. Then, we have $u = p(y)$. The reason is because, by playing action a_Ex, Player A ensures that almost surely a stopping state of value $p(y)$ is reached. Furthermore, Player B can ensure the same thing by playing action b_Ex. Then, playing action a_Ex for Player A and action b_Ex for Player B is optimal w.r.t. $(\mathcal{F}, E)$. This is similar if $p(y) \notin [0,1]$ and $p(z) \in [0,1]$.

- Assume now $p(y), p(z) \in [0,1]$. If $p(y) < p(z)$, then we have $u = \frac{3p(y)+p(z)}{4}$, and as before playing action a_Ex for Player A and action b_Ex

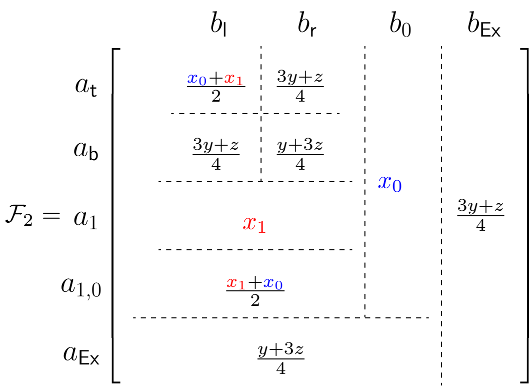

Figure 9.7: The game form $\mathcal{F}_2$.

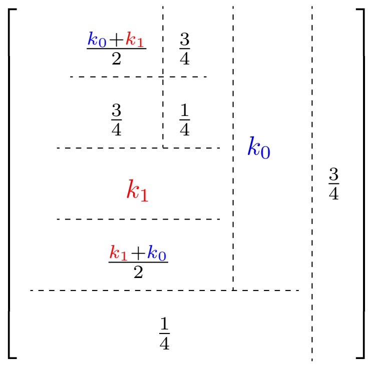

Figure 9.8: The game form $\mathcal{F}_2$ in a specific environment.

for Player B is optimal w.r.t. $(\mathcal{F}, E)$.

- Assume now that $p(y), p(z) \in [0,1]$ and $p(y) \geq p(z)$. If $p(x_0) \in [0,1]$ or $p(x_1) \in [0,1]$, we can also exhibit an optimal GF-strategy for both players.

- Let us now assume that $p(y), p(z) \in [0,1]$ with $p(y) \geq p(z)$. For simplicity, we assume that $p(y) = 1$ and $p(z) = 0$. In that case, we have:

$$\frac{1}{4} = \frac{3p(z) + p(y)}{4} \leq u \leq \frac{p(z) + 3p(y)}{4} = \frac{3}{4}$$

  Assume also that $p(x_0), p(x_1) \notin [0,1]$. Assume first that n_{x_0} is odd. In that case, $u = \frac{1}{4}$. A Player-B strategy playing positionally action b_0 achieves this value in the game $\mathcal{G}_{(\mathcal{F}_2,E)}$. Thus, the Player-A GF-strategy a_{Ex} is optimal w.r.t. $(\mathcal{F}_2, E)$.

  Assume now that n_{x_0} is even. Then, in any case, we have $u = \frac{3}{4}$. Indeed, for all $\varepsilon > 0$, a Player-A positional strategy playing action a_{t} with probability $1-\varepsilon$ and action a_{b} with probability ε has value at least $\frac{3}{4}-\varepsilon$ in the game $\mathcal{G}_{(\mathcal{F}_2,E)}$. Therefore, playing action b_{Ex} is always optimal for Player B. It will also be the case for Player A if $\mathsf{Sz}(E) = 0$, i.e. if $c = e$. Indeed, in that case, we have either $n_{x_0} = n_{x_1} = 0$, in which case $u = \frac{3}{4}$ and the Player-A GF-strategy a_1 is optimal w.r.t. $(\mathcal{F}_2, E)$; or $n_{x_0} = n_{x_1} = 1$, in which case we are in the scope of the previous case of this item since n_{x_0} is odd. However, if $\mathsf{Sz}(E) = 1$, there may be an issue for Player A, as we show below.

Consider now a relevant environment $E = \langle c, e, p \rangle$ of size 1 such that $c := 0$, $p(y) := 1$, $p(z) := 0$, $p(x_0) = k_0$ and $p(x_1) = k_1$. We do have $\mathsf{Sz}(E) = 1$. What we obtain is depicted in Figure 9.8. In that case, we have $u := \chi_{\mathcal{G}_{\mathcal{F}_2,E}}(q_{\mathsf{init}}) = \frac{3}{4}$, as argued above in the last item. Consider any Player-A GF-strategy $\sigma_{\mathsf{A}} \in \Sigma_{\mathsf{A}}(\mathcal{F}_2)$. If it plays action a_{b} or action a_{Ex} with positive probability, then the strategy $\mathsf{s}_{\mathsf{A}}^{(\mathcal{F}_2,E)}(\sigma_{\mathsf{A}})$ does not dominate the valuation $v^u_{(\mathcal{F}_2,E)}$. However, if it does not play these actions with positive probability, the strategy $\mathsf{s}_{\mathsf{A}}^{(\mathcal{F}_2,E)}(\sigma_{\mathsf{A}})$ has value 0 in $\chi_{\mathcal{G}_{(\mathcal{F}_2,E)}}$ from q_{init}. Indeed, by positionally playing action b_{l}, Player B ensures that: 1) surely, no stopping state is seen, and 2) almost surely, the state k_1 (of color 1) is seen infinitely often, while $c = 0$. □

Let us now consider the proof of Proposition 9.12.

Proof. Let $n \geq 1$. The inclusion $\mathsf{ParO}(n-1) \supseteq \mathsf{ParO}(n)$ is direct. Lemma 9.13 ensures that these two sets are not equal when n is even. A similar lemma could be stated in the case where n is odd. □

We can now proceed to the proof of Theorem 9.10. Note that we will refer several times to game forms depicted in Chapter 6.

Proof. Let us prove all these items one by one. Recall that we only consider standard finite game forms.

1. Consider any standard finite game form $\mathcal{F}$. By the characterization of Proposition 9.8, and since no relevant environment has size -1, we have $\mathcal{F} \in \mathsf{ParO}(0)$.

2. This is given by Proposition 9.12.

3. This a direct consequence of the previous item.

4. This is given by Proposition 9.11.

5. By definition, if a game form is uniquely maximizable, it is also finitely maximizable. In addition, by Proposition 6.5, any (deterministic) determined game form $\mathcal{F}$ is maximized w.r.t. Player $\mathsf{C} \in \{\mathsf{A}, \mathsf{B}\}$ by $\mathsf{Act}_{\mathsf{C}}$, which is finite. In addition, any game form $\mathcal{F}$ where Player A plays alone is maximized by $\mathsf{Act}_{\mathsf{A}}$ w.r.t. Player A. This is similar for Player B.

6. The game form depicted in Figure 6.1 is determined, but is not uniquely maximizable since, depending on the valuation, Player A should deterministically play on either of her three available actions.

   Furthermore, the matching pennies interaction, depicted in Figure 6.11 is uniquely maximizable, but not determined.

7. The game form depicted in Figure 6.1 mentioned in the previous item is determined, but not turn-based.

   Furthermore, all deterministic turn-based game forms are determined. However, any turn-based game form that is not deterministic (recall Definition 1.11) is not determined.

8. A game form where Player A plays alone and chooses an action among two that leads to two different outcomes (for instance depicted in the middle of Figure 1.1) is turn-based, but not uniquely maximizable.

   In addition, the game form depicted in Figure 6.12 is uniquely maximizable, but not turn-based.

9. The game form depicted in Figure 6.12 mentioned in the previous item is both determined and uniquely maximizable but not turn-based.

   A trivial game form that is not deterministic is both turn-based and uniquely maximizable but not determined.

   Finally, the turn-based deterministic game form with two possible outcomes described in the previous item (that is depicted in the middle of Figure 1.1) is both turn-based and determined but not uniquely maximizable.

□

9.4 Discussion, open questions and future work

This chapter was devoted to the study of the classes of game forms we have defined in the previous chapters of this part. As mentioned in the introduction of this chapter, we have given two kinds of results, some related to decidability/complexity, and the others related to comparing the different classes of game forms. As stated above, we leave unanswered Open Question 9.1: we do not know whether or not it is decidable that a standard finite deterministic game form is finitely maximizable w.r.t. Player A.

One can notice that, when encoding different problems in the first order theory of the reals, we did not look carefully at the exact complexity it entailed. However, as stated in this chapter (below Theorem 9.5), first order theory of the reals formulas can be decided in doubly exponential time, though one needs to be careful with this statement since the precise complexity is doubly exponential is some parameters, and polynomial in others. However, interestingly, this complexity of deciding those formulas belonging to the existential theory of the reals (i.e. with only existential quantifiers), can be done in polynomial space, as stated below Theorem 9.5. The FO-$\mathbb{R}$ formulas we have exhibited in this chapter all use both universal and existential quantifiers. However, we believe that the fact that a standard finite deterministic game form is in $\mathsf{ParO}(1)$ can be encoded in the existential theory of the reals, as stated in the conjecture below.

Conjecture 9.14. *The fact that a standard finite deterministic game form is in* $\mathsf{ParO}(1)$ *can be encoded, with a polynomial size formula, with* $\exists$-$\mathbb{R}$ *formulas.*

This would show that deciding whether a standard finite game form belongs to $\mathsf{ParO}(1)$ *can be done in polynomial space.*

A natural future work is to study the different properties that the different classes of game forms we have defined in this part enjoy. We have given some in Chapter 6. For instance, a standard deterministic game form with at least one underlying action set that is finite is determined if and only if it is semi-determined for either of the players (recall Proposition 6.8). However, there are still a lot of open questions relating these properties, and in particular maximizable game forms. We give one below, but many others could be inquired.

Open Question 9.2. *Given a standard finite deterministic game form* $\mathcal{F}$*, does it hold that if* $\mathcal{F}$ *is uniquely (or finitely) maximizable w.r.t. Player* A*, then* $\mathcal{F}$ *is finitely maximizable w.r.t. Player* B.

Finally, consider any standard finite deterministic game form ensuring any of the properties defined in this part, e.g. being positionally optimizable. Assume that we replace one of its outcomes with a standard finite deterministic game form that is itself positionally optimizable. This would require a formal definition, in particular we would have to handle properly the number of rows and columns of the new game form. Assuming that it is properly defined, we believe that the obtained standard finite deterministic game form would still be positionally optimizable. In fact, we believe that such a transformation would preserve all classes of game forms defined in this part. Hence, we believe that it constitutes the most relevant future work to inquire, since this allows to effectively build new well-behaved game forms from already existing ones.

9.5 Appendix

9.5.1 . Proof of Proposition 9.8

We prove the result for Player A, the arguments are similar for Player B. In the remainder of this section, we fix a set of outcomes O, a standard finite game form $\mathcal{F} = \langle \mathsf{Act_A}, \mathsf{Act_B}, \mathsf{O}, \varrho \rangle_\mathsf{s} \in \mathsf{Form}(\mathsf{O})$ (that need not be deterministic).

Definition 9.10. *Consider some* $n \geq 1$*. Consider an (arbitrary) environment* $E = \langle c, e, p \rangle \in \mathsf{Env}(\mathsf{O})$ *with* $p : \mathsf{O} \to \{q_{\mathsf{init}}\} \cup [0, 1] \cup K_e$ *with* $\mathsf{Sz_A}(E) = n$*. Assume that no outcome* $o \in \mathsf{O}$ *is such that* $p(o) = q_{\mathsf{init}}$*. Let* $\tilde{e} := \mathsf{Even}(e) \geq 2$*. By definition, we have* $n = \tilde{e} - c$*. Let us define a relevant environment* $E' = \langle c', e', p' \rangle \in \mathsf{Env}(\mathsf{O})$ *of size* $n - 1$.

We let:

$$c' := \begin{cases} 0 & \text{if } c \text{ is even} \\ 1 & \text{otherwise} \end{cases}$$

We also let $\mathsf{Im}_p := \{n \in [\![c+1, \tilde{e}-1]\!] \mid \exists o \in \mathsf{O},\ p(o) = k_i\}$ *and* $c = a_0 < a_1 < a_2 < \ldots < a_k$ *be such that* $\mathsf{Im}_p = \{a_1, a_2, \ldots, a_k\}$*. In particular, we have* $\tilde{e} - 1 \geq c + k$*. Let us define the function* $f_E : \{a_0, \ldots, a_k\} \to \mathbb{N}$ *such that, letting* $f_E(a_0) := c'$*, for all* $1 \leq i \leq k$*:*

$$f_E(a_i) := \begin{cases} f_E(a_{i-1}) & \text{if } a_{i-1} \equiv a_i \mod 2 \\ f_E(a_{i-1}) + 1 & \text{otherwise} \end{cases}$$

Since for all $1 \leq i \leq k$*, we have* $f_E(a_i) \leq f_E(a_{i-1}) + 1$*, we have* $f_E(a_k) \leq c' + k$*. We set* $e' := f_E(a_k)$*. Note also that, for all* $1 \leq i \leq k$*, we have* $f(a_k) \geq c'$*. With this choice, we have* $\mathsf{Sz}(E') = e' - c' \leq k \leq \tilde{e} - 1 - c = \mathsf{Sz}_{\mathsf{A}}(E) - 1 = n - 1$*.*

We can now define the function p'*. We let* $u := \chi_{\mathcal{G}_Y}(q_{\mathsf{init}})$ *for* $Y := (\mathcal{F}, E)$*. Then, for all* $o \in \mathsf{O}$*, we let:*

$$p'(o) := \begin{cases} p(o) \in [0,1] & \text{if } p(o) \in [0,1] \\ u \in [0,1] & \text{if } p(o) = \tilde{e} \\ k_{f_E(\max(c,n))} \in K_{e'} & \text{if } p(o) = k_n \in K_{\tilde{e}-1} \end{cases}$$

Let us now show a useful property about the function f_E from Definition 9.10.

Lemma 9.15. *Consider an (arbitrary) environment* $E = \langle c, e, p \rangle \in \mathsf{Env}(\mathsf{D})$ *and the function* $f_E : \{a_0, \ldots, a_k\} \to \mathbb{N}$ *from Definition 9.10. It ensures the following: for all* $0 \leq i \leq k$*,* a_i *and* $f_E(a_i)$ *have the same parity. In addition, for all* $X \subseteq \{a_0, \ldots, a_k\}$*,* $\max X$ *has the same parity as* $\max f_E(X)$*.*

Proof. Let us show by induction on $i \in [\![0, k]\!]$ the following property $\mathcal{P}(i)$: a_i and $f(a_i)$ have the same parity. The property $\mathcal{P}(0)$ straightforwardly holds since $a_0 = c$ and $f_E(a_0) = c'$ have the same parity. Assume now that $\mathcal{P}(i-1)$ holds for some $1 \leq i \leq k$. If a_i has the same parity as a_{i-1}, then $f_E(a_i) = f_E(a_{i-1})$ which has the same parity as a_{i-1}. Similarly, if a_i does not have the same parity as a_{i-1}, then $f_E(a_i) = f_E(a_{i-1}) + 1$. Hence, a_i and $f(a_i)$ have the same parity. Overall, the property $\mathcal{P}(i)$ holds for all $i \in [\![0, k]\!]$. The second part of the lemma comes from the fact that f_E is non-decreasing. □

Let us now show that the value of local parity games with an arbitrary environment E is equal to the value of the parity game with the corresponding relevant environment.

Lemma 9.16. *Consider an (arbitrary) environment* $E = \langle c, e, p \rangle \in \mathsf{Env}(\mathsf{O})$ *and the relevant environment* $E' = \langle c', e', p' \rangle$ *from Definition 9.10. Let* $Y := (\mathcal{F}, E)$ *and* $Y' := (\mathcal{F}, E')$*. Then,* $\chi_{\mathcal{G}_Y}(q_{\mathsf{init}}) \leq \chi_{\mathcal{G}_{Y'}}(q_{\mathsf{init}})$*.*

Proof. We let Q_y and $q_Y \in Q_Y$ (resp. $Q_{Y'}$ and $q_{Y'}$) denote the set of non-stopping states in $\mathcal{C}_Y$ and the state $q_{\mathsf{init}} \in Q_Y$ (resp. $\mathcal{C}_{Y'}$ and the state $q_{\mathsf{init}} \in Q_{Y'}$).

Recall the function f_E from Definition 9.10. We let $f := f_E$ and we also let $f(\tilde{e}) := d$ with d of the same parity than e such that $d \geq e' + 1$. Then, we define an alternate game $\mathcal{G}_Y^f = \langle \mathcal{C}_Y^f, \mathsf{Parity}_{[\![0,d]\!]} \rangle$ which differs from the game $\mathcal{G}_Y$ only in the coloring function and in the objective. The coloring function col^f in the arena $\mathcal{C}_Y^f$ is such that, for all states q in $\mathcal{C}_Y^f$, we have:

$$\mathsf{col}^f(q) := f(\max(c, \mathsf{col}(q)))$$

where col refers to the coloring function in the arena $\mathcal{C}_Y$. We let $W_Y := (\mathsf{col}^\omega)^{-1}[\mathsf{Parity}_{0,e}] \subseteq (Q_Y)^\omega$ (resp. $W_Y^f := ((\mathsf{col}^f)^\omega)^{-1}[\mathsf{Parity}_{0,d}] \subseteq (Q_Y^f)^\omega$, $W_{Y'} := ((\mathsf{col}')^\omega)^{-1}[\mathsf{Parity}_{0,e'}] \subseteq (Q_{Y'})^\omega$) be the Player-$\mathsf{A}$ winning set of infinite paths in $\mathcal{G}_Y$ (resp. $\mathcal{G}_Y^f$, $\mathcal{G}_{Y'}$). Let also $g_Y := (\mathsf{Parity}_{[\![0,e]\!]})_{\mathcal{C}_Y}$, $g_Y^f := (\mathsf{Parity}_{[\![0,d]\!]})_{\mathcal{C}_Y^f}$ and $g_{Y'} := (\mathsf{Parity}_{[\![0,e']\!]})_{\mathcal{C}_{Y'}}$.

Let us show that the value of both games $\mathcal{G}_Y^f$ and $\mathcal{G}^Y$ is the same. Consider a pair of strategies $(\mathsf{s_A}, \mathsf{s_B}) \in \mathsf{S}_\mathsf{A}^{\mathcal{C}_Y} \times \mathsf{S}_\mathsf{B}^{\mathcal{C}_Y} = \mathsf{S}_\mathsf{A}^{\mathcal{C}_Y^f} \times \mathsf{S}_\mathsf{B}^{\mathcal{C}_Y^f}$. First, we have:

$$\sum_{x \in [0,1]} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_Y, q_Y}[Q_Y^* \cdot x] \cdot x = \sum_{x \in [0,1]} \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_Y^f, q_Y}[Q_Y^* \cdot x] \cdot x$$

In addition, since all states k_i loop back to q_Y in $\mathcal{C}_Y$ and $\mathcal{C}_Y^f$, we have:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_Y, q_Y}[Q_Y^\omega \setminus (Q_Y^* \cdot q_Y)^\omega] = 0 = \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_Y^f, q_Y}[(Q_Y^f)^\omega \setminus ((Q_Y^f)^* \cdot q_Y)^\omega]$$

Furthermore, consider some infinite path $\rho \in (Q_Y^* \cdot q_Y)^\omega$ visiting infinitely often the central state q_Y. Let $X_\rho := \{\mathsf{col}[\mathsf{InfOft}(\rho)] \cap [\![c, e]\!]\}$ denote the set of colors, at least c, seen infinitely often in ρ in $\mathcal{G}_Y$. Note that $X_\rho \neq \emptyset$ since $\mathsf{col}(q_Y) = c$. We also let $X_\rho^f := \{\mathsf{col}^f[\mathsf{InfOft}(\rho)] \cap [\![c', d]\!]\}$ denote the set of colors, at least c', seen infinitely often in ρ in $\mathcal{G}_Y^f$. By definition of col^f, we have $X_\rho^f = f[X_\rho]$. Therefore, we have, by Lemma 9.15 and by definition of d:

$$\begin{aligned} \rho \in W_Y &\Leftrightarrow \max X_\rho \text{ is even} \\ &\Leftrightarrow \max f(X_\rho) \text{ is even} \\ &\Leftrightarrow \max X_\rho^f \text{ is even} \\ &\Leftrightarrow \rho \in W_Y^f \end{aligned}$$

It follows that:

$$\mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_Y, q_Y}[Q_Y^\omega \cap W_Y] = \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_Y^f, q_Y}[(Q_Y^f)^\omega \cap W_Y^f]$$

Overall, we obtain:

$$\mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_Y, q_Y}[g_Y] = \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_Y^f, q_Y}[g_Y^f]$$

As this holds for all pair of strategies $(\mathsf{s_A}, \mathsf{s_B}) \in \mathsf{S}_\mathsf{A}^{\mathcal{C}_Y} \times \mathsf{S}_\mathsf{B}^{\mathcal{C}_Y} = \mathsf{S}_\mathsf{A}^{\mathcal{C}_Y^f} \times \mathsf{S}_\mathsf{B}^{\mathcal{C}_Y^f}$, it follows that $\chi_{\mathcal{G}_Y}(q_Y) = \chi_{\mathcal{G}_Y^f}(q_Y)$.

Let us now relate the games $\mathcal{G}_Y^f$ and $\mathcal{G}_{Y'}$ by using Theorem 2.3 (item 1.b). First, we denote by col' the coloring function in the arena $\mathcal{C}_{Y'}$ Then, by Theorem 2.3 (item 1.b), almost-optimal strategies in $\mathcal{G}_Y^f$ and $\mathcal{G}_{Y'}$ can be found among $(\llbracket c', d \rrbracket, \mathsf{col}^f)$-uniform strategies in $\mathcal{G}_Y^f$ and among $(\llbracket c', e' \rrbracket, \mathsf{col}')$-uniform strategies in $\mathcal{G}_{Y'}$, with $d > e'$. Therefore, any strategy $(\llbracket c', d \rrbracket, \mathsf{col}^f)$-uniform strategy in $\mathcal{C}_Y^f$ can be seen as a strategy in $\mathcal{C}_{Y'}$. However, a $(\llbracket c', e' \rrbracket, \mathsf{col}^f)$-uniform strategy in $\mathcal{C}_{Y'}$ can be seen as a strategy in $\mathcal{C}_Y^f$ up to defining what it does once a state $(k_{\tilde{e}})$ colored with d occurs.

Consider any Player-A strategy $\mathsf{s}_\mathsf{A} \in \mathsf{S}_\mathsf{A}^{\mathcal{C}_Y^f}$ that is $(\llbracket c', d \rrbracket, \mathsf{col}^f)$-uniform in $\mathcal{C}_Y^f$. As mentioned above, it can be seen as a Player-A strategy in $\mathcal{C}_{Y'}$. Let $\delta > 0$ and consider a $(\llbracket c', e' \rrbracket, \mathsf{col}')$-uniform Player-$\mathsf{B}$ strategy in $\mathcal{C}_{Y'}$. Consider the Player-B strategy $\mathsf{s}_\mathsf{B}^\delta \in \mathsf{S}_\mathsf{B}^{\mathcal{C}_Y^f}$ that mimics the strategy s_B as long as $k_{\tilde{e}}$ is not seen. Once it is seen, the Player-B strategy $\mathsf{s}_\mathsf{B}^\delta$ switches to a δ-optimal strategy in the game $\mathcal{G}_Y^f$. By definition of $\mathsf{s}_\mathsf{B}^\delta$, and since $u = \chi_{\mathcal{G}_Y}(q_Y) = \chi_{\mathcal{G}_Y^f}(q_Y)$, we have:

$$\mathbb{E}_{\mathcal{C}_Y^f, q_Y}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}^\delta}[g_Y^f \cdot \mathbb{1}_{(Q_Y^f)^* \cdot k_{\tilde{e}}}] = \mathbb{P}_{\mathcal{C}_Y^f, q_Y}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}^\delta}[(Q_Y^f)^* \cdot k_{\tilde{e}}] \cdot (u + \delta)$$

Since the outcomes leading to $k_{\tilde{e}}$ in $\mathcal{C}_y^f$ lead to the stopping state $u \in [0, 1]$ in $\mathcal{C}_{Y'}$, and all outcomes leading to a stopping state in $\mathcal{C}_y^f$ lead to the same stopping state in $\mathcal{C}_{Y'}$, it follows that:

$$\mathbb{E}_{\mathcal{C}_Y^f, q_Y}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}^\delta}[f_Y \cdot \mathbb{1}_{(Q_Y)^* \cdot (k_{\tilde{e}} \cup [0,1])}] \leq \mathbb{E}_{\mathcal{C}_{Y'}, q_{Y'}}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}}[f_{Y'} \cdot \mathbb{1}_{(Q_{Y'})^* \cdot [0,1]}] + \delta$$

In addition, we have:

$$\mathbb{P}_{\mathcal{C}_Y^f, q_Y}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}^\delta}[(Q_Y \setminus \{k_{\tilde{e}}\})^\omega] = \mathbb{P}_{\mathcal{C}_{Y'}, q_{Y'}}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}}[(Q_{Y'})^\omega]$$

In addition, consider any outcome $o \in \mathsf{O}$ such that $p'(o) \notin [0, 1]$ or equivalently such that $p(o) \notin [0, 1] \cup \{k_{\tilde{e}}\}$. Then, we have $p'(o) = k_{f(\max(c,n))}$ where $n \in \llbracket 0, \tilde{e} - 1 \rrbracket$ is such that $p(o) = k_n$. Then, we have $\mathsf{col}' \circ p'(o) = f(\max(c, n)) = \mathsf{col}^f \circ p(o)$. Furthermore, we have $\mathsf{col}'(q_{Y'}) = c' = \mathsf{col}^f(q_Y)$. It follows that, assuming that $[0, 1] \cup \{k_{\tilde{e}}\}$ is not reached in $\mathcal{C}_Y^f$ and that $[0, 1]$ is not reached in $\mathcal{C}_{Y'}$, the colors seen with s_A and $\mathsf{s}_\mathsf{B}^\delta$ are the same than the colors seen with s_A and s_B in $\mathcal{C}_Y^f$ (because the strategies $\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}, \mathsf{s}_\mathsf{B}^\delta$ only depend on the colors seen). Therefore:

$$\mathbb{P}_{\mathcal{C}_Y^f, q_Y}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}^\delta}[W_Y^f \cap (Q_Y \setminus \{k_{\tilde{e}}\})^\omega] = \mathbb{P}_{\mathcal{C}_{Y'}, q_{Y'}}^{\mathsf{s}_\mathsf{A}, \mathsf{s}_\mathsf{B}}[W_{Y'} \cap (Q_{Y'})^\omega]$$

Overall, we obtain that:

$$\begin{aligned}
\chi_{\mathcal{G}_Y}[\mathsf{s_A}](q_Y) \leq \mathbb{E}^{\mathsf{s_A},\mathsf{s_B^\delta}}_{\mathcal{C}^f_Y,q_Y}[g^f_Y] &= \mathbb{E}^{\mathsf{s_A},\mathsf{s_B^\delta}}_{\mathcal{C}^f_Y,q_Y}[g^f_Y \cdot \mathbb{1}_{(Q_Y)^* \cdot (k_{\tilde{e}} \cup [0,1])}] \\
&+ \mathbb{P}^{\mathsf{s_A},\mathsf{s_B^\delta}}_{\mathcal{C}^f_Y,q_Y}[W^f_Y \cap (Q_Y \setminus ([0,1] \cup \{k_{\tilde{e}}\}))^\omega] \\
&\leq \mathbb{E}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_{Y'},q_{Y'}}[g_{Y'} \cdot \mathbb{1}_{(Q_{Y'})^* \cdot [0,1]}] + \delta \\
&+ \mathbb{P}^{\mathsf{s_A},\mathsf{s_B}}_{\mathcal{C}_{Y'},q_{Y'}}[W_{Y'} \cap (Q_{Y'})^\omega] \\
&= \mathbb{E}^{\mathsf{s_A},\mathsf{s_B^\delta}}_{\mathcal{C}_{Y'},q_{Y'}}[g_{Y'}]
\end{aligned}$$

Since this holds for all $\delta > 0$ and since Player-B $(\llbracket c', e' \rrbracket, \mathsf{col}')$-uniform strategies are enough to play almost optimally against the strategy $\mathsf{s_A}$, by Theorem 2.3 (item 1.b), it follows that $\chi_{\mathcal{G}^f_Y}[\mathsf{s_A}](q_Y) \leq \chi_{\mathcal{G}_{Y'}}[\mathsf{s_A}](q_{Y'}) \leq \chi_{\mathcal{G}_{Y'}}(q_{Y'})$. Since this holds for all Player-A $(\llbracket c, d \rrbracket, \mathsf{col})$-uniform strategies $\mathsf{s_A} \in \mathsf{S}^{\mathcal{C}_Y}_{\mathsf{A}}$, and since, by Theorem 2.3 (item 1.b), $(\llbracket c, d \rrbracket, \mathsf{col})$-uniform strategies achieve the value of the game $\mathcal{G}^f_Y$, it follows that $\chi_{\mathcal{G}_Y}(q_Y) = \chi_{\mathcal{G}^f_Y}(q_Y) \leq \chi_{\mathcal{G}_{Y'}}(q_{Y'})$. □

We can now proceed to the proof of Proposition 9.8.

Proof. Consider some $n \in \mathbb{N}$. First, consider the case where $n = 0$. Then, an environment $E = \langle c, e, p \rangle$ with $\mathsf{Sz_A}(E) = 0$ is such that $c = e$ is even. That is, the corresponding parity game $\mathcal{G}_{(\mathcal{F},E)}$ is, from Player A's point-of-view, a safety game. Hence, positional optimal strategies exists for Player A in $\mathcal{G}_{(\mathcal{F},E)}$ by Theorem 4.5. This is similar for Player B. Furthermore, there is no relevant environment of size at most -1. Hence, the equivalence holds for $n = 0$.

Assume now that $n \geq 1$. Consider any relevant environment E with $\mathsf{Sz}(E) \leq n - 1$. Then, we have $\mathsf{Sz_A}(E) \leq n$. Therefore, if $\mathcal{F}$ is positionally maximizable w.r.t. Player A up to n, then there is an optimal GF-strategy in $\mathcal{F}$ for Player A w.r.t. $(\mathcal{F}, E)$.

Consider now an (arbitrary) environment $E = \langle c, e, p \rangle \in \mathsf{Env}(\mathsf{O})$ with $\mathsf{Sz_A}(E) = n \geq 1$. Let $\tilde{e} := \mathsf{Even}(e)$. Clearly, the outcomes in O leading, w.r.t. p, to q_{init} of color c, can be redirected to k_c of color c (which then loops back to q_{init}) without changing the outcome of the parity game induced by E. Hence, without loss of generality, we assume that no outcome $o \in \mathsf{O}$ is such that $p(o) = q_{\mathsf{init}}$. We can therefore consider the relevant environment $E' \in \mathsf{Env}(\mathsf{O})$ from Definition 9.10 of size $\mathsf{Sz}(E') = n-1$. Let $Y := (\mathcal{F}, E)$ and $Y' := (\mathcal{F}, E')$ and assume that there is a Player-A GF-strategy $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathcal{F})$ that is optimal w.r.t. Y'. Let us show that σ_A is also optimal w.r.t. to Y. We let q_Y denote the state q_{init} in the game $\mathcal{G}_Y$ and $q_{Y'}$ denote the state q_{init} in the game $\mathcal{G}_{Y'}$. Let $u := \chi_{\mathcal{G}_Y}(q_Y)$ and $u' := \chi_{\mathcal{G}_{Y'}}(q_{Y'})$. By Lemma 9.16, we have $u \leq u'$. We want to apply Lemma 8.2.

Let us show that the Player-A positional strategy $\mathsf{s}^Y_\mathsf{A}(\sigma_\mathsf{A})$ ensures item (i) of this lemma, i.e. that it dominates the valuation v^u_Y in the game $\mathcal{G}_Y$. This amount to show that $\mathsf{val}[\langle \mathcal{F}, v_Y^u \circ p \rangle](\sigma_\mathsf{A}) \geq u$. Note that, again by Lemma 8.2, we have that $\mathsf{val}[\langle \mathcal{F}, v_{Y'}^{u'} \circ p' \rangle](\sigma_\mathsf{A}) \geq u'$. Furthermore, we have:

$$v_Y^u \circ p + (u' - u) \geq v_{Y'}^{u'} \circ p'$$

Therefore, by Lemma 1.10, for any Player-B GF-strategy $\sigma_\mathsf{B} \in \mathcal{D}(\mathsf{Act}_\mathsf{B})$, we have:

$$\begin{aligned} \mathsf{out}[\langle \mathcal{F}, v_Y^u \circ p \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) + u' - u &\geq \mathsf{out}[\langle \mathcal{F}, v_Y^{u'} \circ p' \rangle](\sigma_\mathsf{A}, \sigma_\mathsf{B}) \\ &\geq \mathsf{val}[\langle \mathcal{F}, v_Y^{u'} \circ p' \rangle](\sigma_\mathsf{A}) \geq u' \end{aligned}$$

Since this holds for all Player-B GF-strategies $\sigma_\mathsf{B} \in \Sigma_\mathsf{B}(\mathcal{F})$, it follows that

$$\mathsf{val}[\langle \mathcal{F}, v_Y^u \rangle](\sigma_\mathsf{A}) \geq u$$

Consider now item (ii) from Lemma 8.2. Consider some Player-B action $b \in \mathsf{Act}_\mathsf{B}$ and assume that $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[0,1]} \rangle](\sigma_\mathsf{A}, b) = 0$. If we have $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[k_{\tilde{e}}]} \rangle](\sigma_\mathsf{A}, b) > 0$, then it follows that $\max(\mathsf{Color}(\mathcal{F}, p, \sigma_\mathsf{A}, b) \cup \{c\})$ is even since $\tilde{e}$ is even and is the highest integer appearing the game $\mathcal{G}_Y$. Assume now that $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[k_{\tilde{e}}]} \rangle](\sigma_\mathsf{A}, b) = 0$. Then, it follows that, by definition of p', we have $\mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{(p')^{-1}[0,1]} \rangle](\sigma_\mathsf{A}, b) = 0$. Therefore, by Lemma 8.2, it follows that $\max(\mathsf{Color}(\mathcal{F}, p', \sigma_\mathsf{A}, b) \cup \{c'\})$ is even. Furthermore, by definition of p':

$$\begin{aligned} &f_E[\{i \in [\![c, \tilde{e} - 1]\!] \mid \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[k_i]} \rangle](\sigma_\mathsf{A}, b) > 0\} \cup \{c\}] \\ =&\{i \in [\![c', e']\!] \mid \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{(p')^{-1}[k_i]} \rangle](\sigma_\mathsf{A}, b) > 0\} \cup \{c'\} \end{aligned}$$

We have $\max\{i \in [\![c', e']\!] \mid \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{(p')^{-1}[k_i]} \rangle](\sigma_\mathsf{A}, b) > 0\} \cup \{c'\} = \max(\mathsf{Color}(\mathcal{F}, p', \sigma_\mathsf{A}, b) \cup \{c'\})$ is even. Therefore, by Lemma 9.15, we have $\max\{i \in [\![c, \tilde{e} - 1]\!] \mid \mathsf{out}[\langle \mathcal{F}, \mathbb{1}_{p^{-1}[k_i]} \rangle](\sigma_\mathsf{A}, b) > 0\} \cup \{c\} = \max(\mathsf{Color}(\mathcal{F}, p, \sigma_\mathsf{A}, b) \cup \{c\})$ is also even.

In fact, the Player-A GF-strategy $\sigma_\mathsf{A} \in \Sigma_\mathsf{A}(\mathcal{F})$ satisfies item (ii) of Lemma 8.2. Therefore, it is optimal w.r.t. Y. □

9.5.2 . Proof of Lemma 9.13

Proof. Consider a relevant environment $E = \langle c, e, p \rangle \in \mathsf{Env}(\mathsf{O}_n)$ and $Y_n := (\mathcal{F}_n, E)$. We let $u := \chi_{\mathcal{G}_{Y_n}}(q_\mathsf{init})$. Note that since E is relevant, we have $c \in \{0, 1\}$ and $p : \mathsf{O}_n \to [0, 1] \cup K_n$. For all $o \in \mathsf{O}$, if $p(o) \notin [0, 1]$, we let $n_o \in [\![0, e]\!]$ denote the integer that t is mapped to w.r.t. p, i.e. the integer ensuring $p(o) = k_{n_o}$. Note that, since E is relevant, for all $o \in \mathsf{O}_n$, we have $n_o \geq c$.

Let us introduce some notations. We let $X_{[0,1]} := \{o \in \mathsf{O}_n \mid p(o) \in [0, 1]\}$, $X_\mathsf{even} := \{o \in \mathsf{O}_n \mid n_o \text{ is even}\}$ and $X_\mathsf{odd} := p^{-1}[K_n] \setminus X_\mathsf{even}$. We let $X_\mathsf{Lp} := p^{-1}[K_n] = X_\mathsf{even} \uplus X_\mathsf{odd}$. Note that we have $\mathsf{O}_n = X_{[0,1]} \uplus X_\mathsf{Lp}$. We define the valuation $v : \mathsf{O}_n \to [0, 1]$ mapping each outcome to its value in the game $\mathcal{G}_{Y_n}$ such that, for all $o \in \mathsf{O}_n$:

$$v(o) := \begin{cases} p(o) & \text{if } o \in X_{[0,1]} \\ u & \text{otherwise} \end{cases}$$

$$\mathcal{T}_n = \begin{array}{c} \\ a_{\mathsf{t}} \\ a_{\mathsf{b}} \end{array} \begin{array}{c} \begin{array}{cc} b_{\mathsf{l}} & b_{\mathsf{r}} \end{array} \\ \left[\begin{array}{c:c} \frac{x_0 + x_1 + \ldots + x_{n-2} + x_{n-1}}{n} & \frac{3y+z}{4} \\ \hdashline \frac{3y+z}{4} & \frac{y+3z}{4} \end{array}\right] \end{array}$$

Figure 9.9: The game form $\mathcal{T}_n$.

That is, $v = v^u_{y_n} \circ p$. For all $x \in \mathsf{O}_n$, we let:

$$w(x) := \begin{cases} 1 & \text{if } x \in X_{\mathsf{even}} \\ 0 & \text{if } x \in X_{\mathsf{odd}} \\ p(x) & \text{if } x \in X_{[0,1]} \end{cases}$$

For all $x \in \mathsf{O}_n$, the value $w(x) \in [0,1]$ corresponds to the value of the game $\mathcal{G}_{Y_n}$ if x is seen (indefinitely, if it is in X_{Lp}). This can be generalized to a pair of outcomes: for all $x, x' \in \mathsf{O}_n$, we let:

$$w(x, x') := \begin{cases} w(x) & \text{if } x, x' \in X_{\mathsf{Lp}} \text{ and } n_x \leq n_{x'} \\ w(x') & \text{if } x, x' \in X_{\mathsf{Lp}} \text{ and } n_{x'} \preceq n_x \\ p(x) & \text{if } x \in X_{[0,1]}, x' \in X_{\mathsf{Lp}} \\ p(x') & \text{if } x' \in X_{[0,1]}, x \in X_{\mathsf{Lp}} \\ \frac{1}{2} \cdot (p(x) + p(x')) & \text{if } x, x' \in X_{[0,1]} \end{cases}$$

Finally, we let $\mathcal{T}_n$ denote the 2×2 game form at the top left of the game form $\mathcal{F}_n$ from Figure 9.6. It is depicted in Figure 9.9. We let $u' := \mathsf{val}[\langle \mathcal{T}_n, v \rangle]$. Now, there are several cases:

- Assume that $y, z \in X_{\mathsf{Lp}}$. Then, $u \in \{0, 1\}$, as in the proof of Proposition 9.11, and playing $a_{\mathsf{Ex}} \in \mathsf{Act}^n_{\mathsf{A}}$ (resp. $b_{\mathsf{Ex}} \in \mathsf{Act}^n_{\mathsf{B}}$) is optimal for Player A (resp. B) w.r.t. Y_n.

- Assume now that $y \in X_{\mathsf{Lp}}$ and $z \in X_{[0,1]}$. Then, $u = p(z)$, as argued in the proof sketch, and playing $a_{\mathsf{Ex}} \in \mathsf{Act}^n_{\mathsf{A}}$ (resp. $b_{\mathsf{Ex}} \in \mathsf{Act}^n_{\mathsf{B}}$) is optimal for Player A (resp. B) w.r.t. Y_n. This is similar if $y \in X_{[0,1]}$ and $z \in X_{\mathsf{Lp}}$.

- Assume that $p(y), p(z) \in [0,1]$ and $p(y) \leq p(z)$. In that case, we have $u = \frac{3p(y)+p(z)}{4}$ and playing action a_{Ex} (resp. b_{Ex}) is optimal for Player A (resp. B) w.r.t. Y_n.

- Let us now assume that $y, z \in X_{[0,1]}$ and $p(z) < p(y)$. In that case, we have:
$$\frac{3p(z) + p(y)}{4} \leq u \leq \frac{p(z) + 3p(y)}{4}$$
We argue that Player B always has an optimal GF-strategy w.r.t. Y_n.
  - If $u = \frac{p(z)+3p(y)}{4}$, then positionally playing action b_{Ex} is optimal for Player B. We now assume that $u < \frac{p(z)+3p(y)}{4}$.
  - If, for some even $i \in [\![0, n-1]\!]$, we have $w(x_i) \leq u$, then playing action b_i is optimal. We now assume that for all even $i \in [\![0, n-1]\!]$, we have $w(x_i) > u$ (i.e. x_i is mapped to either a real greater then u or even index). Furthermore, if for any even $i \in [\![2, n-1]\!]$, we have $w(x_i, x_{i-1}) \leq u$, then playing action $b_{i,i-1}$ is optimal. We now assume that for all even $i \in [\![2, n-1]\!]$, we have $w(x_i, x_{i-1}) > u$.
  - It follows that, for any odd $i \in [\![0, n-1]\!]$, $w(x_i) \leq u$ and $w(x_i, x_{i-1}) \leq u$ — otherwise Player A could ensure that the value of the game is more than u by playing the corresponding row, i.e. action a_i or $a_{i,i-1}$. It also implies that $u' \leq u$. Indeed, assume that it is not the case, i.e. $u' > u$. Consider a Player-A GF-strategy $\sigma_{\mathsf{A}} \in \mathcal{D}(\{a_{\mathsf{t}}, a_{\mathsf{b}}\})$ whose value in the game in normal form $\langle \mathcal{T}_n, v \rangle$ is greater than u. Then, the Player-A positional strategy s_{A} in the game $\mathcal{G}_{Y_n}$ that plays the GF-strategy σ_{A} in q_{init} parity dominates the valuation $v^r_{(\mathcal{F}_n, E)}$ for some $r > u$, due to the assumptions we made in the previous item. In fact, $u' \leq u$.

    Note that this implies $X_{\mathsf{Ex}} \neq \emptyset$. Indeed, assume that $X_{\mathsf{Ex}} = \emptyset$. Let $\varepsilon > 0$. Consider a Player-A GF-strategy σ_{A} in the game form $\mathcal{T}_n$ that plays action a_{t} with probability $1 - \varepsilon$ and action a_{b} with probability ε. Then, for some small enough, yet positive, $\varepsilon > 0$ and since $\frac{p(z)+3p(y)}{4} > u$ and $v(x_0), \ldots, v(x_{n-1}) = u$, such a GF-strategy has value more than u in the game in normal form $\langle \mathcal{T}_n, v \rangle$. Hence, $X_{\mathsf{Ex}} = \emptyset$.
  - Let us exhibit a Player-B GF-strategy that is optimal w.r.t. Y_n. Consider a Player-B GF-strategy $\sigma_{\mathsf{B}} \in \mathcal{D}(\{b_{\mathsf{l}}, b_{\mathsf{r}}\})$ that is optimal in that game in normal form $\langle \mathcal{T}_n, v \rangle$. Let s_{B} be the Player-B positional strategy in the game $\mathcal{G}_{Y_n}$ that plays the GF-strategy σ_{B} at q_{init}. Recall that we assume that for any odd $i \in [\![0, n-1]\!]$, $w(x_i) \leq u$ and $w(x_i, x_{i-1}) \leq u$, $\frac{3p(z)+p(y)}{4} \leq u$ and $X_{\mathsf{Ex}} \neq \emptyset$. Therefore, the strategy s_{B} dominates the valuations $v^u_{Y_n}$ and no Player-A action can, against the strategy s_{B}, make the game loop indefinitely on q_{init} while ensuring that the highest color seen is even, almost-surely. In fact, this strategy s_{B} parity dominates the valuations

$v^u_{Y_n}$. Hence, the Player-B GF-strategy σ_B is optimal w.r.t. (Y_n), recall Lemma 8.2.

Let us now consider the case of Player A.

- If $u = \frac{3p(z)+p(y)}{4}$, then playing action a_{Ex} is optimal. We now assume that $\frac{3p(z)+p(y)}{4} < u$.
- It follows that, for all even $i \in [\![0, n-1]\!]$, we have $w(x_i) \geq u$ and for all even $i \in [\![2, k-1]\!]$, we have $w(x_{i,i-1}) \geq u$. Otherwise Player B could play the action b_i or $b_{i,i-1}$ and ensure that the value of the game $\mathcal{G}_{Y_n}$, from q_{init}, is less than u.
- Furthermore, for any odd $i \in [\![1, n-1]\!]$, if either $w(x_i) \geq u$ or $w(x_i, x_{i-1}) \geq u$, then playing action a_i or action $a_{i,i-1}$ is optimal. We now assume that, for all odd $i \in [\![A, k-1]\!]$, $w(x_i) < u$ and $w(x_i, x_{i-1}) < u$.
- Assume that $X_{\mathsf{Ex}} \neq \emptyset$. Then, we have $u' \geq u$. This is analogous to the previous case for Player B. Indeed, assume that $u' < u$. Then, consider a Player-B GF-strategy $\sigma_\mathsf{B} \in \mathcal{D}(\{b_\mathsf{l}, b_\mathsf{r}\})$ whose value in the game in normal form $\langle \mathcal{T}_n, v \rangle$ is less than u. Then, a Player-B positional strategy s_B in the game $\mathcal{G}_{Y_n}$ that plays the GF-strategy σ_B in q_{init} parity dominates the valuation $v^r_{Y_n}$ for some $r < u$, due to the assumptions we made in the previous item and since we assume that $\frac{3p(z)+p(y)}{4} < u$. In fact, $u' \geq u$.

  Therefore, in the case where $X_{\mathsf{Ex}} \neq \emptyset$, Player A has an optimal GF-strategy w.r.t. Y_n that consists in playing optimally (in $\mathcal{D}(a_\mathsf{t}, a_\mathsf{b})$) in the game in normal form $\mathsf{val}[\langle \mathcal{T}_n, v \rangle]$, due to the assumptions of the previous item.
- Let us now assume that $X_{\mathsf{Ex}} = \emptyset$. Since, for all even $i \in [\![0, n-1]\!]$, we have $w(x_i) \geq u$, it follows that n_{x_i} is even and at least equal to c. Similarly, since for all odd $i \in [\![1, n-1]\!]$, we have $w(x_i) < u$. That is, m_{x_i} is odd or less than c. In addition, for all even $i \in [\![2, n-1]\!]$, we have $w(x_{i,i-1}) \geq u$. That is, $n_{x_i} > n_{x_{i-1}}$. Similarly, for all odd $i \in [\![1, n-1]\!]$, we have $w(x_i, x_{i-1}) < u$, it follows that $n_{x_i} > n(x_{i-1})$. We have $c \leq n_{x_0} < n_{x_1} < \ldots < n_{x_{n-2}} < n_{x_{n-1}} \leq e$. This cannot happen if $e - c \leq n - 2$.

Let us now exhibit a relevant environment E of size $n-1$ w.r.t. which Player A has no optimal strategy. Consider the environment where $c := 0$, $e := n-1$, $p(z) := 0$, $p(y) := 1$ and for all $i \in [\![0, n-1]\!]$, we have $p(x_i) := k_i$. We do have $\mathsf{Sz}(E) = n-1$. Furthermore, the value u of the parity game $\mathcal{G}_{Y_n}$ from q_{init} is equal to $\frac{p(z)+3p(y)}{4} = \frac{3}{4}$. Indeed, Player B ensures that $u \leq \frac{3}{4}$ by playing action

b_{Ex}. Furthermore, for all $\varepsilon > 0$, Player A can play with probability $1-\varepsilon$ action a_{t} and with probability ε action a_{b}, thus ensuring that $u \geq \frac{3}{4} - \varepsilon$.

Consider now an arbitrary Player-A GF-strategy σ_{A} and the corresponding positional strategy s_{A} in the game $\mathcal{G}_{Y_n}$. If σ_{A} plays with positive probability action a_{b} or action a_{Ex}, then the value of s_{A} is less that u since $\frac{p(y)+3p(z)}{4} < \frac{p(z)+3p(y)}{4}$ (Player B can play the action b_{r}). However, if σ_{A} does not play these actions with positive probability, then consider the Player B strategy s_{B} that plays action b_{l} with probability 1. It ensures that the value of the strategy s_{A} is 0. The reason why is the following: the highest index that the variables $x_0, x_1, \ldots, x_{n-1}$ are mapped to w.r.t. p is $n-1$ and it is odd. Furthermore, all variables $x_1, x_3, \ldots, x_{n-1}$ are mapped to odd indices. Finally, the highest index that the variables (x_1, x_0) are mapped to is 1 and it is odd. This is also the case for (x_3, x_2), ...,(x_{n-1}, x_{n-2}). Hence, the highest color seen infinitely often with s_{A} and s_{B} is almost-surely odd, and surely no stopping state is reached.. $\square$

Conclusion

The goal of this dissertation was to give significant insight on how concurrent games behave. Our first contribution is conceptual: we introduced the notion of arbitrary game forms, which generalize standard game forms. Usually, in the literature, standard game forms are the underlying structure of the (local) interactions of the players in concurrent games. A priori, using arbitrary game forms instead of standard game forms as local interactions makes concurrent games harder to handle. However, we have provided in this dissertation several new results that hold even with non-standard local interactions. The most notable example of this fact is the new version of Blackwell determinacy (i.e. Theorem 2.2). However, Theorems 3.1, 3.12 and 3.17, which we believe are also important results on concurrent games, are also relevant examples. Remarkably, the proofs of the above-mentioned results are not made any harder by the fact that non-standard game forms may appear as local interactions. By contrast, the proof of Theorem 4.11, and especially the intermediate definitions used to prove it, are more intricate when considering arbitrary game forms than when considering standard finite game forms. This shows that, although manipulating arbitrary game forms may prove tricky, significant results can still be established on concurrent games when using them as local interactions.

We would now like to highlight some of the aspects we believe are the most important takeaways from this dissertation. First, let us discuss the generalization of Blackwell determinacy, stated in Theorem 2.3. Interestingly, both the original result on the determinacy of Blackwell games [12] and Borel determinacy [8] are logical consequences of this generalization of Blackwell determinacy. Furthermore, as mentioned above, another benefit of this generalization is that it holds even with arbitrary game forms. In addition, as discussed in Chapter 2, this generalization extends Martin's determinacy of Blackwell in two directions (items 1.a and 1.b in Chapter 2). We believe that both of these directions are interesting; we have provided in this dissertation applications of both of them. Notably, these applications have been established with the help of the generalization stated as Theorem 2.3 and, a priori, the original result on Blackwell determinacy by Martin would not have been enough to establish them.

Second, in Chapter 3, we have introduced the notion of finite-choice strategies (Definition 3.22). Recall, these are strategies that, at every state of a game, play only finitely many different GF-strategies. We believe that this notion is very useful when studying concurrent parity games, especially standard finite ones. For instance, in such a setting with a parity objective, as stated in Corollary 3.38, any value achieved by a finite-choice strategy can also be achieved by a positional one. Said otherwise, if a positional strategy cannot achieve a value, then no finite-choice (and in particular, no finite-memory) strategy can achieve it either. In fact, we have extensively used Corollary 3.38 in Part II to establish that infinite choice is required to achieve some optimal values in several parity games.

Third, although we have already done it in the introduction and in the beginning of Part III, we would like to highlight once more the local-global transfers that we established in Part III. Recall that our approach consisted in defining classes of concurrent arenas by restricting the game forms that could be used as local interactions. These sets of game forms are defined such that the arenas built from them ensure some desirable properties. We have established several NSC-transfers, which all correspond to underlying necessary and sufficient conditions, as mentioned in page 248. Also, we have established merely sufficient conditions on game forms for the arenas built on them to behave in a desirable way. We believe that the various results that we have established in Part III show that this approach is viable to build concurrent arenas that are well-behaved by design.

Future leads

We have already discussed at the end of each chapter (except Chapter 1) natural future work and open questions on the technical contents of the chapters. To conclude this dissertation, we would like to mention possible research prospects beyond what has been studied in this dissertation.

More objectives. First, when dealing with specific win/lose objectives, in this dissertation, we have only considered parity objectives. A very natural research lead would consist in exploring concurrent games with other kinds of win/lose objectives. The main characteristics of parity objectives are twofold. First, these objectives are prefix-independent. Second, they are qualitative in the sense that what matters is what colors are seen infinitely often regardless of how often these colors are seen. We believe that the most interesting questions would arise when considering quantitative objectives, i.e objectives where it matters how often colors occur. In this case, we believe that the easiest objectives to consider would be related to discounted sum w.r.t. a discount factor smaller than 1, either as payoff functions or win/lose objectives defined by threshold. The reason why studying these specific objectives should be trackable is because the discounted sum payoff functions are upper semi-continuous (recall Definition 4.2). In particular, one could apply Proposition 4.1.

We believe that more challenging questions would arise when studying mean-payoff objectives. Contrary to discounted sum, mean-payoff objectives are prefix-independent. Therefore, several general results that we have shown in this dissertations can be applied to mean-payoff objectives, such as Theorem 3.12 or Theorem 3.37. We believe that the most interesting question would be to look for a similar result than the one stated in Theorem 8.3 in the case of mean-payoff objectives. However, the fact that these objectives are quantitative instead of qualitative seems a tricky issue to handle.

Non-antagonistic games. In this dissertation, we have only considered antagonistic objectives for the two players involved in the game. It could be interesting to allow for non-antagonistic player preferences over the traces, i.e. infinite sequence of colors. In such a setting, we would not consider (subgame) optimal strategies anymore but rather (subgame perfect) Nash equilibria. Not many results are known in this framework. For instance, consider a finite concurrent arena where both players have (non-antagonistic) reachability objectives. In this rather simple context, it is not known whether there always exist (ε-)Nash equilibria. This is partly discussed in the second paragraph of the second page of [81]. We believe that the notion of finite-choice strategies may prove useful also in this setting.

Finally, one could also be interested in investigating even more involved questions, related to games with more than two players. Let us mention that, in the early times of this PhD, we have explored questions related to Nash equilibria with deterministic strategies in multi-player games (with arbitrarily many payers). The purpose of this work was to extend in a concurrent setting what is done in [82]. This is unpublished, but opens the way to various interesting research directions.

Main results

Bibliography

[1] P. Vandenhove, Strategy complexity of zero-sum games on graphs. (Complexité des stratégies des jeux sur graphes à somme nulle). PhD thesis, University of Mons, Belgium, 2023.

[2] M. Randour, Synthesis in Multi-Criteria Quantitative Games. PhD thesis, Université de Mons, Belgium, 2014.

[3] J. v. Neumann, "Zur theorie der gesellschaftsspiele," Mathematische annalen, vol. 100, no. 1, pp. 295–320, 1928.

[4] J. von Neumann and O. Morgenstern, Theory of Games and Economic Behavior. Princeton: Princeton Univ. Press, 1944.

[5] J. F. Nash Jr, "Equilibrium points in n-person games," Proceedings of the national academy of sciences, vol. 36, no. 1, pp. 48–49, 1950.

[6] J. Nash, "Non-cooperative games," Annals of mathematics, pp. 286–295, 1951.

[7] D. Gale and F. M. Stewart, 13. Infinite Games with Perfect Information, pp. 245–266. Princeton: Princeton University Press, 1953.

[8] D. A. Martin, "Borel determinacy," Annals of Mathematics, vol. 102, pp. 363–371, 1975.

[9] L. S. Shapley, "Stochastic games," Proceedings of the national academy of sciences, vol. 39, no. 10, pp. 1095–1100, 1953.

[10] H. Everett, "Recursive games," Annals of Mathematics Studies – Contributions to the Theory of Games, vol. 3, pp. 67–78, 1957.

[11] D. Blackwell, "Infinite g_δ-games with imperfect information," Applicationes Mathematicae, vol. 1, no. 10, pp. 99–101, 1969.

[12] D. A. Martin, "The determinacy of Blackwell games," The Journal of Symbolic Logic, vol. 63, no. 4, pp. 1565–1581, 1998.

[13] J. M. Smith and G. R. Price, "The logic of animal conflict," Nature, vol. 246, no. 5427, pp. 15–18, 1973.

[14] A. Ehrenfeucht, "An application of games to the completeness problem for formalized theories," Fund. Math, vol. 49, no. 129-141, p. 13, 1961.

[15] L. H. Landweber, "Finite state games-a solvability algorithm for restricted second-order arithmetic," Notices Amer. Math. Soc, vol. 14, pp. 129–130, 1967.

[16] J. R. Büchi and L. H. Landweber, "Solving sequential conditions by finite-state strategies," Transactions of the American Mathematical Society, vol. 138, pp. 295–311, 1969.

[17] J. R. Büchi, "Using determinancy of games to eliminate quantifiers," in International Symposium on Fundamentals of Computation Theory, 1977.

[18] Y. Gurevich and L. Harrington, "Trees, automata, and games," in Proceedings of the fourteenth annual ACM symposium on Theory of computing, pp. 60–65, 1982.

[19] D. Harel and A. Pnueli, "On the development of reactive systems," in Logics and Models of Concurrent Systems - Conference proceedings, Colle-sur-Loup (near Nice), France, 8-19 October 1984 (K. R. Apt, ed.), vol. 13 of NATO ASI Series, pp. 477–498, Springer, 1984.

[20] E. M. Clarke and E. A. Emerson, "Design and synthesis of synchronization skeletons using branching-time temporal logic," in Logics of Programs, Workshop, Yorktown Heights, New York, USA, May 1981 (D. Kozen, ed.), vol. 131 of Lecture Notes in Computer Science, pp. 52–71, Springer, 1981.

[21] M. Y. Vardi and P. Wolper, "An automata-theoretic approach to automatic program verification (preliminary report)," in Proceedings of the Symposium on Logic in Computer Science (LICS '86), Cambridge, Massachusetts, USA, June 16-18, 1986, pp. 332–344, IEEE Computer Society, 1986.

[22] W. Thomas, "On the synthesis of strategies in infinite games," in Annual Symposium on Theoretical Aspects of Computer Science, pp. 1–13, Springer, 1995.

[23] W. Thomas, Handbook of Formal Languages, Volume 3: Beyond Words. Springer, 1997.

[24] Z. Manna and A. Pnueli, The temporal logic of reactive and concurrent systems - specification. Springer, 1992.

[25] K. R. Apt and E. Grädel, eds., Lectures in Game Theory for Computer Scientists. Cambridge University Press, 2011.

[26] N. Fijalkow, N. Bertrand, P. Bouyer-Decitre, R. Brenguier, A. Carayol, J. Fearnley, H. Gimbert, F. Horn, R. Ibsen-Jensen, N. Markey, B. Monmege, P. Novotný, M. Randour, O. Sankur, S. Schmitz, O. Serre, and M. Skomra, "Games on graphs," CoRR, vol. abs/2305.10546, 2023.

[27] K. Chatterjee, M. Jurdzinski, and T. A. Henzinger, "Quantitative stochastic parity games," in Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004, New Orleans, Louisiana, USA, January 11-14, 2004, pp. 121–130, SIAM, 2004.

[28] B. Karelovic and W. Zielonka, "Nearest fixed points and concurrent priority games," in Fundamentals of Computation Theory - 20th International Symposium, FCT 2015, Gdańsk, Poland, August 17-19, 2015, Proceedings (A. Kosowski and I. Walukiewicz, eds.), vol. 9210 of Lecture Notes in Computer Science, pp. 381–393, Springer, 2015.

[29] L. de Alfaro, T. A. Henzinger, and F. Y. C. Mang, "The control of synchronous systems," in CONCUR 2000 - Concurrency Theory, 11th International Conference, University Park, PA, USA, August 22-25, 2000, Proceedings (C. Palamidessi, ed.), vol. 1877 of Lecture Notes in Computer Science, pp. 458–473, Springer, 2000.

[30] L. de Alfaro, T. A. Henzinger, and F. Y. C. Mang, "The control of synchronous systems, part II," in CONCUR 2001 - Concurrency Theory, 12th International Conference, Aalborg, Denmark, August 20-25, 2001, Proceedings (K. G. Larsen and M. Nielsen, eds.), vol. 2154 of Lecture Notes in Computer Science, pp. 566–582, Springer, 2001.

[31] L. de Alfaro, T. Henzinger, and O. Kupferman, "Concurrent reachability games," Theoretical Computer Science, vol. 386, no. 3, pp. 188–217, 2007.

[32] L. de Alfaro and R. Majumdar, "Quantitative solution of omega-regular games," Journal of Computer and System Sciences, vol. 68, pp. 374–397, 2004.

[33] S. K. S. Frederiksen, Semi-algebraic tools for stochastic games. PhD thesis, PhD thesis, Aarhus Universitet, Aarhus, Denmark, 7 2015, 2015.

[34] M. Kwiatkowska, G. Norman, D. Parker, G. Santos, and R. Yan, "Probabilistic model checking for strategic equilibria-based decision making," LIPIcs, Leibniz-Zentrum für Informatik, 2022. To appear; Available as arXiv:2206.15148.

[35] A. Gibbard, "Manipulation of voting schemes: a general result," Econometrica: journal of the Econometric Society, pp. 587–601, 1973.

[36] V. Gurvich, "The solvability of positional games in pure strategies," USSR Computational Mathematics and Mathematical Physics, vol. 15, no. 2, pp. 74–87, 1975.

[37] A. P. Maitra and W. D. Sudderth, Discrete gambling and stochastic games, vol. 32. Springer Science & Business Media, 2012.

[38] B. Bordais, P. Bouyer, and S. Le Roux, "From local to global determinacy in concurrent graph games," in 41st IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2021, December 15-17, 2021, Virtual Conference (M. Bojanczyk and C. Chekuri, eds.), vol. 213 of LIPIcs, pp. 41:1–41:14, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021.

[39] B. Bordais, P. Bouyer, and S. Le Roux, "Optimal strategies in concurrent reachability games," in 30th EACSL Annual Conference on Computer Science Logic, CSL 2022, February 14-19, 2022, Göttingen, Germany (Virtual Conference) (F. Manea and A. Simpson, eds.), vol. 216 of LIPIcs, pp. 7:1–7:17, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[40] B. Bordais, P. Bouyer, and S. Le Roux, "Playing (almost-)optimally in concurrent büchi and co-büchi games," in 42nd IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science, FSTTCS 2022, December 18-20, 2022, IIT Madras, Chennai, India (A. Dawar and V. Guruswami, eds.), vol. 250 of LIPIcs, pp. 33:1–33:18, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022.

[41] B. Bordais, P. Bouyer, and S. L. Roux, "Subgame optimal strategies in finite concurrent games with prefix-independent objectives," in Foundations of Software Science and Computation Structures - 26th International Conference, FoSSaCS 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2023, Paris, France, April 22-27, 2023, Proceedings (O. Kupferman and P. Sobocinski, eds.), vol. 13992 of Lecture Notes in Computer Science, pp. 541–560, Springer, 2023.

[42] M. Löwe, "Theory of measure and integration." Available On Webpage. Consulted 13/02/23 at https://www.uni-muenster.de/Stochastik/loewe/measure.pdf.

[43] C. Baier and J.-P. Katoen, Principles of Model Checking. MIT Press, 2008. http://mitpress.mit.edu/catalog/item/default.asp?ttype=2&tid=11481.

[44] J. von Neumann and O. Morgenstern, Theory of Games and Economic Behavior. Princeton University Press, 1944.

[45] C. Calabro, “Extension of minimax to infinite matrices.” Available On Webpage, 2004. Consulted 13/02/23 at https://cseweb.ucsd.edu/~ccalabro/essays/minimax.pdf.

[46] M. Sion, “On general minimax theorems.,” 1958.

[47] L. de Alfaro and T. A. Henzinger, “Concurrent omega-regular games,” in 15th Annual IEEE Symposium on Logic in Computer Science, Santa Barbara, California, USA, June 26-29, 2000, pp. 141–154, IEEE Computer Society, 2000.

[48] S. Kiefer, R. Mayr, M. Shirmohammadi, and P. Totzke, “Strategy complexity of reachability in countable stochastic 2-player games,” CoRR, vol. abs/2203.12024, 2022.

[49] K. Chatterjee, “Concurrent games with tail objectives,” Theor. Comput. Sci., vol. 388, no. 1-3, pp. 181–198, 2007.

[50] K. Chatterjee, L. de Alfaro, and T. A. Henzinger, “The complexity of quantitative concurrent parity games,” in Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006, pp. 678–687, ACM Press, 2006.

[51] P. Bouyer, S. Le Roux, Y. Oualhadj, M. Randour, and P. Vandenhove, “Games where you can play optimally with arena-independent finite memory,” in 31st International Conference on Concurrency Theory, CONCUR 2020, September 1-4, 2020, Vienna, Austria (Virtual Conference), pp. 24:1–24:22, 2020.

[52] L. de Alfaro, Formal Verification of Probabilistic Systems. PhD thesis, Stanford University, 1997.

[53] S. Le Roux, “Time-aware uniformization of winning strategies,” in Beyond the Horizon of Computability - 16th Conference on Computability in Europe, CiE 2020, Fisciano, Italy, June 29 - July 3, 2020, Proceedings (M. Anselmo, G. D. Vedova, F. Manea, and A. Pauly, eds.), vol. 12098 of Lecture Notes in Computer Science, pp. 193–204, Springer, 2020.

[54] J.-F. Le Gall, Brownian motion, martingales, and stochastic calculus. Springer, 2016.

[55] K. Chatterjee, L. Doyen, H. Gimbert, and T. A. Henzinger, “Randomness for free,” Inf. Comput., vol. 245, pp. 3–16, 2015.

[56] K. Chatterjee, R. Majumdar, and M. Jurdzinski, “On Nash equilibria in stochastic games,” in Computer Science Logic, 18th International Workshop, CSL 2004, 13th Annual Conference of the EACSL, Karpacz, Poland, September 20-24, 2004, Proceedings (J. Marcinkowski and A. Tarlecki, eds.), vol. 3210 of Lecture Notes in Computer Science, pp. 26–40, Springer, 2004.

[57] A. Mashiah-Yaakovi, “Correlated equilibria in stochastic games with Borel measurable payoffs,” Dyn. Games Appl., vol. 5, no. 1, pp. 120–135, 2015.

[58] H. Gimbert and F. Horn, “Solving simple stochastic tail games,” in Proceedings of the Twenty-First Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2010, Austin, Texas, USA, January 17-19, 2010 (M. Charikar, ed.), pp. 847–862, SIAM, 2010.

[59] S. Kiefer, R. Mayr, M. Shirmohammadi, P. Totzke, and D. Wojtczak, “How to play in infinite MDPs (invited talk),” in 47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference) (A. Czumaj, A. Dawar, and E. Merelli, eds.), vol. 168 of LIPIcs, pp. 3:1–3:18, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

[60] L. de Alfaro, T. A. Henzinger, and O. Kupferman, “Concurrent reachability games,” Theor. Comput. Sci., vol. 386, no. 3, pp. 188–217, 2007.

[61] L. S. Shapley and R. Snow, “Basic solutions of discrete games,” Contributions to the Theory of Games, vol. 1, no. 24, pp. 27–27, 1950.

[62] H. Gimbert and E. Kelmendi, “Two-player perfect-information shift-invariant submixing stochastic games are half-positional,” CoRR, vol. abs/1401.6575, 2014.

[63] S. Le Roux, “From winning strategy to nash equilibrium,” Math. Log. Q., vol. 60, no. 4-5, pp. 354–371, 2014.

[64] R. A. Purves and W. D. Sudderth, “Perfect information games with upper semicontinuous payoffs,” Mathematics of Operations Research, vol. 36, no. 3, pp. 468–473, 2011.

[65] J. Flesch, P. J.-J. Herings, J. Maes, and A. Predtetchinski, “Individual upper semicontinuity and subgame perfect epsilon-equilibria in games with almost perfect information,” Economic Theory, vol. 73, no. 2-3, pp. 695–719, 2022.

[66] H. Gimbert and W. Zielonka, “Games where you can play optimally without any memory,” in Proc. 25th International Conference on Concurrency

Theory (CONCUR'14), vol. 3653 of Lecture Notes in Computer Science, pp. 428–442, Springer, 2005.

[67] S. Kiefer, R. Mayr, M. Shirmohammadi, and D. Wojtczak, "Parity objectives in countable MDPs," in 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017, pp. 1–11, IEEE Computer Society, 2017.

[68] J. Filar and K. Vrieze, Competitive Markov decision processes. Springer Science & Business Media, 2012.

[69] B. Bordais, P. Bouyer, and S. Le Roux, "Optimal strategies in concurrent reachability games," CoRR, vol. abs/2110.14724, 2021.

[70] K. Chatterjee, L. de Alfaro, and T. A. Henzinger, "Trading memory for randomness," in 1st International Conference on Quantitative Evaluation of Systems (QEST 2004), 27-30 September 2004, Enschede, The Netherlands, pp. 206–217, IEEE Computer Society, 2004.

[71] E. Boros, O. Cepek, and V. Gurvich, "Separable discrete functions: Recognition and sufficient conditions," Discret. Math., vol. 342, no. 5, pp. 1275–1292, 2019.

[72] B. Bordais, P. Bouyer, and S. Le Roux, "From local to global determinacy in concurrent graph games," CoRR, vol. abs/2107.04081, 2021.

[73] M. Develin and S. Payne, "Discrete bidding games," Electron. J. Comb., vol. 17, no. 1, 2010.

[74] G. Avni, T. A. Henzinger, and D. Zikelic, "Bidding mechanisms in graph games," in 44th International Symposium on Mathematical Foundations of Computer Science, MFCS 2019, August 26-30, 2019, Aachen, Germany (P. Rossmanith, P. Heggernes, and J. Katoen, eds.), vol. 138 of LIPIcs, pp. 11:1–11:13, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[75] M. L. Fredman and L. Khachiyan, "On the complexity of dualization of monotone disjunctive normal forms," J. Algorithms, vol. 21, no. 3, pp. 618–628, 1996.

[76] T. Eiter, K. Makino, and G. Gottlob, "Computational aspects of monotone dualization: A brief survey," Discret. Appl. Math., vol. 156, no. 11, pp. 2035–2049, 2008.

[77] T. Eiter and G. Gottlob, "Hypergraph transversal computation and related problems in logic and AI," in Logics in Artificial Intelligence, European Conference, JELIA 2002, Cosenza, Italy, September, 23-26, Proceedings (S. Flesca, S. Greco, N. Leone, and G. Ianni, eds.), vol. 2424 of Lecture Notes in Computer Science, pp. 549–564, Springer, 2002.

[78] A. Tarski, A decision method for elementary algebra and geometry. Springer, 1998.

[79] J. Renegar, "On the computational complexity and geometry of the first-order theory of the reals, part III: quantifier elimination," J. Symb. Comput., vol. 13, no. 3, pp. 329–352, 1992.

[80] J. F. Canny, "Some algebraic and geometric computations in PSPACE," in Proceedings of the 20th Annual ACM Symposium on Theory of Computing, May 2-4, 1988, Chicago, Illinois, USA (J. Simon, ed.), pp. 460–467, ACM, 1988.

[81] P. Bouyer, N. Markey, and D. Stan, "Stochastic equilibria under imprecise deviations in terminal-reward concurrent games," in Proceedings of the Seventh International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2016, Catania, Italy, 14-16 September 2016 (D. Cantone and G. Delzanno, eds.), vol. 226 of EPTCS, pp. 61–75, 2016.

[82] S. Le Roux and A. Pauly, "Equilibria in multi-player multi-outcome infinite sequential games," Inf. Comput., vol. 276, p. 104557, 2021.