

Mean-payoff Games: Extensions and Variations

LCCC Workshop
Lund

J.-F. Raskin
Université Libre de Bruxelles, U.L.B.


based on joint works with K. Chatterjee (IST Austria), L. Doyen (ENS Cachan)
T. Henzinger (IST Austria), A. Rabinovitch (U Tel Aviv), M. Randour (U Mons) and Y. Velner (U Tel Aviv)

Plan

- ★ **Synthesis** and 2 player zero-sum **games** on graphs
- ★ Mean-payoff games (as an example of **quantitative** games):
 - ★ **Mean-payoff games** in 1 dimension
 - ★ **Extensions:** k dimensions - recent results
 - ★ **Variations: window objectives**
 - ★ Motivations and definitions
 - ★ Algorithm for 1 dimension
 - ★ Overview of results
- ★ **Conclusion**

Games for Synthesis (of Reactive Systems)

☞ support the design process with **automatic synthesis**




- Sys is constructed by an algorithm
- Sys is **correct** by construction
- Underlying theory: **2-player zero-sum games**
- Env is **adversarial** (worst-case assumption)

Winning strategy = Correct Sys

**Preliminaries:
2-player Zero-sum Games
on Graphs**


2-player Zero-sum Games on Graphs



(Finite) directed graph

Two types of vertices
(Player 1 and Player 2 vertices)

2-player Zero-sum Games on Graphs



How to play ?


One token is placed on initial vertex

Players play for an infinite number of rounds:

- in each round: the player that owns the vertex moves the token to an adjacent vertex


Outcome=infinite path

2-player Zero-sum Games on Graphs




I

2-player Zero-sum Games on Graphs




$1 \rightarrow 2$

2-player Zero-sum Games on Graphs




1 → 2 → 1

2-player Zero-sum Games on Graphs




1 → 2 → 1 → 4

2-player Zero-sum Games on Graphs




$1 \rightarrow 2 \rightarrow 1 \rightarrow 4 \rightarrow 5$

2-player Zero-sum Games on Graphs




$1 \rightarrow 2 \rightarrow 1 \rightarrow 4 \rightarrow 5 \rightarrow 3$

2-player Zero-sum Games on Graphs




1 → 2 → 1 → 4 → 5 → 3 → 4

2-player Zero-sum Games on Graphs



1 → 2 → 1 → 4 → 5 → 3 → 4 → 5

2-player Zero-sum Games on Graphs



1 → 2 → 1 → 4 → 5 → 3 → 4 → 5 → 4 → ...

2-player Zero-sum Games on Graphs

Who is winning ? winning conditions

➡ $Win_1 \subseteq V^\omega$: Set of good outcomes (paths) for Player 1


➡ $Win_2 = V^\omega \setminus Win_1$ (zero sum)

Examples of winning conditions:

- $Win_1 = \{ \pi \mid \pi \text{ visits } \text{Good} \}$
Reachability winning condition
- $Win_1 = \{ \pi \mid \pi \text{ visits } \text{Good} \text{ infinitely often} \}$
Büchi winning condition


Strategies

Unfolding of the game graph




Strategies

Unfolding of the game graph




Strategies

Unfolding of the game graph

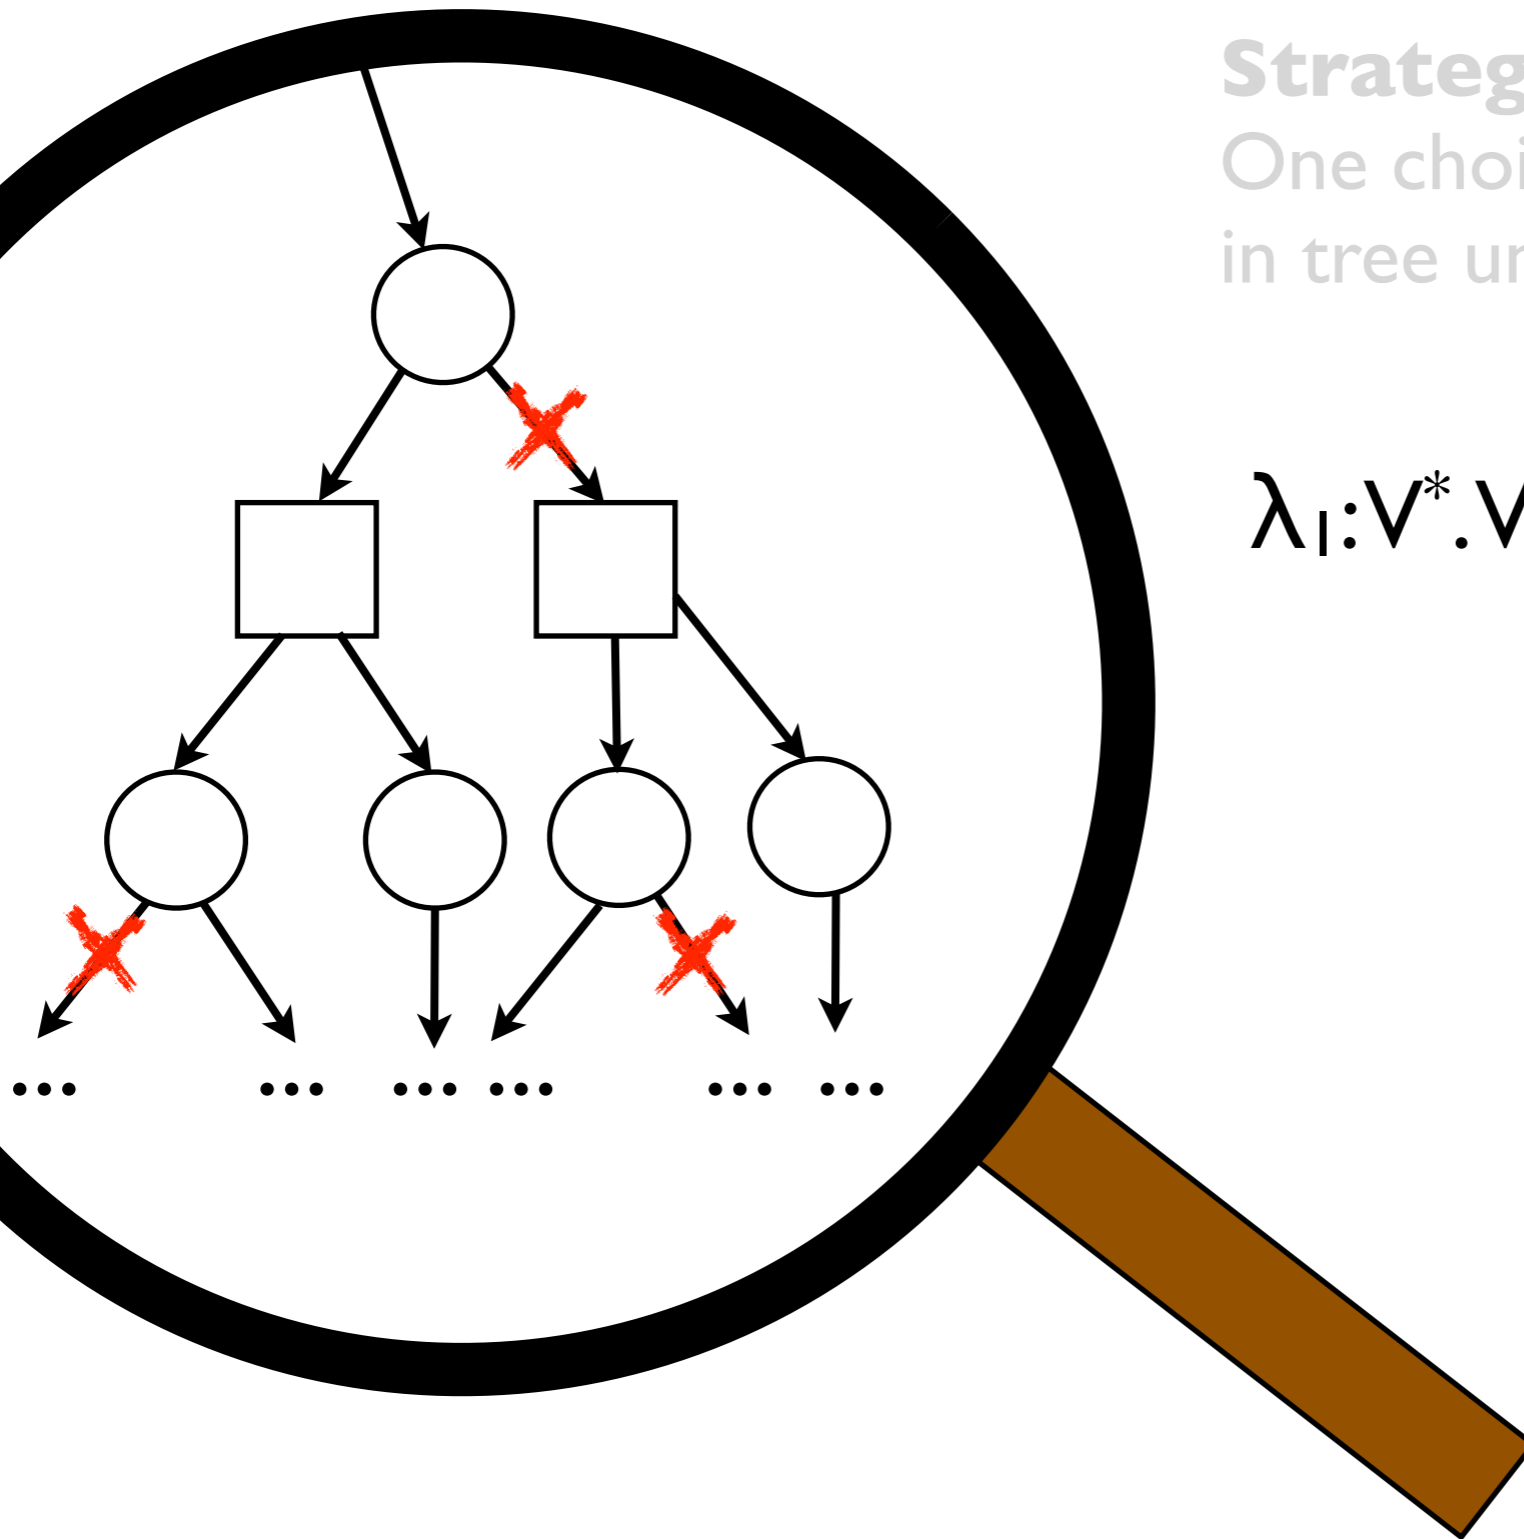


Strategies

Strategy for Player I =
One choice in each node of Player I
in tree unfolding




Strategies



Strategy for Player I =
One choice in each node of Player I
in tree unfolding

$$\lambda_I: V^*.V_I \rightarrow \text{edge}$$

Strategies



Strategy for Player I =
One choice in each node of Player I
in tree unfolding

$$\lambda_I: V^*.V_I \rightarrow \text{edge}$$


Strategy is **winning**
(for Player I),
if **all** branches of the resulting
tree in the winning condition

Types of strategies

(Player I) strategy:

$\lambda_I: V^*.V_I \rightarrow \text{edge}$.


$\Sigma_I = \text{set of strategies of Player I}$



Memoryless strategy:

$\lambda_{I,m}: V_I \rightarrow \text{edge}$.


$\Sigma_{I,m} = \text{set of memoryless strategies of Player I}$



Finite-memory strategy:

$\lambda_{I,f}: V^*.V_I \rightarrow \text{edge}$ but **regular** (Moore machine)

$\Sigma_{I,f} = \text{set of finite memory strategies of Player I}$



Randomized strategy:

$\lambda_{I,m}: V^*.V_I \rightarrow \text{Dist}(\text{edge})$.

$\Sigma_{I,m} = \text{set of randomized strategies of Player I}$




Decision Problem

Given:

- a game graph G
- a winning condition $Win_1 \subseteq V^\omega$
- **decide** if Player 1 has a **winning strategy**
- **determinacy**:
either Player 1 has a winning for Win_1
or Player 2 has a winning strategy for $Win_2 = V^\omega \setminus Win_1$
It is true for a large class of objectives,
e.g. ω -regular objectives

2-player Zero-sum Games on Graphs



One token is placed on initial vertex


We play an infinite number of rounds:

- in each round: the player that owns the vertex that the token is on chooses the next vertex

Classical model for the synthesis of reactive systems:
Player I=system and Player II=environment
For embedded systems, we need **quantities** !


Quantitative Objectives: Mean-Payoff Games

Mean-Payoff Games [EM79]



defined on
weighted directed graphs

Mean-Payoff Games [EM79]



- : positions of **maximizer=system**
- : positions of **minimizer=environment**

Edges are labelled with rewards

(1,4) (4,5) (5,4) ... (4,5) (5,4) ... =play
 4 3 -1 3 -1 ...



$$= \mathbf{Lim\ Sup}_{n \rightarrow +\infty} \sum_{i=1, i=n} r_i / n$$

$$= \mathbf{MP}((1,4) (4,5) (5,4) \dots (4,5) (5,4) \dots) = 1$$

Win = { play | MP(play) ≥ **c** }

Note: **not** ω -regular.

Mean-Payoff Games [EM79]




: positions of **maximizer=system**




: positions of **minimizer=environment**

Edges are labelled with rewards


$(1,4)$ $(4,5)$ $(5,4)$... $(4,5)$ $(5,4)$... =play
 4 3 -1 3 -1 ...



Mean-Payoff Games [EM79]



Lim Sup - Lim Inf do **not** define the same set of plays



Mean-payoff Games

Theorem [EM79, Jur98, ZP97, GZ09]

(i) *In mean-payoff games, the two players can play optimally with memoryless strategies*

(ii) *The winner can be decided in $\mathbf{NP} \cap \mathbf{coNP}$ and in pseudo-polynomial time*

Rem: Those results hold no matter if $\lim \sup / \lim \inf$ is used in the definition and the winner is the same for the two definitions

Mean-payoff Games

Theorem [EM79, Jur98, ZP97, GZ09]

(i) *In mean-payoff games, the two players can play optimally with memoryless strategies*

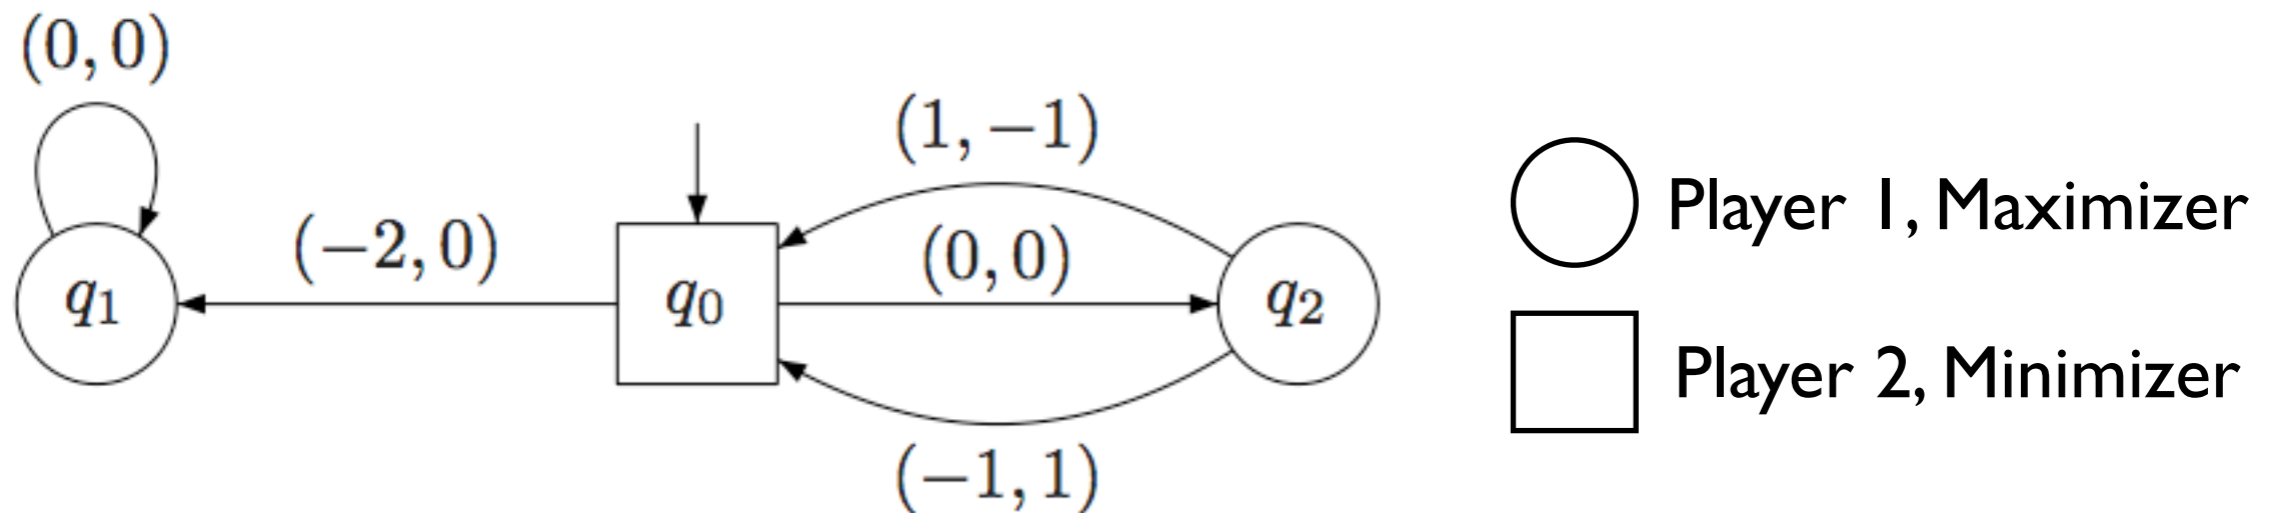
(ii) *The winner can be decided in $\text{NP} \cap \text{coNP}$ and in pseudo-polynomial time*

Rem: Those results hold no matter if \limsup / \liminf definition and the winner is the same

Open question: are MP games solvable in PTime?

Quantitative Objectives: Multi-dimension Extensions


Multi-dim. Mean-Payoff Games (MMPs)



$$? \exists \lambda_I \text{ s. t. } \mathbf{Outcome}(q_0, \lambda_I) \models \mathbf{MP}_{\text{inf.}} \geq (0,0)$$

Multi-dimension objectives (for Player 1)
=
conjunction of one-dimension objectives

Multi-dim. Mean-Payoff Games (MMPs)




? $\exists \lambda_I$ s. t. **Outcome**(q_0, λ_I) \models $MP_{\text{inf.}} \geq (0,0)$

- Player I has a winning strategy.
- Player I may need **infinite** memory !
- Player II can play memoryless

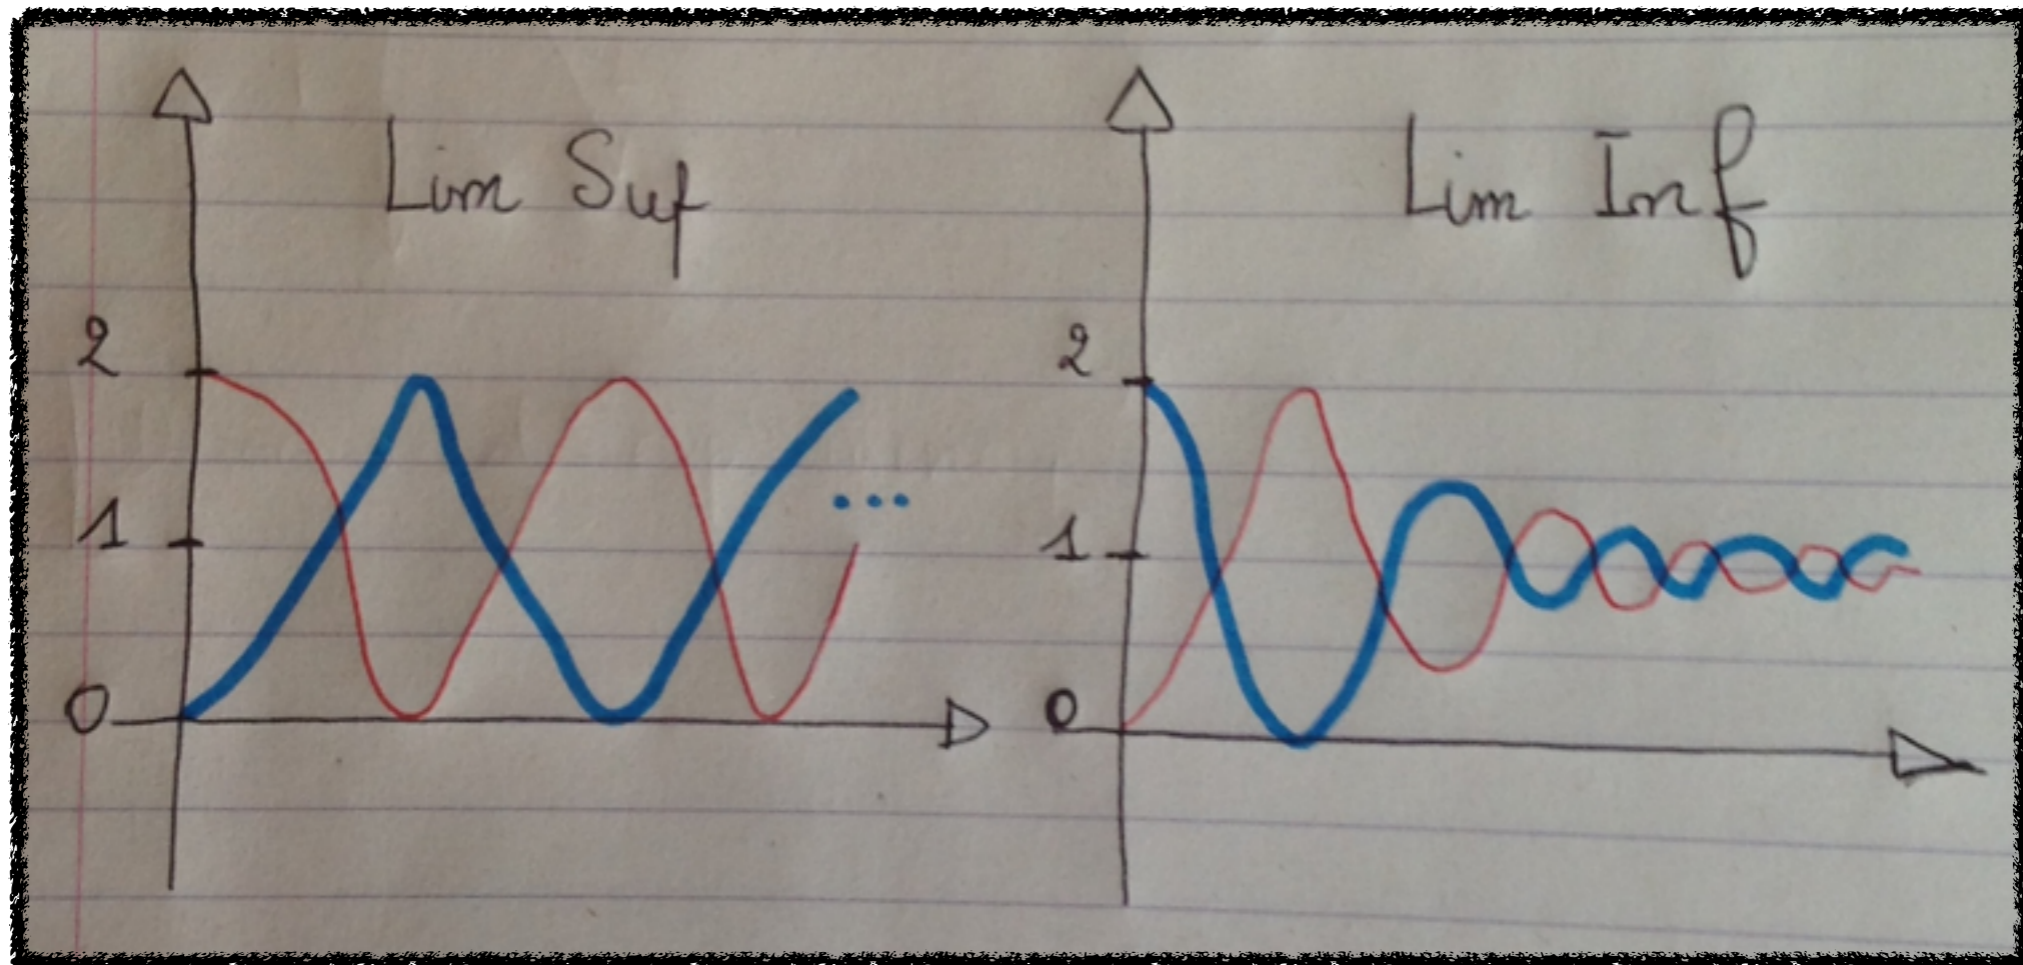

gMPGs - Infinite Memory

To play optimally gMPGs, infinite memory is necessary



- ◆ $(2, 2)$ for Lim Sup MP
- ◆ $(1, 1)$ is achievable for Lim Inf MP
- ◆ None of the two is achievable with finite memory

gMPGs - Infinite Memory




Results for Extensions

	Opt. Stg. Player 1	Opt. Stg. Player 2	Complexity Decision
MP	Memoryless	Memoryless	$NP_n \text{co} NP$
MMPG - Sup	Infinite	Memoryless	$NP_n \text{co} NP$
MMPG - Inf	Infinite	Memoryless	coNP-C
MMPG - Mix	Infinite	Memoryless	coNP-C

Results for Extensions

	Opt. Stg. Player 1	Opt. Stg. Player 2	
MP	Memoryless		$NP_n \text{co} NP$
MMPG - Su		Memoryless	$NP_n \text{co} NP$
MMPG - Inf	<u>Infinite</u>	Memoryless	coNP-C
MMPG - Mix	Infinite	Memoryless	coNP-C

ϵ -optimality achievable with finite memory strategies



Variations on MP: Window Objectives

Window Objectives


- Space for new definitions as classical objectives have **drawbacks**:
 - ① complexity of MP is open
 - ② MMP is sensitive to lim inf. vs. lim sup.
- ➔ **Window objectives**: look at the payoff through a **local finite window sliding** over the play
- **conservative approximations** of MP
- ensure good properties within a bounded time horizon
- algorithms and complexity results:
 - **PTIME-C** for 1 dim. fixed and polynomial window size
 - **EXPTIME-C** for k dim. fixed window size
 - ... and more

Window Objectives - Definitions

Idea: look at payoffs through a local finite window
⇒ mean should be above zero **within** window size


Window Objectives - Definitions

Idea: look at payoffs through a local finite window
⇒ mean should be above zero **within** window size




Window Objectives - Definitions

Idea: look at payoffs through a local finite window
⇒ mean should be above zero **within** window size




Window Objectives - Definitions

Idea: look at payoffs through a local finite window
⇒ mean should be above zero **within** window size




Window Objectives - Definitions

Idea: look at payoffs through a local finite window
⇒ mean should be above zero **within** window size




Window Objectives - Definitions

Idea: look at payoffs through a local finite window
⇒ mean should be above zero **within** window size



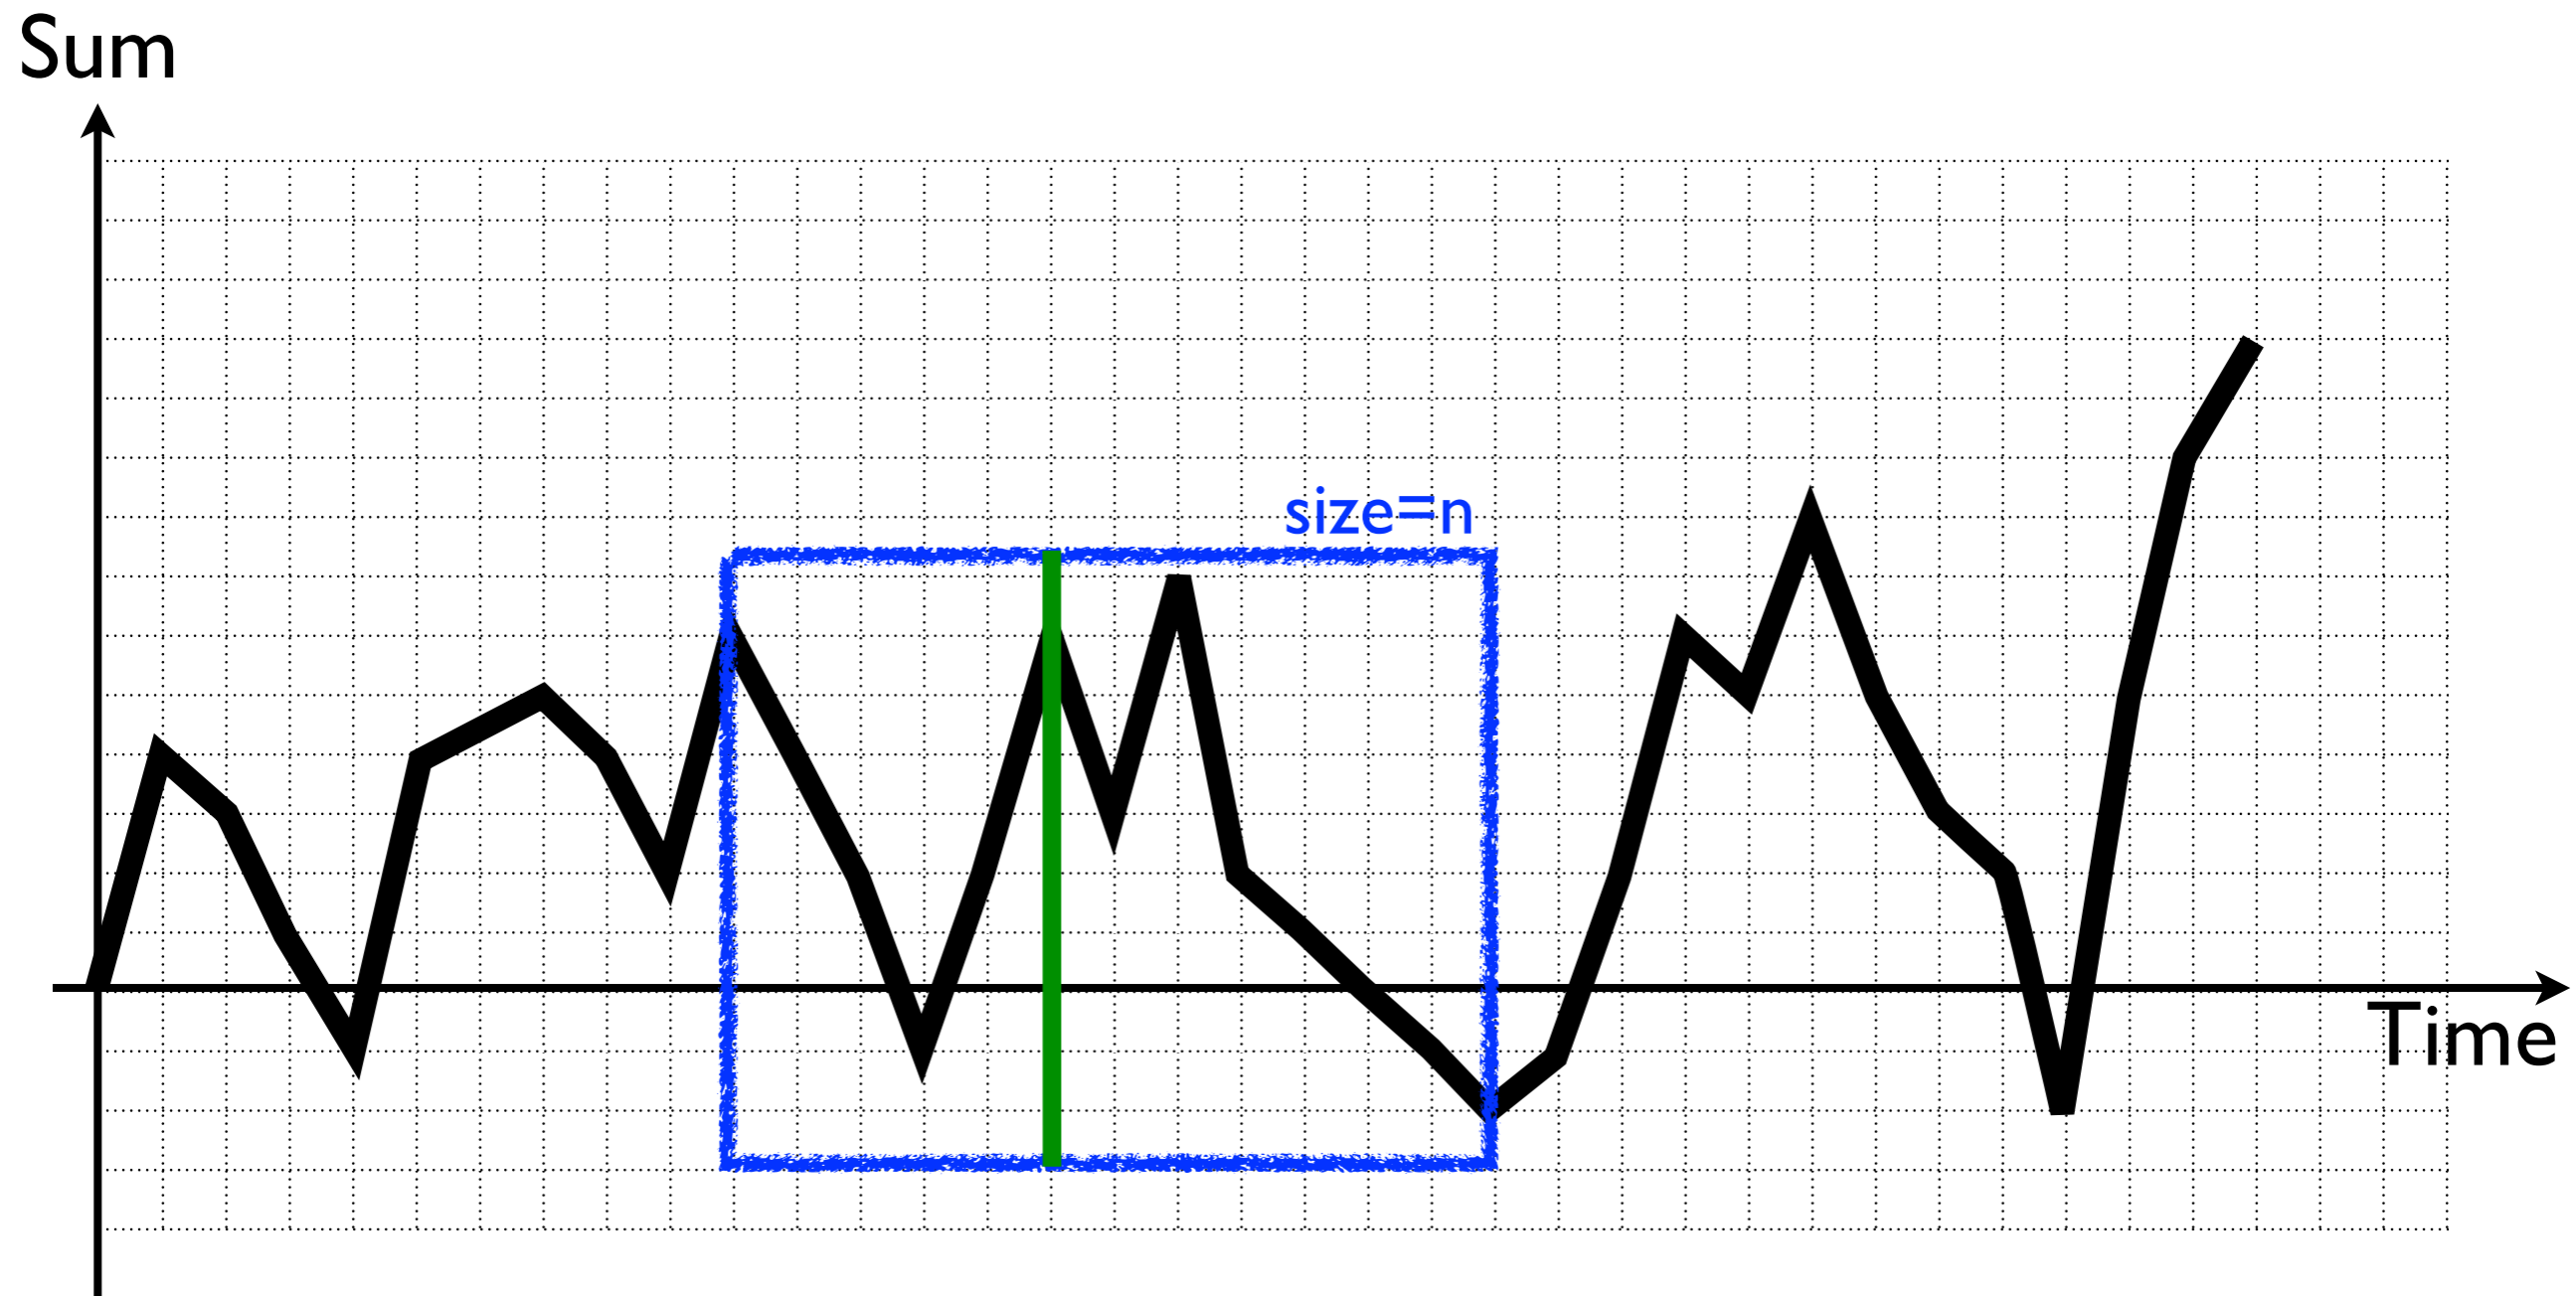
Window Objectives - Definitions

Idea: look at payoffs through a local finite window
⇒ mean should be above zero **within** window size



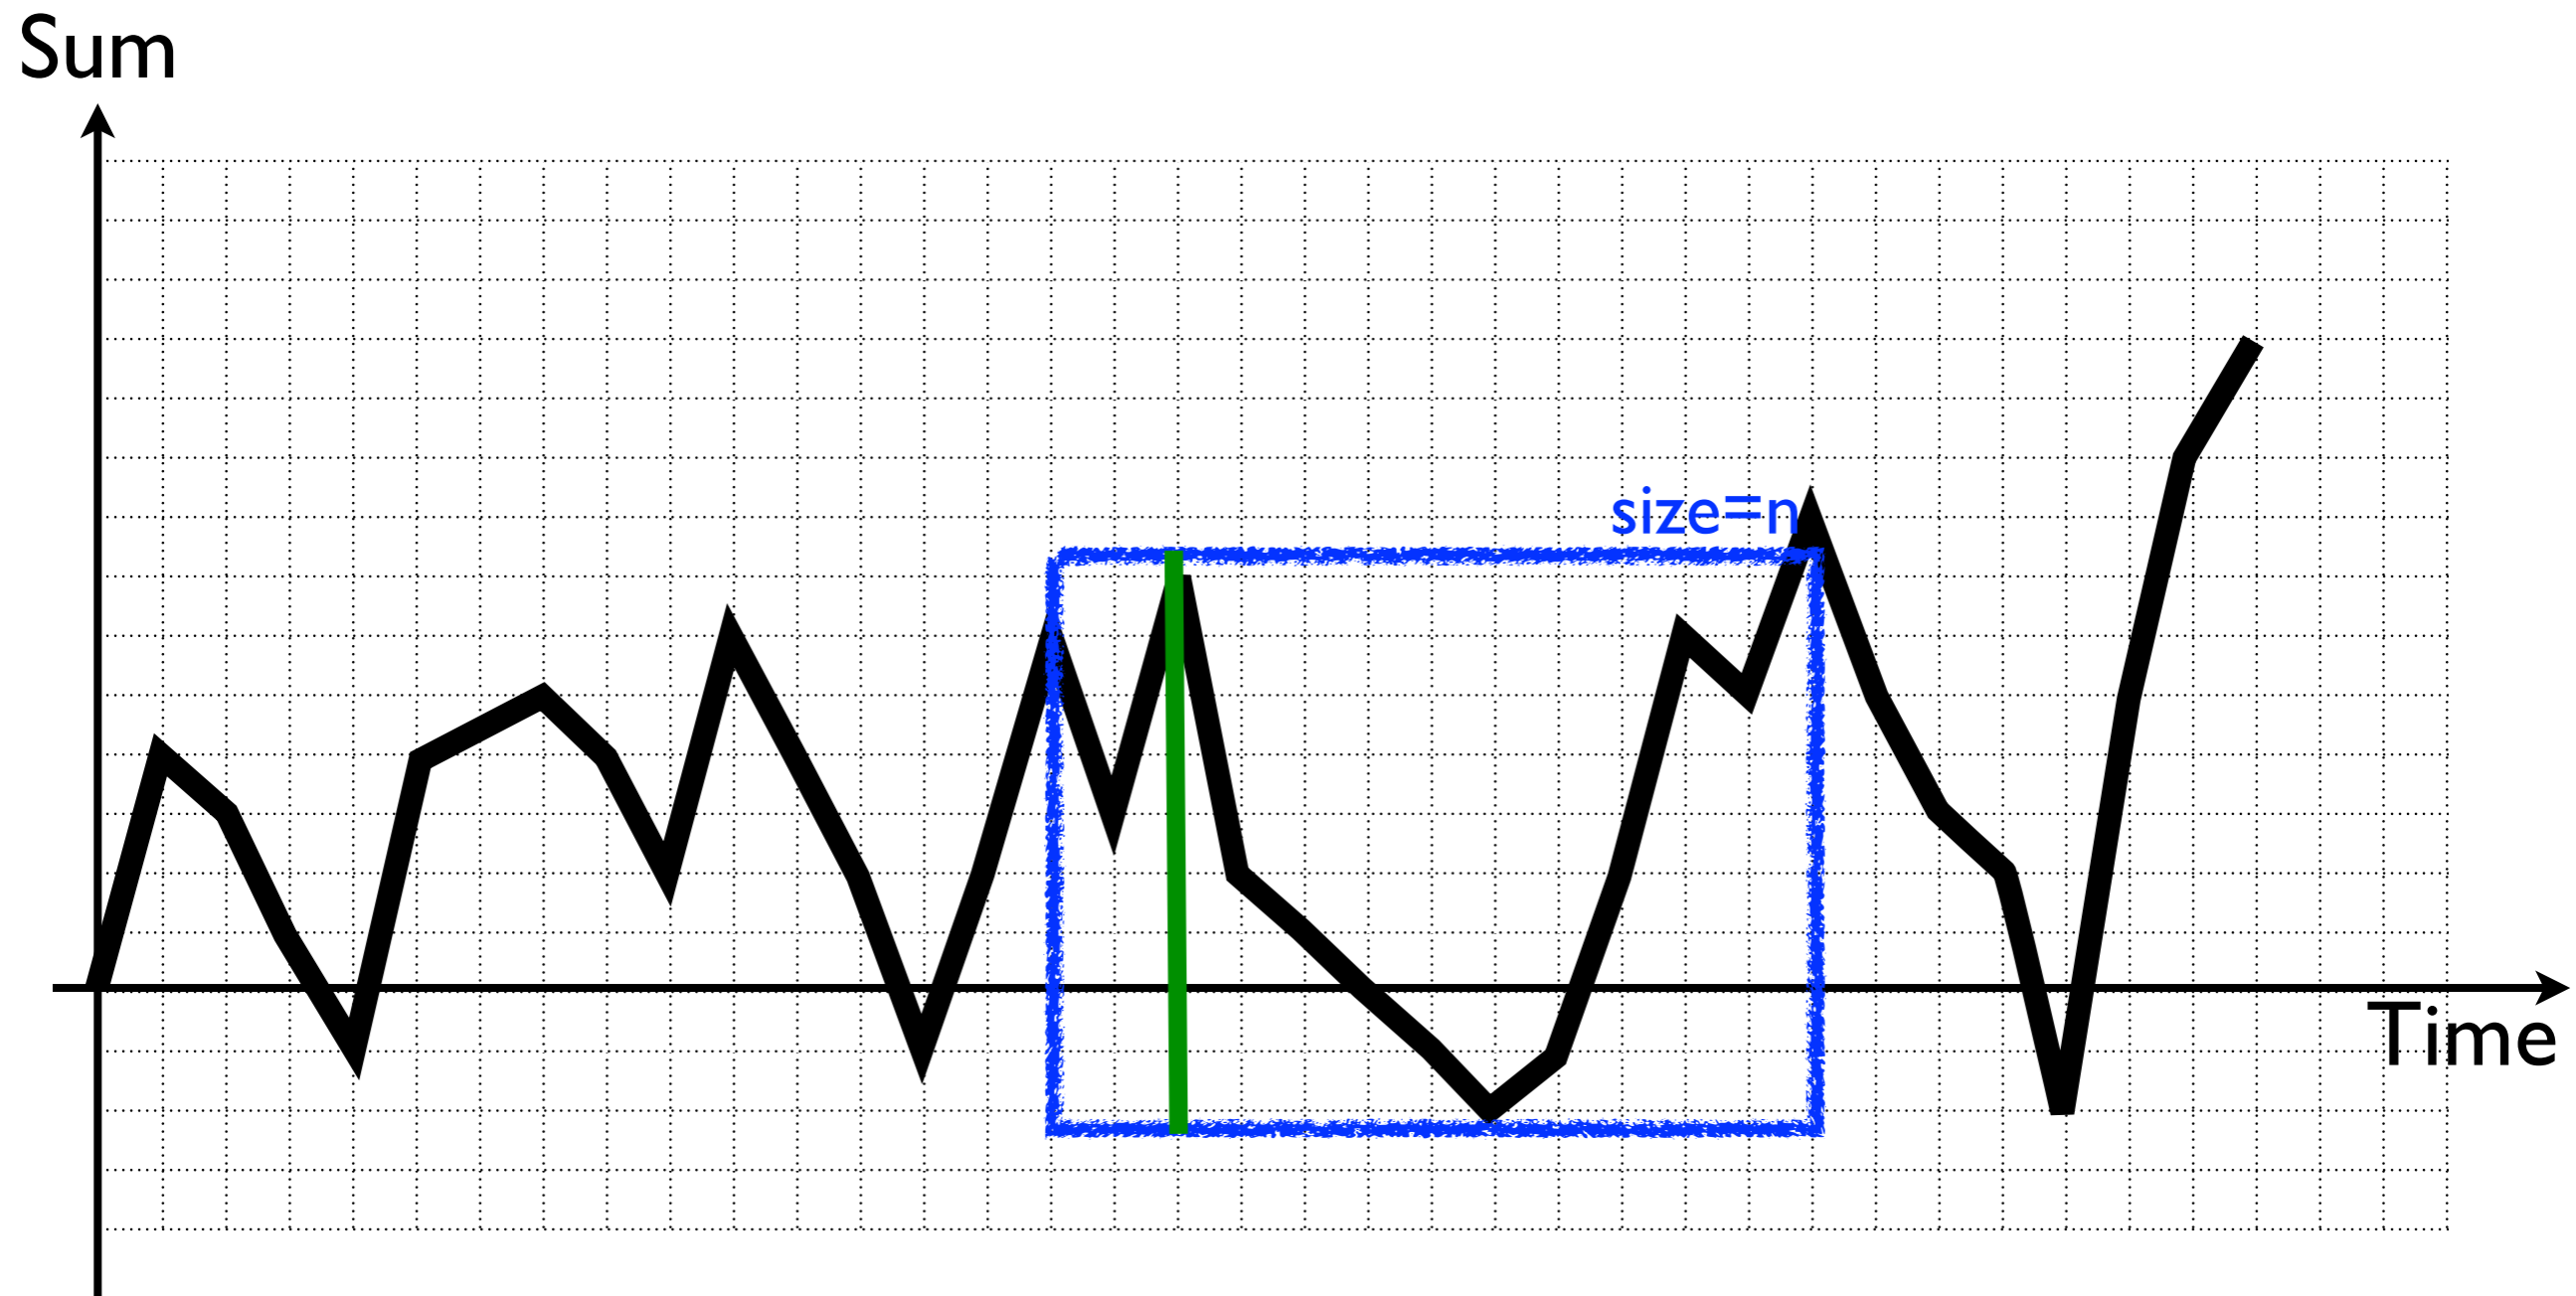
Window Objectives - Definitions

Idea: look at payoffs through a local finite window
⇒ mean should be above zero **within** window size




Window Objectives - Definitions

Idea: look at payoffs through a local finite window
⇒ mean should be above zero **within** window size




Window Objectives - Definitions

Idea: look at payoffs through a local finite window
⇒ mean should be above zero **within** window size



Window Objectives - Definitions

Idea: look at payoffs through a local finite window
⇒ mean should be above zero **within** window size




Formal Definitions

Formal Definitions

Good Window of size n


play ρ satisfies **GW**(n) at position $i \geq 0$
iff $\exists j: i \leq j \leq i+n$ s.t. $\text{Sum}(\rho(i..j)) \geq 0$,
noted $\rho(i..\infty) \models \mathbf{GW}(n)$



Formal Definitions

Good Window of size n

play ρ satisfies **GW**(n) at position $i \geq 0$
iff $\exists j: i \leq j \leq i+n$ s.t. $\text{Sum}(\rho(i..j)) \geq 0$,
noted $\rho(i..\infty) \models \mathbf{GW}(n)$



Window opens at i


successfully closes at j

Otherwise, it unsuccessfully closes at $i+n$

Formal Definitions

Good Window of size n

play ρ satisfies **GW**(n) at position $i \geq 0$
iff $\exists j: i \leq j \leq i+n$ s.t. $\text{Sum}(\rho(i..j)) \geq 0$,
noted $\rho(i..\infty) \models \mathbf{GW}(n)$



Direct Fixed Window of size n :


play ρ satisfies **DFW**(n) iff for all $i \geq 0$, $\rho(i..\infty) \models \mathbf{GW}(n)$

$$\mathbf{DFW}(n) \\ \equiv \square \mathbf{GW}(n)$$

Formal Definitions

Good Window of size n

play ρ satisfies **GW**(n) at position $i \geq 0$
iff $\exists j: i \leq j \leq i+n$ s.t. $\text{Sum}(\rho(i..j)) \geq 0$,
noted $\rho(i..\infty) \models \mathbf{GW}(n)$



Direct Fixed Window of size n:

play ρ satisfies **DFW**(n) iff for all $i \geq 0$, $\rho(i..\infty) \models \mathbf{GW}(n)$

$$\mathbf{DFW}(n) \\ \equiv \square \mathbf{GW}(n)$$

Fixed Window of size n


play ρ satisfies **FixedW**(n) iff
there exists $i \geq 0$, $\rho(i..\infty) \models \mathbf{DFW}(n)$

$$\mathbf{FixedW}(n) \\ \equiv \diamond \square \mathbf{GW}(n)$$

Formal Definitions

Good Window of size n

play ρ satisfies **GW**(n) at position $i \geq 0$
iff $\exists j: i \leq j \leq i+n$ s.t. $\text{Sum}(\rho(i..j)) \geq 0$,
noted $\rho(i..\infty) \models \mathbf{GW}(n)$



Direct Fixed Window of size n:


play ρ satisfies **DFW**(n) iff for all $i \geq 0$, $\rho(i..\infty) \models \mathbf{GW}(n)$

$$\mathbf{DFW}(n) \\ \equiv \square \mathbf{GW}(n)$$

Fixed Window of size n

play ρ satisfies **FixedW**(n) iff
there exists $i \geq 0$, $\rho(i..\infty) \models \mathbf{DFW}(n)$


$$\mathbf{FixedW}(n) \\ \equiv \diamond \square \mathbf{GW}(n)$$



Formal Definitions

Good Window of size n

play ρ satisfies **GW**(n) at position $i \geq 0$
 iff $\exists j: i \leq j \leq i+n$ s.t. $\text{Sum}(\rho(i..j)) \geq 0$,
 noted $\rho(i..\infty) \models \mathbf{GW}(n)$



Direct Fixed Window of size n:

play ρ satisfies **DFW**(n) iff for all $i \geq 0$, $\rho(i..\infty) \models \mathbf{GW}(n)$

$$\mathbf{DFW}(n) \equiv \square \mathbf{GW}(n)$$

Fixed Window of size n


play ρ satisfies **FixedW**(n) iff
 there exists $i \geq 0$, $\rho(i..\infty) \models \mathbf{DFW}(n)$

$$\mathbf{FixedW}(n) \equiv \diamond \square \mathbf{GW}(n)$$


Bounded Window

play ρ satisfies **BW** iff
 there exists $n \geq 1$, $\rho \models \mathbf{FixedW}(n)$

$$\mathbf{BW} \equiv \exists n \cdot \mathbf{FixedW}(n)$$




Examples



- ➔ Fixed window is satisfied for $\text{size} \geq 2$
- ➔ Direct fixed window is not satisfied for any size !


Examples



Player I wins mean-payoff

... but he loses for every window objectives

Examples



For $n=3$: memory needed
For $n=4$: no memory needed


Relation with “classical” objectives

Relation between MP/TP and W.O.

If the answer to one of window mean-payoff problems is YES, then the answer to the mean-payoff threshold problem for threshold zero is also YES.


Relation between MP/TP and W.O.

If the answer to one of window mean-payoff problems is YES, then the answer to the mean-payoff threshold problem for threshold zero is also YES.




Relation between MP/TP and W.O.

If the answer to one of window mean-payoff problems is YES, then the answer to the mean-payoff threshold problem for threshold zero is also YES.




Relation between MP/TP and W.O.

If the answer to one of window mean-payoff problems is YES, then the answer to the mean-payoff threshold problem for threshold zero is also YES.



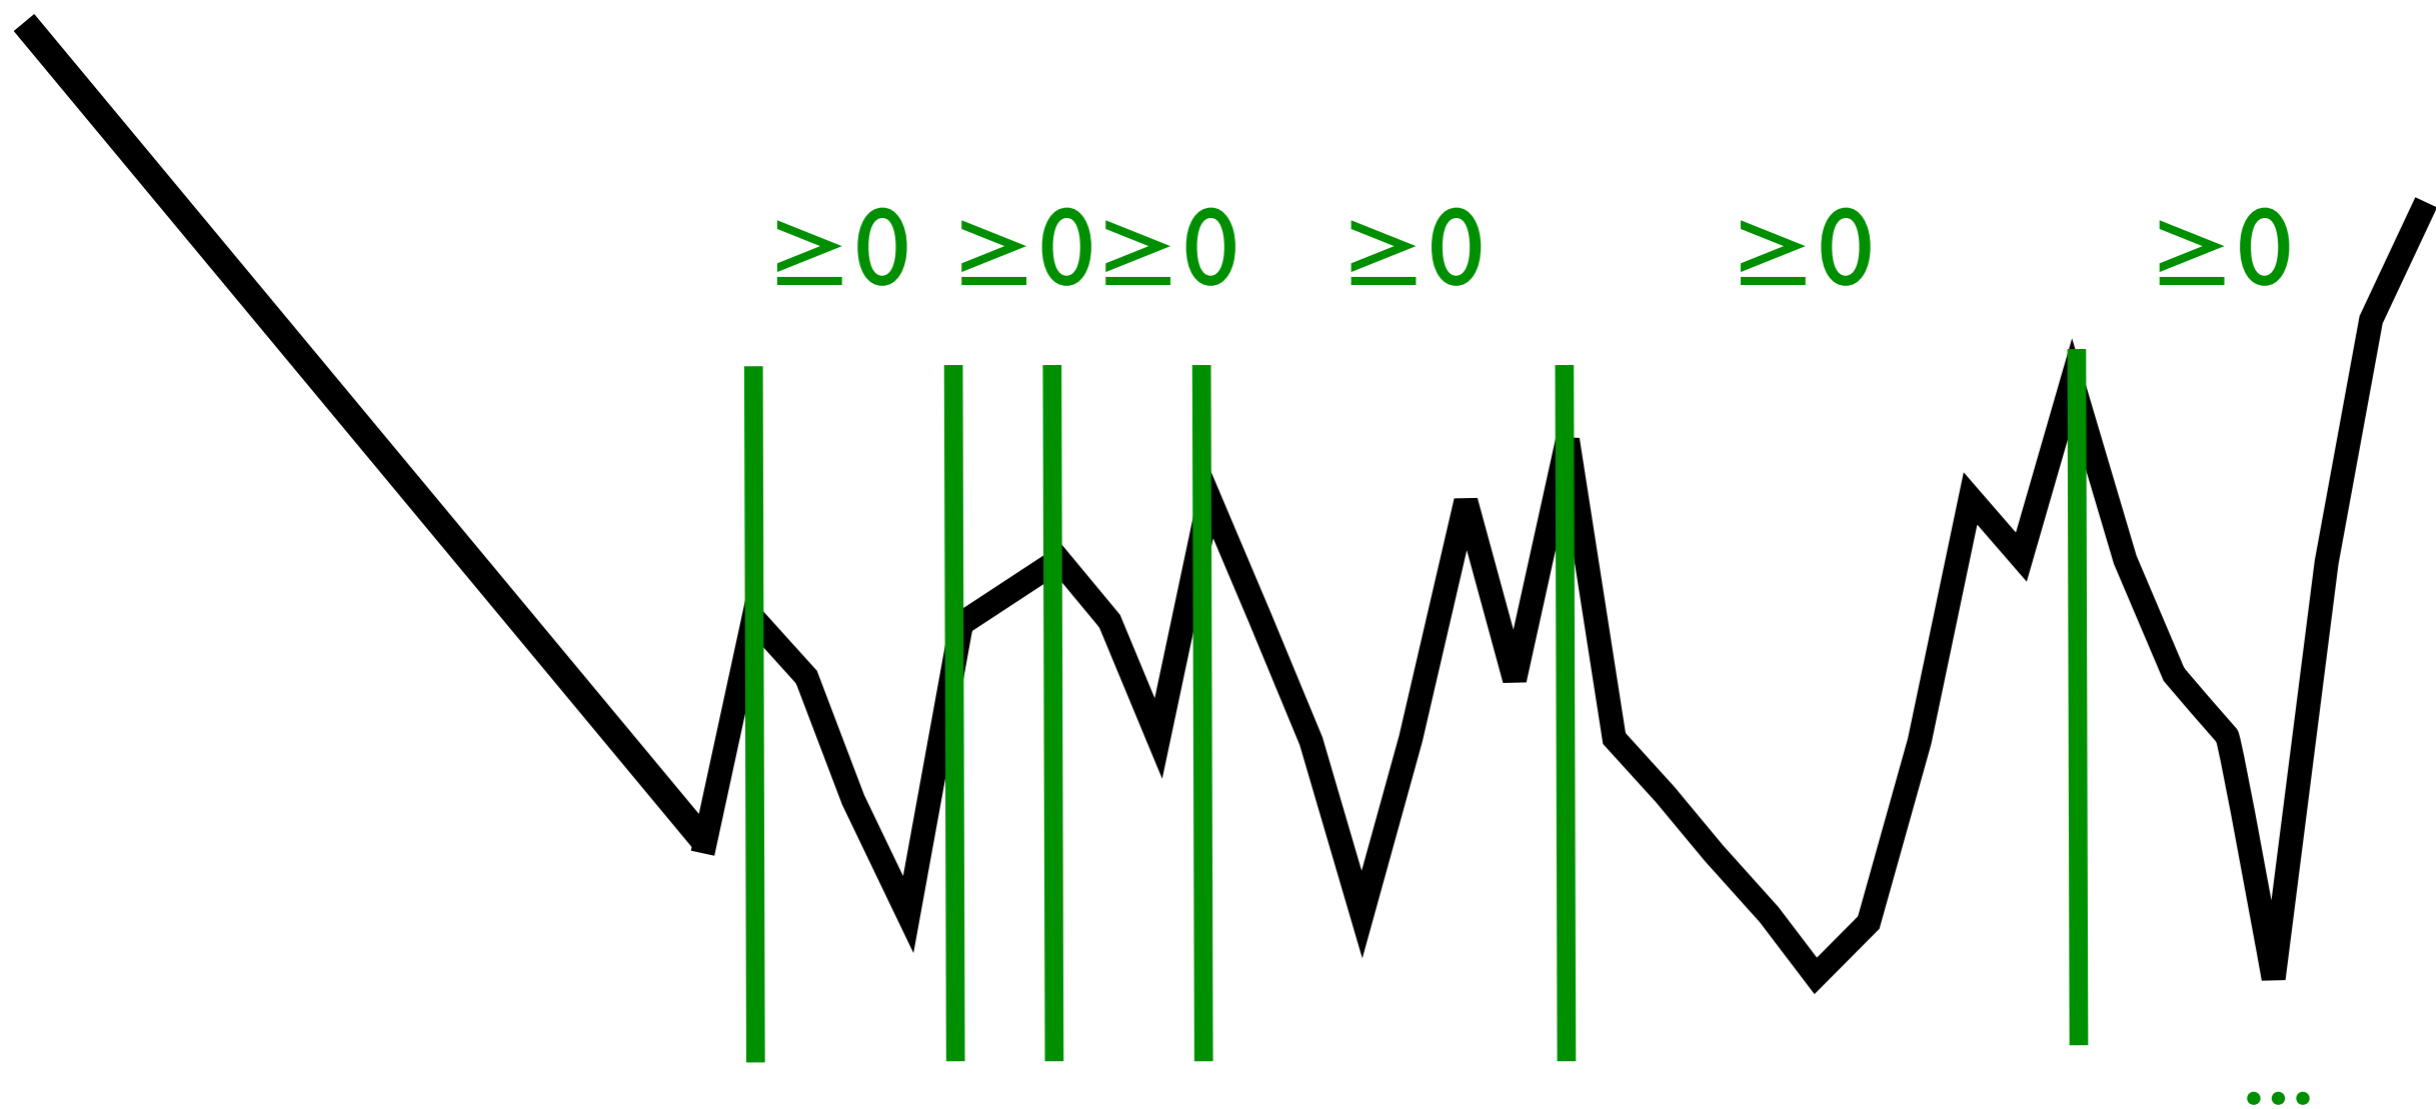
Relation between MP/TP and W.O.

If the answer to one of window mean-payoff problems is YES, then the answer to the mean-payoff threshold problem for threshold zero is also YES.



Relation between MP/TP and W.O.

If the answer to one of window mean-payoff problems is YES, then the answer to the mean-payoff threshold problem for threshold zero is also YES.

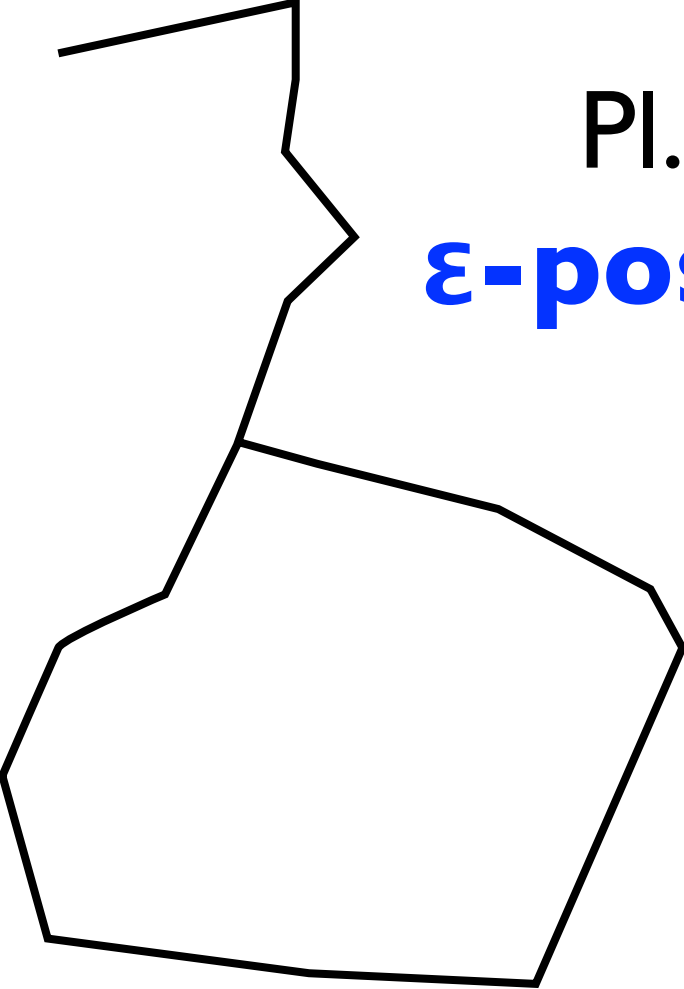


Relation between MP/TP and W.O.

*If there exists $\varepsilon > 0$ s.t. the answer to the mean-payoff threshold problem for threshold ε is YES, then the answer to the **BW** problem is also YES.*

Relation between MP/TP and W.O.

*If there exists $\varepsilon > 0$ s.t. the answer to the mean-payoff threshold problem for threshold ε is YES, then the answer to the **BW** problem is also YES.*



Pl. I can force
 ε -positive cycles



Pl. I can win the **DFW**(n) objective
for n **large enough**
i.e. $n = (|S| - 1) \cdot (1 + |S| \cdot W)$

Proof uses cycle decomposition of outcomes

Relation between MP/TP and W.O.

Theorem

- 1. If the answer to the one of window mean-payoff problems is YES, then the answer to the mean-payoff threshold problem for threshold zero is also YES.*
- 2. If there exists $\varepsilon > 0$ s.t. the answer to the mean-payoff threshold problem for threshold ε is YES, then the answer to the **BW** problem is also YES.*

\Rightarrow Window objectives can be seen as ε -approximations of mean-payoff for large enough windows

Solving **FixedW**(n)
for 1 dim.

Algorithm for **FixedW**(n)

Idea:

Follow the structure of definition
+
recurse on subgames

Algorithm for **FixedW(n)**


1) compute $W_1 = \{ s \mid s \models \langle \langle 1 \rangle \rangle \mathbf{GW}(n) \}$

i.e. the set of states from which Pl. 1 can force a positive sum within n steps (dynamic programming) $\mathbf{O}(n \cdot |G| \cdot \log(W))$

Algorithm for **FixedW(n)**

1) compute $W_1 = \{ s \mid s \models \langle \langle I \rangle \rangle \text{ GW}(n) \}$


i.e. the set of states from which Pl. 1 can force a positive sum within n steps (dynamic programming) $O(n \cdot |G| \cdot \log(W))$



Algorithm for **FixedW(n)**

1) compute $W_1 = \{ s \mid s \models \langle \langle 1 \rangle \rangle \text{ GW}(n) \}$

i.e. the set of states from which Pl. 1 can force a positive sum within n steps (dynamic programming) $O(n \cdot |G| \cdot \log(W))$




Algorithm for FixedW(n)

1) compute $W_1 = \{ s \mid s \models \langle\langle I \rangle\rangle \mathbf{GW}(n) \}$

i.e. the set of states from which Pl. 1 can force a positive sum within n steps (dynamic programming) $\mathbf{O}(n \cdot |G| \cdot \log(W))$

2) compute $W_2 = \{ s \mid s \models \langle\langle I \rangle\rangle \square W_1 \}$ winning for $\mathbf{DFW}(n)$

*i.e. the set of states from which Pl. 1 can win the **direct window objective** for size n . Complexity: $\mathbf{O}(|G|)$*




Algorithm for FixedW(n)

1) compute $W_1 = \{ s \mid s \models \langle\langle I \rangle\rangle \mathbf{GW}(n) \}$

i.e. the set of states from which Pl. 1 can force a positive sum within n steps (dynamic programming) $\mathbf{O}(n \cdot |G| \cdot \log(W))$

2) compute $W_2 = \{ s \mid s \models \langle\langle I \rangle\rangle \square W_1 \}$ winning for $\mathbf{DFW}(n)$

i.e. the set of states from which Pl. 1 can win the direct window objective for size n . Complexity: $\mathbf{O}(|G|)$




Algorithm for FixedW(n)

1) compute $W_1 = \{ s \mid s \models \langle\langle I \rangle\rangle \mathbf{GW}(n) \}$

i.e. the set of states from which Pl. 1 can force a positive sum within n steps (dynamic programming) $\mathbf{O}(n \cdot |G| \cdot \log(W))$

2) compute $W_2 = \{ s \mid s \models \langle\langle I \rangle\rangle \square W_1 \}$ winning for $\mathbf{DFW}(n)$

i.e. the set of states from which Pl. 1 can win the direct window objective for size n . Complexity: $\mathbf{O}(|G|)$



Algorithm for **FixedW(n)**

1) compute $W_1 = \{ s \mid s \models \langle\langle I \rangle\rangle \mathbf{GW}(n) \}$

i.e. the set of states from which Pl. 1 can force a positive sum within n steps (dynamic programming) $\mathbf{O}(n \cdot |G| \cdot \log(W))$

2) compute $W_2 = \{ s \mid s \models \langle\langle I \rangle\rangle \square W_1 \}$

i.e. the set of states from which Pl. 1 can win the direct window objective for size n . Complexity: $\mathbf{O}(|G|)$

3) compute $W_3 = \{ s \mid s \models \langle\langle I \rangle\rangle \diamond W_2 \}$ winning for **FixedW(n)**

*i.e. the **attractor** of W_2 is clearly winning for Pl. 1. Complexity: $\mathbf{O}(|G|)$*

Those states are **winning for Player 1**

So they should be **avoided at all cost by player 2**

Algorithm for FixedW(n)

1) compute $W_1 = \{ s \mid s \models \langle\langle I \rangle\rangle \mathbf{GW}(n) \}$

i.e. the set of states from which Pl. 1 can force a positive sum within n steps (dynamic programming) $\mathbf{O}(n \cdot |G| \cdot \log(W))$

2) compute $W_2 = \{ s \mid s \models \langle\langle I \rangle\rangle \square W_1 \}$

i.e. the set of states from which Pl. 1 can win the direct window objective for size n . Complexity: $\mathbf{O}(|G|)$

3) compute $W_3 = \{ s \mid s \models \langle\langle I \rangle\rangle \diamond W_2 \}$

i.e. the attractor of W_2 is clearly winning for Pl. 1. Complexity: $\mathbf{O}(|G|)$

4) Pl.2 should avoid W_3 at all cost: **recurse** on $S \setminus W_3$

The number of recursive calls is bounded by $\mathbf{O}(|G|)$

Overall complexity: *bounded by $\mathbf{O}(|G|^3 \cdot n \cdot \log(W))$*

Solving **FixedW**(n)

Theorem

In two-player one-dimension games:

*(a) the fixed arbitrary window mean-payoff problem is decidable in time **polynomial** in the size of the game **and** the window size*

*(b) the fixed polynomial-size window mean-payoff problem is **P-complete***

*(c) Both players **require memory**, and memory of size **linear** in the size of the game and the size of the window is sufficient*

Main Results

	one-dimension			k -dimension		
	complexity	\mathcal{P}_1 mem.	\mathcal{P}_2 mem.	complexity	\mathcal{P}_1 mem.	\mathcal{P}_2 mem.
$\underline{\text{MP}} / \overline{\text{MP}}$	$\text{NP} \cap \text{coNP}$	mem-less		$\text{coNP-c.} / \text{NP} \cap \text{coNP}$	infinite	mem-less
$\underline{\text{TP}} / \overline{\text{TP}}$	$\text{NP} \cap \text{coNP}$	mem-less		undec. (Thm. 1)	-	-
WMP: fixed polynomial window	P-c. (Thm. 2)	mem. req. $\leq \text{linear}(S \cdot l_{\max})$ (Thm. 2)		PSPACE-h. (Thm. 4)	exponential (Thm. 4)	
WMP: fixed arbitrary window	P ($ G , l_{\max}$) (Thm. 2)			EXP-easy (Thm. 4)		
WMP: bounded window problem	NP \cap coNP (Thm. 3)	mem-less (Thm. 3)	infinite (Thm. 3)	NPR-h. (Thm. 5)	-	-

More details in
<http://arxiv.org/abs/1209.3234>

The Complexity of Multi-Mean-Payoff and Multi-Energy Games^{*,**}

Yaron Velner¹, Krishnendu Chatterjee², Laurent Doyen³, Thomas A. Henzinger², Alexander Rabinovich¹, and Jean-François Raskin⁴

¹ The Blavatnik School of Computer Science, Tel Aviv University, Israel
² IST Austria (Institute of Science and Technology Austria)
³ LSV, ENS Cachan & CNRS, France
⁴ Département d'Informatique, Université Libre de Bruxelles (U.L.B.)

Abstract. In mean-payoff games, the objective of the protagonist is to ensure that the limit average of an infinite sequence of numeric weights is nonnegative. In energy games, the objective is to ensure that the running sum of weights is always nonnegative. Multi-mean-payoff and multi-energy games replace individual weights by tuples, and the limit average (resp. running sum) of each coordinate must be (resp. remain) nonnegative. These games have applications in the synthesis of resource-bounded processes with multiple resources. We prove the finite-memory determinacy of multi-energy games and show the inter-reducibility of multi-mean-payoff and multi-energy games for finite-memory strategies: while the previously best known upper bound was EXPSpace, and no lower bound was known, we give an optimal coNP-complete bound. For memoryless strategies, we show that the problem of deciding the existence of a winning strategy for the protagonist is NP-complete. Finally we present the first solution of multi-mean-payoff games with infinite-memory strategies. We show that multi-mean-payoff games with mean-payoff-sup objectives can be decided in $NP \cap coNP$, whereas multi-mean-payoff games with mean-payoff-inf objectives are coNP-complete.

Keywords: *Games on graphs; mean-payoff objectives; energy objectives; multi-dimensional objectives.*

1 Introduction

Graph games and multi-objectives. Two-player games on graphs are central in many applications in computer science. For example, in the synthesis problem, implementations are obtained from winning strategies in games with a qualitative specification [22, 21, 1]. In these applications, the specification determines which player wins, and the game is played in a natural model.

arXiv:1209.3234v1 [cs.GT] 14 Sep 2012

and in

Looking at Mean-Payoff and Total-Payoff through Windows

Krishnendu Chatterjee^{1,*}, Laurent Doyen², Mickael Randour^{3,†}, and Jean-François Raskin^{4,‡}

¹ IST Austria (Institute of Science and Technology Austria)

² LSV - ENS Cachan, France

³ Computer Science Department, Université de Mons (UMONS), Belgium

⁴ Département d'Informatique, Université Libre de Bruxelles (U.L.B.), Belgium

Abstract. We consider two-player games played on weighted directed graphs with mean-payoff and total-payoff objectives, which are two classical quantitative objectives. While for single dimensional objectives all results for mean-payoff and total-payoff coincide, we show that in contrast to multi-dimensional mean-payoff games that are known to be coNP-complete, multi-dimensional total-payoff games are undecidable. We introduce conservative approximations of these objectives, where the payoff is considered over a local finite window sliding along a play, instead of the whole play. For single dimension, we show that (i) if the window size is polynomial, then the problem can be solved in polynomial time, and (ii) the existence of a bounded window can be decided in $NP \cap coNP$, and is at least as hard as solving mean-payoff games. For multiple dimensions, we show that (i) the problem with fixed window size is EXPTIME-complete, and (ii) there is no primitive-recursive algorithm to decide the existence of a bounded window.

1 Introduction

Mean-payoff and total-payoff games. Two-player mean-payoff and total-payoff games are played on finite weighted directed graphs (in which every edge has an integer weight) with two types of vertices: in player-1 vertices, player 1 chooses the successor vertex from the set of outgoing edges; in player-2 vertices, player 2 does likewise. The game results in an infinite path (called a *play*) through the graph. The mean-payoff (resp. total-payoff) value of a play is the long-run average (resp. sum) of the edge-weights along the

see on arXiv:

<http://arxiv.org/abs/1302.4248>

Conclusion

- ▶ **Quantitative games:** energy games (+generalizations by Larsen, Bouyer,...), mean-payoff games, total-payoff games, ... is an active research area, part of a larger program...
- ▶ **From quality to quantity:** broad effort in order to lift boolean verification/synthesis to quantitative verification/synthesis, e.g.:
 - ▶ quantitative languages (Henzinger et al.) def. by **weighted automata**
- ▶ In this talk, we have shown:
 - ▶ MP games can be extended to multi-dim.
 - ▶ Space for alternative objectives: e.g. **window objectives**
 - ▶ “approximations” of MP/TP
 - ▶ natural definitions and interesting algorithms