

# The Complexity of Nash Equilibria in Limit-Average Games

Michael Ummels

`ummels@lsv.ens-cachan.fr`


LSV, CNRS & ENS Cachan

(Joint Work with Dominik Wojtczak)

CONCUR 2011, Aachen

# Limit-Average Games


What is a limit-average game? [Let's play!](#)



Rewards:

# Limit-Average Games


What is a limit-average game? [Let's play!](#)



Rewards: 1  
0  
0

# Limit-Average Games

What is a limit-average game? [Let's play!](#)




Rewards:

1	0
0	1
0	0

# Limit-Average Games

What is a limit-average game? [Let's play!](#)




Rewards:

	1	0	1
0	1	0	
0	0	0	

# Limit-Average Games

What is a limit-average game? [Let's play!](#)




Rewards:

	1	0	1	0
0	1	0	1	
0	0	0	0	

# Limit-Average Games


What is a limit-average game? [Let's play!](#)



	1	0	1	0	2
Rewards:	0	1	0	1	0
	0	0	0	0	0

# Limit-Average Games


What is a limit-average game? [Let's play!](#)



	1	0	1	0	2	0
Rewards:	0	1	0	1	0	0
	0	0	0	0	0	4

# Limit-Average Games


What is a limit-average game? [Let's play!](#)



	1	0	1	0	2	0	2
Rewards:	0	1	0	1	0	0	0
	0	0	0	0	0	4	0

# Limit-Average Games

What is a limit-average game? [Let's play!](#)




$$\text{Rewards: } \begin{matrix} 1 & 0 & 1 & 0 & 2 & 0 & 2 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 4 & 0 \end{matrix} \begin{pmatrix} 2 & 0 \\ 0 & 0 \\ 0 & 4 \end{pmatrix}^{\omega} \quad (\text{Mean}) \text{ payoff: } \begin{pmatrix} 1 \\ 0 \\ 2 \end{pmatrix}$$

# Strategies and Payoffs

**Strategy:**  $\sigma: (\text{States} \cdot \text{Actions}^{\text{Players}})^* \cdot \text{States} \rightarrow \mathcal{D}(\text{Actions})$

A strategy profile induces a probability measure on plays.




**Expected payoff:**  $\text{payoff}(\bar{\sigma}) := \int_{\text{Plays}} \text{payoff}(\pi) d \Pr^{\bar{\sigma}}(\pi) = (1, 2, 0)$

# Strategies and Payoffs

Strategy profile:  $\bar{\sigma}: (\text{States} \cdot \text{Actions}^{\text{Players}})^* \cdot \text{States} \rightarrow \mathcal{D}(\text{Actions})^{\text{Players}}$

A strategy profile induces a probability measure on plays.




Expected payoff:  $\text{payoff}(\bar{\sigma}) := \int_{\text{Plays}} \text{payoff}(\pi) d\text{Pr}^{\bar{\sigma}}(\pi) = (1, 2, 0)$

# Nash Equilibria

**Definition:** A strategy profile is a **Nash equilibrium** if no player can gain from unilaterally switching to a different strategy.

**Question:** Do Nash equilibria always exist?

**Example:** *Hide or Run* (de Alfaro and Henzinger 2000)



**Answer:** No!


With the strategy of running in each step with probability  $\varepsilon$ , Player 1 can achieve payoff  $1 - \varepsilon$ . Hence, any equilibrium must have payoff  $(1, -1)$ .

# Nash Equilibria

**Definition:** A strategy profile is a **Nash equilibrium** if no player can gain from unilaterally switching to a different strategy.

**Question:** Do Nash equilibria always exist?

**Example:** *Hide or Run* (de Alfaro and Henzinger 2000)



**Answer:** No!


Let  $(\sigma_1, \sigma_2)$  be a Nash equilibrium with payoff  $(1, -1)$ , and let  $\sigma_1(\text{run} \mid s^k) = p > 0$  for the smallest such  $k$ .

By throwing with probability 1, Player 2 can achieve payoff at least  $p - 1$ .

# Example

But there may also be more than one equilibrium...

**Example:** Terminal-reward game (nonzero rewards only on terminals).




Nash equilibrium where Player 1 receives payoff 0.

**Observation:** Randomised strategies are more powerful.

# Example

But there may also be more than one equilibrium...

**Example:** Terminal-reward game (nonzero rewards only on terminals).




Nash equilibrium where Player 1 receives payoff 0.

**Observation:** Randomised strategies are more powerful.

# Example

But there may also be more than one equilibrium...

**Example:** Terminal-reward game (nonzero rewards only on terminals).



Nash equilibrium where Player 1 receives payoff 1.

**Observation:** Randomised strategies are more powerful.

# The Problem NE

**Goal:** Compute a Nash equilibrium with a “good” payoff.

**The problem NE:** Given a game  $\mathcal{G}$  and  $\bar{x}, \bar{y} \in (\mathbb{Q} \cup \{\pm\infty\})^k$ , decide whether the game has a Nash equilibrium with expected payoff  $\geq \bar{x}$  and  $\leq \bar{y}$ .

**Special case:** Decide whether there exists a Nash equilibrium at all.

**Variants of the Problem:**

- ▶ Rand. strategies:  $(\text{States} \cdot \text{Actions}^{\text{Players}})^* \cdot \text{States} \rightarrow \mathcal{D}(\text{Actions})$
- ▶ Pure strategies:  $(\text{States} \cdot \text{Actions}^{\text{Players}})^* \cdot \text{States} \rightarrow \text{Actions}$
- ▶ Memoryless strategies:  $\text{States} \rightarrow \mathcal{D}(\text{Actions})$
- ▶ Pure memoryless strategies:  $\text{States} \rightarrow \text{Actions}$

# Results

The Complexity of NE in limit-average games:

	Pure	Randomised
Memoryful	NP-complete*	Undecidable
Memoryless	NP-complete	PSPACE**

\* in P for terminal-reward games

\*\* hard for NP and SqrtSum

All lower bounds hold for the special case!

# Pure memoryless Nash equilibria

Deciding the existence of a pure memoryless NE with payoff  $\geq \bar{x}$  and  $\leq \bar{y}$ :

- ▶ Guess a pure memoryless strategy profile  $\bar{\sigma}$ .
- ▶ For each player  $i$ :
  1. Compute the payoff  $r_i$  of  $\bar{\sigma}$  for player  $i$ .
  2. Compute the maximal payoff  $z_i$  player  $i$  can achieve by herself.
  3. Check whether  $x_i \leq r_i \leq y_i$  and  $z_i \leq r_i$ .

1. and 2. are doable in polynomial time (cycle analysis).

## Theorem

NE for pure memoryless strategies is in NP.

NP-hardness can be proved by a reduction from HAMILTONIAN-CYCLE (for games with three players).

# Results

The Complexity of NE in limit-average games:

	Pure	Randomised
Memoryful	NP-complete*	Undecidable
Memoryless	NP-complete	PSPACE**

\* in P for terminal-reward games

\*\* hard for NP and SqrtSum

All lower bounds hold for the special case!

# Memoryless Nash Equilibria

Deciding the existence of a memoryless NE with payoff  $\geq \bar{x}$  and  $\leq \bar{y}$ :

- ▶ Guess the **support**  $S$  of a memoryless strategy profile  $\bar{\sigma}$ .
- ▶ Evaluate an existential first-order sentence  $\psi$  (which is polynomial-time computable from  $\mathcal{G}$ ,  $\bar{x}$ ,  $\bar{y}$  and  $S$ ) over  $\mathfrak{R} = (\mathbb{R}, +, \cdot, 0, 1)$ .

$\psi$  states that there exists a memoryless Nash equilibrium  $\bar{\sigma}$  with payoff  $\geq \bar{x}$  and  $\leq \bar{y}$  whose support is precisely  $S$ .

**Note:** The existential theory of the Reals is decidable in PSPACE.

## Theorem

NE is in NPSPACE, and thus in PSPACE, for memoryless strategies.

NP-hardness is proved using a reduction from SAT.

# The Square-Root-Sum Problem

**Square-Root-Sum Problem (SQRT-SUM):** Given  $d_1, \dots, d_n \in \mathbb{N}$  and  $k \in \mathbb{N}$ , decide whether  $\sum_{i=1}^n \sqrt{d_i} \geq k$ .

The precise complexity of SQRT-SUM is not known:

- ▶ Known to be in PSPACE (actually in the 4th level of the counting hierarchy).
- ▶ No non-trivial lower bounds known.

**Open Problem (since 1970s):** Does SQRT-SUM lie inside the polynomial hierarchy? Is it in NP?

## Theorem

There is a polynomial-time reduction from SQRT-SUM to NE for memoryless strategies in terminal-reward games (with eight players).

# Results

The Complexity of NE in limit-average games:

	Pure	Randomised
Memoryful	NP-complete*	Undecidable
Memoryless	NP-complete	PSPACE**

\* in P for terminal-reward games

\*\* hard for NP and SqrtSum

All lower bounds hold for the special case!

# Pure Strategies

Define  $\text{pval}_i^{\mathcal{G}}(s) = \inf_{\bar{\sigma}} \sup_{\tau} \text{payoff}_i^{\mathcal{G}(s)}(\bar{\sigma}[\sigma_i \mapsto \tau])$ , where  $\bar{\sigma}$  ranges over all pure strategy profiles and  $\tau$  ranges over all strategies of Player  $i$ .

**Observation:**  $\bar{\sigma}$  can only be a pure Nash equilibrium if  $\text{payoff}_i(\bar{\sigma}) \geq \text{pval}_i^{\mathcal{G}}(s)$  for each player  $i$  and for all states  $s$  reachable with  $\bar{\sigma}$ .

In fact, this condition is sufficient for having a pure equilibrium.

## Lemma


If  $\bar{\sigma}$  is a pure strategy profile such that  $\text{payoff}_i(\bar{\sigma}) \geq \text{pval}_i^{\mathcal{G}}(s)$  for each player  $i$  and all states  $s$  reachable with  $\bar{\sigma}$ , then there exists a pure Nash equilibrium  $\bar{\sigma}^*$  such that  $\text{payoff}(\bar{\sigma}^*) = \text{payoff}(\bar{\sigma})$ .

**NP algorithm:** Guess lower bound  $\bar{z} \geq \bar{x}$  on payoff and check for existence of good play that avoids states  $s$  with  $\text{pval}_i^{\mathcal{G}}(s) > z_i \rightsquigarrow$  graph problem.

# Finding a Path

**MULTI-MEAN-WEIGHT:** Given a multi-weighted graph  $G$ , vertex  $v_0$ , and  $\bar{x}, \bar{y} \in (\mathbb{Q} \cup \{\pm\infty\})^k$ , does there exist an infinite path from  $v_0$  with mean-weight  $\geq \bar{x}$  and  $\leq \bar{y}$ ?

Example (Chatterjee-Doyen-Henzinger-Raskin 2010)



There is an infinite path with mean weight  $(1, 1)$ , but no ultimately periodic path achieves this mean weight.

Linear programming comes to our rescue...

## Theorem

MULTI-MEAN-WEIGHT is decidable in polynomial time.

# Results

The Complexity of NE in limit-average games:

	Pure	Randomised
Memoryful	NP-complete*	Undecidable
Memoryless	NP-complete	PSPACE**

\* in P for terminal-reward games

\*\* hard for NP and SqrtSum

All lower bounds hold for the special case!

# Undecidability

Undecidability is shown by a reduction from the non-halting problem for two-counter machines.

## Theorem

For every two-counter machine  $\mathcal{M}$  one can effectively construct a terminal-reward game  $\mathcal{G}$  (with 14 players) such that the computation of  $\mathcal{M}$  is infinite iff  $\mathcal{G}$  has a (randomised) Nash equilibrium with payoff  $\geq 0$  for Player 1.

By combining  $\mathcal{G}$  with a game that has no Nash equilibrium, we get:

## Corollary

The existence of a (randomised) Nash equilibrium is undecidable (not r.e.) for terminal-reward games.

Future work:

- ▶ Games with few players
- ▶ Discounted games
- ▶  $\epsilon$ -equilibria