The Power of Counting Steps in Quantitative Games

Sougata Bose
University of Liverpool, UK

Rasmus Ibsen-Jensen
University of Liverpool, UK

David Purser
University of Liverpool, UK

Patrick Totzke
University of Liverpool, UK

Pierre Vandenhove
LaBRI, Université de Bordeaux, France

Abstract

We study deterministic games of infinite duration played on graphs and focus on the strategy complexity of quantitative objectives. Such games are known to admit optimal memoryless strategies over finite graphs, but require infinite-memory strategies in general over infinite graphs.

We provide new lower and upper bounds for the strategy complexity of *mean-payoff* and *total-payoff* objectives over infinite graphs, focusing on whether *step-counter strategies* (sometimes called *Markov strategies*) suffice to implement winning strategies. In particular, we show that over finitely branching arenas, three variants of lim sup mean-payoff and total-payoff objectives admit winning strategies that are based either on a step counter or on a step counter and an additional bit of memory. Conversely, we show that for certain lim inf total-payoff objectives, strategies resorting to a step counter and finite memory are not sufficient. For step-counter strategies, this settles the case of all classical quantitative objectives up to the second level of the Borel hierarchy.

2012 ACM Subject Classification Theory of computation → Automata over infinite objects

Keywords and phrases Games on graphs, Markov strategies, quantitative objectives, infinite-state systems

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2024.13

Related Version *Full Version*: `https://arxiv.org/abs/2406.17482` [3]

Funding *P. Vandenhove* is funded by ANR project *G4S* (ANR-21-CE48-0010-01). *S. Bose* and *P. Totzke* acknowledge support from the EPSRC (EP/V025848/1, EP/X042596/1). Collaboration initiated via Royal Society International Exchanges grant (IES\R1\201211).

1 Introduction

Two-player (zero-sum, turn-based, perfect-information) games on graphs are an established formalism in formal verification, especially for *reactive synthesis* [1, 13]. They are used to model the interaction between a system, trying to satisfy a given *specification*, against an uncontrollable environment, assumed to act antagonistically as a worst case. We can model the system and its environment as two opposing players, called *Player 1* and *Player 2* respectively, who move a token through the graph of possible system configurations (called the *arena*). The specification is modelled as a winning condition (called *objective* henceforth), which is a set of all those interactions that the system player deems acceptable. The main algorithmic task when using this approach for formal verification is *solving* such games:

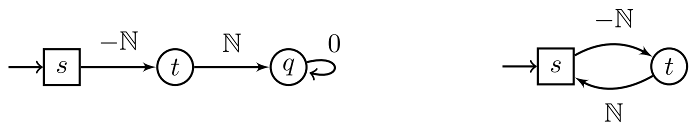

(a) The arena $\mathcal{A}_1$.

(b) The arena $\mathcal{A}'_1$.

Figure 1 Arenas implementing the "match the number" game. Circles designate vertices controlled by Player 1 and squares designate Player 2. The edge labels indicate that for every $i \in \mathbb{N}$ there is a distinct edge with weight $-i$ from s to t, and $+i$ from t to q or from t to s. For $\mathcal{A}_1$, consider the objective "sum of weights exceeds 0". Player 1 can always match and thus win, but needs unbounded memory. The arena $\mathcal{A}'_1$ shows a repeated version for the lim sup *mean*-payoff objective.

given an arena, an objective, and an initial vertex, decide whether the system player has a *winning strategy*, which corresponds to a controller for the system that guarantees that the specification holds no matter the behaviour of the environment. Additionally, reactive synthesis aims to *synthesise* (compute a representation of) a winning strategy if one exists.

Strategy complexity. To synthesise winning strategies, it is useful to know what kind of resources "suffice", i.e., are needed to implement a winning strategy, should one exist. This naturally depends on the model used for the interaction (the size and topology of the arena) and on the specification (the type of objective and whether probabilistic or absolute guarantees are required). We assume that strategies make decisions based on some internal memory, that stores and updates an abstraction of the past play.

The simplest strategies are those that are *memoryless*, meaning they base their decisions solely on the current arena vertex. Games on finite arenas where memoryless strategies are sufficient to win can usually be solved in NP ∩ coNP [28] and winning strategies effectively synthesised. This is true for *parity*, *discounted-payoff* [31], *mean-payoff* [11], and *total-payoff* [8, 16] objectives. Even beyond finite graphs, memoryless strategies may suffice in more general contexts, such as for parity objectives over arenas of arbitrary cardinality [12, 33], or *discounted-payoff objectives* over finitely branching arenas [26, Corollary 2.1].[1] For concurrent (stochastic) *reachability* games on finite arenas, memoryless strategies also suffice [2, 21].

Generally more powerful than memoryless strategies are *finite-memory* strategies, which refer to strategies that can be implemented with a finite-state (Mealy) machine. A canonical class of languages over infinite words, and standard for defining objectives in games, are the ω-regular languages [30, 17]. One of the celebrated related results about reactive synthesis is the *finite-memory determinacy* of ω-regular games [6, 30, 18], which means that if there is a winning strategy in a game on a finite arena and with an ω-regular objective, there is one that can be implemented with a simple finite-state machine (whose size can be bounded). This implies that games with ω-regular objectives can be solved and that strategies can be synthesised, since it bounds the search space for winning strategies. Remarkably, the existence of winning finite-memory strategies for ω-regular games even holds over arbitrary infinite arenas [33]. When finite-memory strategies are sufficient, one of the main questions is usually to *minimise* their size, i.e., to find winning strategies with as few memory states as possible [10, 7, 9, 5, 4].

[1] Thus we consider the strategy complexity in discounted-payoff games as settled for the setting we consider. On infinitely branching arenas, step-counter strategies are insufficient (see Figure 1a).

Already very simple games require infinite memory to win. This especially holds for quantitative objectives, which ask that the aggregate of individual edge weights along a play exceeds some threshold. For instance, consider a game where the environment picks a number and then the controller has to pick a larger one (see Figure 1a). In order to win, Player 1 has to remember the (per se unbounded) initial challenge and no finite memory structure would be sufficient to do so. This objective is not ω-regular since it is built upon an infinite alphabet. We seek to understand for different classes of games, what kind of infinite-memory structures are sufficient for winning strategies.

A natural, arguably the simplest, type of infinite memory structure is a *step counter*: it only remembers how many steps have elapsed since the start of the game. The availability of such a counter is a reasonable assumption for practical applications, as most embedded devices have access to the current time, which suffices when each step takes a fixed amount of time. A *step-counter strategy* is one that, in addition to the current arena vertex, has access to the number of steps elapsed. Notice that in the game in Figure 1a, a step counter does not provide any relevant information (every path to vertex t has length one). Therefore, step-counter strategies do not suffice for Player 1. An important ingredient for these counterexamples is that the underlying arena is infinitely branching (and uses arbitrary weights). For many classes of games on *finitely* branching arenas, strategies based on a step counter and additional finite memory are close to being the simplest kinds of strategies sufficient to win. Examples are especially prevalent in stochastic games. For instance, in the "Big Match" (a concurrent mean-payoff game on a finite arena), neither a step counter nor finite memory is sufficient to play ε-optimally, yet a step counter *together with* one bit is [19]. The same is true for the "Bad Match", which can be presented as a Büchi (repeated reachability) game [23, 32, 22]. This upper bound holds generally for concurrent Büchi games on finite arenas [22].

Quantitative objectives. Objectives based on numerical weights are commonly called *quantitative objectives*. These are defined using *quantitative payoff functions*, which combine any finite sequence of weights into an aggregate number. The three more common ones are the discounted-payoff [31], mean-payoff [15, 11], and total-payoff functions [14, 8]. Every payoff function induces four variants of objectives, depending on whether we consider the $\limsup$ or $\liminf$, and on whether we ask that the limit is larger or strictly larger than a threshold. For total payoff, it is also relevant to distinguish the use of real values or ∞ as a threshold. We give an example to describe informally how we denote such objectives: $\overline{\mathsf{MP}}_{\geq 0}$ refers to the set of infinite sequences of rational numbers that achieve a value ≥ 0 for the $\limsup$ variant (the line is above MP) of the mean-payoff function (specified by letters MP). Over infinite arenas, the four variants are not equivalent and infinite-memory strategies are needed for at least one of the players (see [29, Example 8.10.2] and [27]).

To study the strategy complexity for different quantitative objectives, we classify them according to which level of the *Borel hierarchy* they belong to (which also ensures that the games we consider are determined [24]). In the first level of the hierarchy lie the *open* and *closed* objectives (i.e., the sets respectively in $\mathbf{\Sigma}_1^0$ and $\mathbf{\Pi}_1^0$), for which there exist recent characterisations of the sufficient memory structures over finite or infinite arenas [9, 5]. We build on this to establish upper bounds for more complex objectives. All variants of mean-payoff and total-payoff objectives are on the second or third level of the Borel hierarchy. Ohlmann and Skrzypczak [27] study objectives through their topological properties and provide a characterisation of the *prefix-independent* $\mathbf{\Sigma}_2^0$ objectives for which memoryless strategies suffice for Player 1 over arbitrary arenas. They show in particular that memoryless strategies suffice for Player 1 for the quantitative objectives $\underline{\mathsf{MP}}_{>0}$ and $\underline{\mathsf{TP}}_{>-\infty}$, even over infinitely branching arenas. Over stochastic games, quantitative (in particular lim inf mean-payoff) objectives on infinite arenas generally do not have (ε-)optimal strategies based on a step counter, even for finitely branching Markov decision processes [25].

Our contributions. We settle the strategy complexity over infinite, deterministic games for the mean-payoff and total-payoff objectives up to the second level of the Borel hierarchy. In particular, we show for which of these, step-counter strategies are sufficient for Player 1. Our upper bounds all allow for arenas with arbitrary weights, while our strongest lower bounds only use weights -1, 0, and 1. Our results are as follows and summarised in Table 1.

- For $\underline{\mathsf{TP}}_{>0}$ and $\underline{\mathsf{TP}}_{\geq 0}$, strategies using a step counter and an arbitrary amount of finite memory do not suffice, even over acyclic finitely branching arenas (Theorem 10, Section 3). The proof rules out finite-memory structures using an application of the *infinite Ramsey theorem* to allow Player 2 to stay winning in a particular infinite arena regardless of the finite-memory structure of Player 1.
- In Section 5, we provide a sufficient condition for when step-counter strategies suffice over finitely branching arenas for prefix-independent objectives in $\mathbf{\Pi}_2^0$, i.e, countable intersections of open sub-objectives (Theorem 16). This implies in particular that step-counter strategies do suffice for $\overline{\mathsf{MP}}_{\geq 0}$ and $\overline{\mathsf{TP}}_{=+\infty}$ (Corollary 17), which is tight in the sense that finite-memory strategies do not suffice for these objectives, even over acyclic finitely branching arenas (Lemma 4). The proof uses carefully constructed expanding "bubbles", so that within each consecutive bubble, Player 1 can satisfy the next open sub-objective. The step counter is used to determine the current bubble.
- In Section 6, we show that for $\overline{\mathsf{TP}}_{\geq 0}$, which is not prefix-independent, strategies using a step-counter and one additional bit of memory suffice (Theorem 20). This is tight in that neither finite-memory strategies nor step-counter strategies suffice, even over acyclic finitely branching arenas (Lemmas 4 and 5). The proof similarly employs bubbles, but an additional bit is needed to keep track of whether a "sub-objective" has been achieved in the current bubble and then switches to stay in the winning region.

Structure. We define the various notions used throughout the paper in Section 2. Section 3 is dedicated to all lower bounds on the strategy complexity of the various objectives, culminating in a lower bound for $\underline{\mathsf{TP}}_{\geq 0}$. Section 4 is devoted to recalling useful results on open and closed objectives, upon which the following sections build. Section 5 proves a sufficient condition for the sufficiency of step-counter strategies for prefix-independent $\mathbf{\Pi}_2^0$ objectives. Section 6 proves an upper bound on the strategy complexity of $\overline{\mathsf{TP}}_{\geq 0}$.

Due to space constraints, some proofs are omitted from this conference version. Complete details for all proofs can be found in the extended version [3].

2 Preliminaries

Given a set X, we write X^* for the set of finite words on X, X^+ for the set of non-empty finite words on X, and X^ω for the set of infinite words on X. For $w \in X^*$, we write $|w|$ for the length of w. For $w \in X^\omega$ and $j \in \mathbb{N}$, we write $w_{\leq j}$ for the finite prefix of length j of w.

Games. We study two-player zero-sum *games*, each given by an *arena* and an *objective*, as defined below. We refer to the two opposing players as Player 1 and Player 2.

Table 1 Results for quantitative objectives up to the second level of the Borel hierarchy for finitely branching arenas. *SC* refers to *step counter*, and *FM* refers to *finite memory*.

Obj.	Description	Class	Strategy complexity
$\underline{\mathsf{MP}}_{>0}$	$\bigcup_{m\geq 1}\bigcup_{i\geq 1}\bigcap_{j\geq i}\{w \mid \mathsf{MP}(w_{\leq j}) \geq \frac{1}{m}\}$	$\mathbf{\Sigma}_2^0$	Memoryless (even over infinitely branching arenas) [27]
$\underline{\mathsf{TP}}_{>-\infty}$	$\bigcup_{m\geq 1}\bigcup_{i\geq 1}\bigcap_{j\geq i}\{w \mid \mathsf{TP}(w_{\leq j}) \geq -m\}$	$\mathbf{\Sigma}_2^0$	
$\underline{\mathsf{TP}}_{>0}$	$\bigcup_{m\geq 1}\bigcup_{i\geq 1}\bigcap_{j\geq i}\{w \mid \mathsf{TP}(w_{\leq j}) \geq \frac{1}{m}\}$	$\mathbf{\Sigma}_2^0$	SC + FM insufficient (Theorem 10)
$\overline{\mathsf{MP}}_{\geq 0}$	$\bigcap_{m\geq 1}\bigcap_{i\geq 1}\bigcup_{j\geq i}\{w \mid \mathsf{MP}(w_{\leq j}) \geq \frac{-1}{m}\}$	$\mathbf{\Pi}_2^0$	SC sufficient (Corollary 17) FM insufficient (Lemma 4)
$\overline{\mathsf{TP}}_{=+\infty}$	$\bigcap_{m\geq 1}\bigcap_{i\geq 1}\bigcup_{j\geq i}\{w \mid \mathsf{TP}(w_{\leq j}) \geq m\}$	$\mathbf{\Pi}_2^0$	
$\overline{\mathsf{TP}}_{\geq 0}$	$\bigcap_{m\geq 1}\bigcap_{i\geq 1}\bigcup_{j\geq i}\{w \mid \mathsf{TP}(w_{\leq j}) \geq \frac{-1}{m}\}$	$\mathbf{\Pi}_2^0$	SC + 1-bit sufficient (Theorem 20) FM insufficient (Lemma 4) SC insufficient (Lemma 5)

An *arena* is a directed graph with two kinds of vertices where edges are labelled by an element of C, a non-empty set of *colours*. Formally, an arena is a tuple $\mathcal{A} = (V, V_1, V_2, E)$ where $V = V_1 \cup V_2$ is a non-empty set of *vertices*, V_1 and V_2 are disjoint, and $E \subseteq V \times C \times V$ is a set of labelled *edges*. Vertices in V_1 and V_2 are respectively controlled by Player 1 and Player 2, which will appear clearly when we define strategies below. We require that for every vertex $v \in V$, there is an edge $(v, c, v') \in E$ (arenas are "non-blocking"). For $e = (v, c, v')$, we write $\mathsf{from}(e)$ for v, $\mathsf{col}(e)$ for c, and $\mathsf{to}(e)$ for v'. An arena is *finite* if V is finite, and *finitely branching* if for every $v \in V$, the set $\{e \in E \mid \mathsf{from}(e) = v\}$ is finite.

A *history* is a finite sequence $h = e_1 \dots e_n \in E^*$ of edges such that for $i \in \{1, \dots, n-1\}$, $\mathsf{to}(e_i) = \mathsf{from}(e_{i+1})$. We write $\mathsf{from}(h)$ for $\mathsf{from}(e_1)$, $\mathsf{to}(h)$ for $\mathsf{to}(e_n)$, and $\mathsf{col}(h)$ for the sequence $\mathsf{col}(e_1)\dots\mathsf{col}(e_n) \in C^*$. For convenience, we assume that for every vertex v, there is a distinct *empty history* λ_v such that $\mathsf{from}(\lambda_v) = \mathsf{to}(\lambda_v) = v$. The set of histories of $\mathcal{A}$ is denoted as $\mathsf{hists}(\mathcal{A})$. For $p \in \{1, 2\}$, we write $\mathsf{hists}_p(\mathcal{A})$ for the set of histories h such that $\mathsf{to}(h) \in V_p$. A *play* is an infinite sequence of edges $\rho = e_1e_2 \dots \in E^\omega$ such that for $i \geq 1$, $\mathsf{to}(e_i) = \mathsf{from}(e_{i+1})$. We write $\mathsf{from}(\rho)$ for $\mathsf{from}(e_1)$ and $\mathsf{col}(\rho)$ for $\mathsf{col}(e_1)\mathsf{col}(e_2)\dots \in C^\omega$. A history h (resp. a play ρ) is said to be *from* v if $v = \mathsf{from}(h)$ (resp. $v = \mathsf{from}(\rho)$).

An *objective* (sometimes called a *winning condition* in the literature) is a set $O \subseteq C^\omega$. An objective O is *prefix-independent* if for all $w \in C^*$, $w' \in C^\omega$, $ww' \in O$ if and only if $w' \in O$.

Strategies. A *strategy of Player p on $\mathcal{A}$* is a function $\sigma\colon \mathsf{hists}_p(\mathcal{A}) \to E$ such that for all $h \in \mathsf{hists}_p(\mathcal{A})$, $\mathsf{from}(\sigma(h)) = \mathsf{to}(h)$. A play $\rho = e_1e_2\dots$ is *consistent with a strategy σ of Player p* if for all finite prefixes h of ρ such that $\mathsf{to}(h) \in V_p$, $\sigma(h) = e_{|h|+1}$. A strategy σ of Player 1 is *winning for objective O from a vertex v* if all plays from v consistent with σ induce a sequence of colours in O. For a fixed objective, the set of vertices of an arena $\mathcal{A}$ from which a winning strategy for Player 1 exists is called the *winning region of Player 1 on $\mathcal{A}$* and is denoted $W_{\mathcal{A},1}$. A strategy σ of Player 1 is *uniformly winning for objective O in $\mathcal{A}$* if σ is winning from every vertex of the winning region of $\mathcal{A}$.

A *memory structure for an arena* $\mathcal{A} = (V, V_1, V_2, E)$ is a tuple $\mathcal{M} = (M, m_0, \delta)$ where M is a set of *memory states*, $m_0 \in M$ is an *initial state*, and $\delta\colon M \times E \to M$ is a *memory update function*. We extend δ to a function $\delta^*\colon M \times E^* \to M$ in a natural way. A memory structure $\mathcal{M}$ is *finite* if M is finite. A strategy σ of Player p on $\mathcal{A}$ is *based on* $\mathcal{M}$ if there exists a function $f\colon V_p \times M \to E$ such that, for all $h \in \mathsf{hists}_p(\mathcal{A})$, $\sigma(h) = f(\mathsf{to}(h), \delta^*(m_0, h))$. We will abusively assume that a strategy based on a memory structure is this function f.

A *memoryless strategy* is a strategy based on a memory structure with a single memory state. A 1-*bit strategy* is a strategy based on a memory structure with two memory states. A *step counter* is a memory structure $\mathcal{S} = (\mathbb{N}, 0, (s, e) \mapsto s + 1)$ that simply counts the number of steps already elapsed in a game. A strategy σ of Player p on $\mathcal{A}$ is a *step-counter strategy* if σ is based on a step counter; in other words, if there is a function $f : V_p \times \mathbb{N} \to E$ such that $\sigma(h) = f(\mathsf{to}(h), |h|)$. This means that σ only considers the current vertex and the number of steps elapsed to make its decisions. Step-counter strategies are sometimes called "Markov strategies" [32, 20].

A *step-counter and finite-memory structure* is a memory structure with state space $M = \mathbb{N} \times \{0, \dots, K-1\}$, initial state $(0, 0)$, and a transition function δ such that $\delta((s, m), e) = (s+1, \delta'((s, m), e))$ for some function $\delta' : M \times E \to \{0, \dots, K-1\}$. Notice that a step counter corresponds to the special case of a step-counter and finite-memory structure with $K = 1$. A *step-counter* + 1-*bit strategy* is a strategy based on a step-counter and finite-memory structure with $K = 2$.

We say that a kind of strategies *suffices for objective O over a class of arenas* if, for all arenas in this class, from all vertices of her winning region, Player 1 has a winning strategy of this kind. We say that a kind of strategies *suffices uniformly for objective O over a class of arenas* if, for all arenas in this class, Player 1 has a uniformly winning strategy of this kind.

For an arena $\mathcal{A} = (V, V_1, V_2, E)$ and a memory structure $\mathcal{M} = (M, m_0, \delta)$, we write $\mathcal{A} \otimes \mathcal{M}$ for the *product between $\mathcal{A}$ and $\mathcal{M}$*. It is the arena (V', V'_1, V'_2, E') such that $V' = V \times M$, $V'_1 = V_1 \times M$, $V'_2 = V_2 \times M$, and $E' = \{((v, m), c, (v', \delta(m, e))) \mid e = (v, c, v') \in E, m \in M\}$. Observe that Player 1 has a winning strategy based on $\mathcal{M}$ from a vertex v in an arena $\mathcal{A}$ if and only if Player 1 has a winning memoryless strategy from vertex (v, m_0) in $\mathcal{A} \otimes \mathcal{M}$.

To simplify reasonings over specific arenas, we show that step counters do not have any use when the arena already *encodes the step count*.

▶ **Lemma 1.** *Let $\mathcal{A} = (V, V_1, V_2, E)$ be an arena, and $v_0 \in V$ be an initial vertex. Assume that for each pair of histories h_1, h_2 from v_0 to some $v \in V$, we have $|h_1| = |h_2|$ (i.e., the arena already "encodes the step count from v_0"). Then, a step-counter and finite-memory strategy with K states of finite memory can be simulated from v_0 by a strategy with only K states of finite memory.*

Proof. By hypothesis on $\mathcal{A}$, there exists $n_v \in \mathbb{N}$ the length of any history from v_0 to v. Let $\sigma' : V_1 \times \mathbb{N} \times M \to E$ be a step-counter and finite-memory strategy with $M = \{0, \dots, K-1\}$, with finite-memory update function $\delta' : M \times E \to \{0, \dots, K-1\}$. Let $\mathcal{M} = (M, 0, \delta)$ be the memory structure with $\delta(m, e) = \delta'((n_{\mathsf{from}(e)}, m), e)$. By construction, the strategy $\sigma : V_1 \times M \to E$ such that $\sigma(v, m) = \sigma'(v, n_v, m)$ behaves exactly like σ' from v_0. ◀

Quantitative objectives. We consider classical quantitative objectives: mean-payoff and total-payoff objectives, as defined below. Let $C \subseteq \mathbb{Q}$ (when colours are rational numbers, we often refer to them as *weights*). For a finite word $w = c_1 \dots c_{|w|} \in C^*$, define $\mathsf{TP}(w) = \sum_{i=1}^{|w|} c_i$ for the *total payoff* of the word, i.e., the sum of the weights it contains. Further, when $|w| \geq 1$, let $\mathsf{MP}(w) = \mathsf{TP}(w)/|w|$ denote the *mean payoff* of the word w, i.e., the mean of the weights it contains. We extend any such aggregate function $X : C^* \to \mathbb{R}$ to infinite words by taking limits: for $w \in C^\omega$, we define $\overline{X}(w) = \limsup_j X(w_{\leq j})$ and $\underline{X}(w) = \liminf_j X(w_{\leq j})$. Fixing a binary relation $\triangleright \subseteq \mathbb{R}^2$ and threshold $r \in \mathbb{Q} \cup \{-\infty, \infty\}$, this naturally defines objectives $\overline{X}_{\triangleright r} = \{w \in C^\omega \mid \overline{X}(w) \triangleright r\}$ and $\underline{X}_{\triangleright r} = \{w \in C^\omega \mid \underline{X}(w) \triangleright r\}$.

In particular, we are interested in the limit infimum/supremum objectives for total and mean payoff.[2] We consider the mean-payoff variants with threshold $r \in \mathbb{Q}$, and the total-payoff variants with $r \in \mathbb{Q} \cup \{-\infty, +\infty\}$. Note that all four mean-payoff objectives and all four total-payoff objectives with ∞ threshold are prefix-independent, but the four total-payoff objectives with threshold in $\mathbb{Q}$ are not prefix-independent.

▶ Remark 2. Our results are generally stated for threshold $r = 0$. This is without loss of generality since the results deal with large classes of arenas, and little modifications to the arenas allow to reduce from an arbitrary rational threshold to threshold 0. ⌟

Topology of objectives. For $w \in C^*$, we write $wC^\omega = \{ww' \mid w' \in C^\omega\}$ for the objective containing all infinite words that start with w (it is sometimes called the *cylinder* or *cone* of w). An objective O is *open* if there is a set $A \subseteq C^*$ such that $O = \bigcup_{w \in A} wC^\omega$. For an open objective O, we say that a finite word $w \in C^*$ *already satisfies* O if $wC^\omega \subseteq O$. If an objective is open, then by definition, any infinite word it contains has a finite prefix that already satisfies it. An objective is *closed* if it is the complement of an open set.

Open and closed objectives are at the first level of the *Borel hierarchy*; the set of open (resp. closed) objectives is denoted $\mathbf{\Sigma}^0_1$ (resp. $\mathbf{\Pi}^0_1$). For $i > 1$, we can define $\mathbf{\Sigma}^0_i$ as all the countable unions of sets in $\mathbf{\Pi}^0_{i-1}$, and $\mathbf{\Pi}^0_i$ as all the countable intersections of sets in $\mathbf{\Sigma}^0_{i-1}$. All the objectives considered in this paper lie in the first three levels of this hierarchy, and we focus on those on the second level.

3 Lower bounds

We provide lower bounds on the size/structure of the memory to build winning strategies, focusing on objectives $\overline{\mathsf{MP}}_{\geq 0}$, $\overline{\mathsf{TP}}_{=+\infty}$, $\overline{\mathsf{TP}}_{\geq 0}$, and $\underline{\mathsf{TP}}_{>0}$, which are the four objectives on the second level of Borel hierarchy for which we want to establish whether step-counters strategies suffice. We mention where our constructions directly work for further objectives.

All lower bounds are based on the simple idea that one player chooses some number and the other must match it. We first observe that on infinitely branching arenas with arbitrary weights, neither finite memory nor a step counter, nor both together, is sufficient. The proof uses the arenas from Figure 1, discussed informally in Section 1 (the missing proofs in this section are available in [3, Appendix A]).

▶ **Lemma 3.** *Over infinitely branching arenas with arbitrary weights, step-counter and finite-memory strategies are not sufficient for Player 1 for objectives* $\overline{\mathsf{MP}}_{>0}$, $\overline{\mathsf{MP}}_{\geq 0}$, $\overline{\mathsf{TP}}_{=+\infty}$, $\underline{\mathsf{TP}}_{>0}$, $\underline{\mathsf{TP}}_{\geq 0}$, $\overline{\mathsf{TP}}_{>0}$ *and* $\overline{\mathsf{TP}}_{\geq 0}$.

We now establish lower bounds over finitely branching arenas. Firstly, the example $\mathcal{A}'_1$ can be made finitely branching and acyclic, as depicted in Figure 2. The resulting arena, $\mathcal{A}_2$, simply unfolds $\mathcal{A}'_1$ so that any edge $(s, -j, t)$ is replaced by a finite path $s^i_0 \to \cdots \to s^i_j \to t^i_0$, and similarly for the responses. This construction works as long as one can discourage (i.e., make losing) the choice to stay on the infinite intermediate chain of vertices and not moving on to a vertex controlled by the opponent. Here, this is achieved by using weights 1 on the chains of Player 2 and weights -1 on the chains of Player 1, which are then compensated by weights twice as large. In practice, edges with weights $i \in \mathbb{N}$ (resp. $-i \in -\mathbb{N}$) can be

[2] We only consider objectives where the threshold is a *lower* bound ($\triangleright \in \{>, \geq\}$); each variant with *upper* bound behaves like a variant with lower bound when we replace each weight c in arenas with its additive inverse $-c$ and switch the sup/inf (for instance, $\overline{\mathsf{MP}}_{<r}$ behaves like $\underline{\mathsf{MP}}_{>r}$ when we invert the weights).

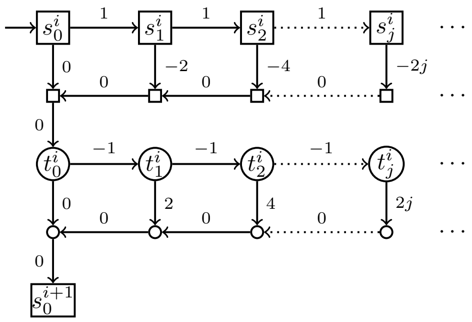

■ **Figure 2** The arena $\mathcal{A}_2$ is acyclic and every vertex has finite in- and out-degree. We recall that circles are controlled by Player 1 and squares by Player 2.

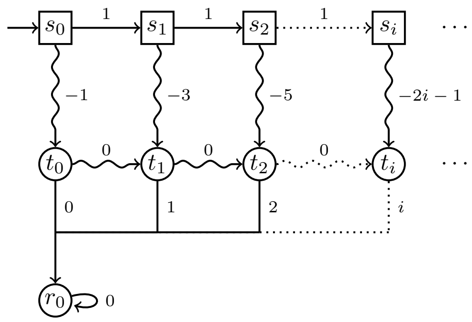

■ **Figure 3** The arena $\mathcal{A}_3$. Arrows $s_i \overset{-2i-1}{\rightsquigarrow} t_i$ are shorthand for paths of length $2i+1$ with edge weights -1, and $t_i \overset{0}{\rightsquigarrow} t_{i+1}$ are shorthand for paths of length 3 with edge weights 0.

replaced by chains of i weights 1 (resp. i weights -1). This allows to obtain lower bounds on the lim sup objectives. The fact that finite-memory strategies are insufficient for variants of the mean-payoff objectives over finitely branching arenas was already discussed in [29, Example 8.10.2] and [27]; we rephrase it here for completeness.

▶ **Lemma 4.** *Over finitely branching arenas, finite-memory strategies are not sufficient for Player 1 for objectives* $\overline{\mathsf{MP}}_{>0}$, $\overline{\mathsf{MP}}_{\geq 0}$, $\overline{\mathsf{TP}}_{=+\infty}$, $\overline{\mathsf{TP}}_{>0}$, *and* $\overline{\mathsf{TP}}_{\geq 0}$.

Notice that although finite memory is insufficient for Player 1 in $\mathcal{A}_2$, a step counter allows her to deduce an upper bound on the previous choice of Player 2 and is therefore sufficient. Indeed, since $\mathcal{A}_2$ is finitely branching and every round starts in a unique initial vertex for that round, Player 1 can (over) estimate that all steps of the history so far were spent by her opponent's choice (steps between s_0^i up to some s_j^i and then leading directly to t_0^{i+1}).

In order to construct an arena in which no step-counter strategy is sufficient, we obfuscate possible histories leading to Player 1's choices by making them the same length (see Figure 3).

▶ **Lemma 5.** *Consider the arena $\mathcal{A}_3$ depicted in Figure 3. Player 1 has a winning strategy, but no winning step-counter strategy for objectives* $\underline{\mathsf{TP}}_{>0}$, $\underline{\mathsf{TP}}_{\geq 0}$, $\overline{\mathsf{TP}}_{>0}$, *and* $\overline{\mathsf{TP}}_{\geq 0}$. *Hence, over finitely branching arenas, step-counter strategies are not sufficient for Player 1 for objectives* $\underline{\mathsf{TP}}_{>0}$, $\underline{\mathsf{TP}}_{\geq 0}$, $\overline{\mathsf{TP}}_{>0}$, *and* $\overline{\mathsf{TP}}_{\geq 0}$.

Proof. Player 1 only makes relevant choices at vertices t_i, and the choice is whether to *delay* (move to t_{i+1}) or *exit* (move to r_0). A winning (finite-memory) strategy for all mentioned objectives is to delay twice and then exit. Indeed, any history leading to t_i has total payoff of at least $-i-1$. By delaying twice and then exiting, Player 1 guarantees that the sink vertex r_0 is reached and the total payoff collected on the way is at least 1.

Conversely, any strategy σ of Player 1 that is based solely on a step counter cannot distinguish histories leading to the same vertex t_i. Let us assume that σ does not choose to avoid r_0 indefinitely, as doing so would result in a negative total payoff, which is losing for her. Then there is at least one vertex t_i from which the strategy exits. Player 2 can exploit this by going there via s_i. The resulting play has a negative total payoff. ◀

We now extend the previous examples to show that even access to both a step counter and finite memory is not sufficient for Player 1. The construction below is stated for the total-payoff objective $\underline{\mathsf{TP}}_{\geq 0}$, and also works for $\underline{\mathsf{TP}}_{>0}$. The main idea is to require Player 1 to delay going to r_0 more than a constant number of times, as dictated by Player 2's initial move.

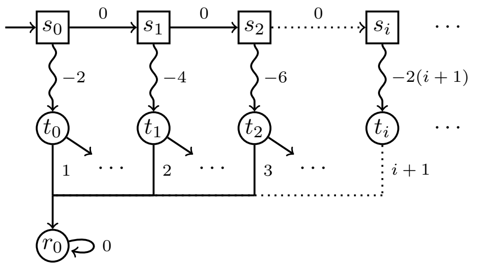

Figure 4 The arena $\mathcal{A}_4$. Arrows $s_i \overset{-2(i+1)}{\rightsquigarrow} t_i$ are shorthand for paths of length $2i+3$ with total payoff $-2(i+1)$. From a vertex t_i, Player 1 either exits to r_0 or moves to the gadget in Figure 5.

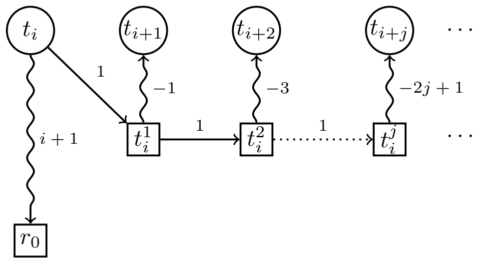

Figure 5 The delay gadget from vertex t_i in arena $\mathcal{A}_4$. The arrows from t_i^j to t_{i+j} are shorthand for paths of length $2j$ and payoff $-2j+1$.

▶ **Definition 6.** *Let $\mathcal{A}_4$ be the arena from Figure 4. It has a similar high-level structure to $\mathcal{A}_3$ with different weights, and with more complex gadgets (Figure 5) between vertices t_i. At each vertex t_i, Player 1 decides between two actions:*

1. *to* exit *to r_0 and gain payoff $i+1$ by doing so, or*
2. *to* delay *to some vertex t_{i+j} where $j > 0$ is chosen by Player 2, and gain payoff $-j+1$.*

Notice that, after Player 2 moved down from vertex s_k, Player 1 can (only) win by delaying at least $k+1$ times (which we show in Lemma 8). We will show that the gadgets allow Player 2 to confuse any strategy of Player 1 that is only based on a step counter and finite memory. Without them, the current vertex t_i together with finite extra memory would allow Player 1 to approximate how many delays she has chosen so far and therefore allow her to win with a finite-memory strategy.[3]

A simple counting argument shows that all paths from s_0 to a vertex t_k have the same length (proof in [3, Appendix A]). By Lemma 1, it implies that a step counter is useless in $\mathcal{A}_4$.

▶ **Lemma 7.** *For every t_k in arena $\mathcal{A}_4$, all paths from s_0 to t_k have the same length.*

The following lemma will be used to argue that Player 1 wins, albeit with infinite memory.

▶ **Lemma 8.** *From a vertex t_i, if Player 2 does not stay forever in a gadget, the strategy σ_k of Player 1 that enters the delay gadget exactly $k \in \mathbb{N}$ times achieves a total payoff of exactly $i+k+1$ in r_0.*

Proof. Assume that Player 2 never stays forever in a gadget (which would be winning for Player 1 for all quantitative objectives considered). The total payoff on the path from t_i to the next vertex t_{i+j} is $-j+1$. Suppose Player 1 delays k times and let $j(1), j(2), \dots, j(k)$ be the lengths of the intermediate paths through gadgets, as chosen by Player 2. That is, the play ends up in vertex t_{i+l} for $l = \sum_{c=1}^{k} j(c)$ and has gained payoff $\sum_{c=1}^{k} ((-j(c)+1)) = -l+k$. After k delays, exiting to r_0 from vertex t_{i+l} gives an immediate payoff of $i+l+1$. The total payoff from t_i to r_0 is thus $(-l+k)+(i+l+1) = i+k+1$. ◀

[3] The idea would be to partition t_i's into (growing) intervals, so that each interval is picked so large that it is safe to exit from any vertex after the interval if the play entered a vertex before or at the start of that interval. A winning strategy is then to keep on delaying to t_i's until vertices in three different intervals have been seen, and then exit. This requires 3 memory states to remember the interval changes.

▶ **Lemma 9.** *Consider the game played on $\mathcal{A}_4$. Then, from vertex s_0,*
1. *Player 1 wins for objective $\underline{\mathsf{TP}}_{\geq 0}$;*
2. *every step-counter and finite-memory strategy of Player 1 is losing for $\underline{\mathsf{TP}}_{\geq 0}$.*

Proof. For point (1), let σ be the Player 1 strategy that, upon observing history $s_0 \xrightarrow{*} s_k \to t_k$, switches to the finite-memory strategy σ_{k+1} from the previous lemma (delay $k+1$ times and then exit). Consider any play consistent with this strategy σ. Either Player 2 never moves to a vertex t_k, and then the total payoff is 0, which is winning for Player 1 for $\underline{\mathsf{TP}}_{\geq 0}$. Otherwise, a vertex t_k is reached (and accordingly, the payoff until reaching it is $-2(k+1)$). Using σ_{k+1}, Player 1 guarantees a lim inf total payoff of at least 0 on any continuation: either Player 2 never leaves some gadget and the total payoff is $+\infty$, or Player 1 exits to r_0 after $k+1$ delays, which adds $k+(k+1)+1 = 2(k+1)$ to the total payoff by Lemma 8. In this second case, the total payoff is therefore $-2(k+1)+2(k+1) = 0$.

For point (2), by Lemmas 1 and 7, it suffices to show that every finite-memory strategy of Player 1 is losing. Consider now any such strategy σ_1 of Player 1 with memory of size $K \in \mathbb{N}$ and memory update function δ. We will show that there exists a strategy σ_2 for Player 2 that is winning against σ_1. Player 2's strategy is determined by 1) the initial choice of t_j it visits and 2) which vertex t_{i+j} to select in the gadgets (Figure 5) when Player 1 delays from vertex t_i. We show the existence of suitable choices by employing an argument based on the infinite Ramsey theorem, as follows.

First, δ defines naturally, for any history $h \in E^*$, a function $\delta_h \colon M \to M$ that specifies how the memory is updated when observing this history (formally, $\delta_h(m) = \delta^*(m, h)$). Further, for every $i \geq 0$ there is a function $f_i \colon M \to \{0,1\}$ that describes for which memory states the strategy σ_1 chooses to delay or exit from t_i (formally, $f_i(m)$ equals 1 if $\sigma_1(t_i, m) = (t_i, i+1, r_0)$, and 0 otherwise). Since $|M| = K \in \mathbb{N}$, there are only finitely many distinct such functions f_i and δ_h. Consider now the edge-labelled graph G consisting of all vertices $t_i, i \geq 0$, and where for any two $i, j \in \mathbb{N}$, the edge between t_i and t_{i+j} is labelled by the pair (f_i, δ_h) where $h = t_i \to t_i^1 \to \cdots \to t_i^j \to t_{i+j}$ is the history through the delay gadget in $\mathcal{A}_4$.

Recall the infinite Ramsey theorem: If one labels all edges of the complete (undirected and countably infinite) graph with finitely many colours, then there exists an infinite monochromatic subgraph. Applying this to our graph G yields an infinite subgraph, say with vertices $t_{\ell(i)}$ identified by $\ell \colon \mathbb{N} \to \mathbb{N}$, where all edges have the same label. W.l.o.g., assume that $\ell(0) \geq K$ and $\ell(i+1) > \ell(i)+1$ for all $i \geq 0$. Based on this, the strategy σ_2 of Player 2 will 1) initially move to $t_{\ell(0)}$ and 2) whenever Player 1 chooses to delay from $t_{\ell(i)}$ then Player 2 moves to vertex $t_{\ell(i+1)}$. Now consider the play ρ consistent with both strategies σ_1 and σ_2. There are two cases. Either along this play Player 1 chooses to exit from some vertex $t_{\ell(j)}, j < K$, or not. If she exits too early (after delaying only $j < K$ times), then the total payoff after exiting is exactly $-2(\ell(0)+1) + (\ell(0)+j+1) = -\ell(0)+j-1$ by Lemma 8, which is < 0 as $\ell(0) \geq K > j$. Hence, the play is won by Player 2. Alternatively, if along the play, Player 1 delays at least K times then, by the pigeonhole principle, there is at least one memory mode that she revisits. More precisely, the play visits vertices $t_{\ell(i)}$ and $t_{\ell(j)}$, $i < j \leq K$ in the same memory mode. Recall that the functions $f_{\ell(i)}$ are all identical for $i \geq 0$. It follows that the play will continue visiting vertices $t_{\ell(k)}, k \geq 0$ only and never exit to r_0. Finally, observe that in any delay gadget from a vertex $t_{\ell(i)}$, the path to vertex $t_{\ell(i+1)}$ has total payoff of $1 - (\ell(i+1) - \ell(i))$. Consequently, the infinite play ρ that visits all $t_{\ell(i)}$ will be such that $\underline{\mathsf{TP}}(\rho) = -\infty$ and is losing for Player 1 for $\underline{\mathsf{TP}}_{\geq 0}$. ◀

▶ **Theorem 10.** *Strategies based on a step counter and finite memory are not sufficient for Player 1 in games with finitely branching arenas and objectives $\underline{\mathsf{TP}}_{\geq 0}$ or $\underline{\mathsf{TP}}_{>0}$.*

Proof. For $\underline{\mathsf{TP}}_{\geq 0}$ this follows directly from Lemma 9. For $\underline{\mathsf{TP}}_{>0}$, just extend the arena by a new initial vertex s_{-1} with sole outgoing edge $s_{-1} \xrightarrow{1} s_0$ to ensure that the play in which Player 2 never moves to a vertex t_i is won by Player 1. ◀

4 Open objectives

The quantitative objectives defined in Section 2 all belong to the second or third level of the Borel hierarchy, and the strategy complexity of such objectives is not yet well understood. However, they use as building blocks objectives from the first level of the Borel hierarchy (i.e., open and closed objectives), for which there already exist characterisations of memory requirements. We recall some of these results for the memory structures that we study.

Step-monotonicity. Let $O \subseteq C^\omega$ be an objective. For two finite words $w_1, w_2 \in C^*$, we write $w_1 \preceq_O w_2$ if for all $w \in C^\omega$, $w_1 w \in O$ implies $w_2 w \in O$ (meaning that the winning continuations of w_1 are included in those of w_2). The relation $\preceq_O$ is a preorder and satisfies that for $w_1, w_2 \in C^*$ and $c \in C$, $w_1 \preceq_O w_2$ implies $w_1 c \preceq_O w_2 c$ (i.e., it is a "congruence"). We write $w_1 \prec_O w_2$ if $w_1 \preceq_O w_2$ but $w_2 \not\preceq_O w_1$. We say that two finite words $w_1, w_2 \in C^*$ are *comparable for* $\preceq_O$ if $w_1 \preceq_O w_2$ or $w_2 \preceq_O w_1$. We extend preorder $\preceq_O$ to histories: we write $h_1 \preceq_O h_2$ if $\mathsf{col}(h_1) \preceq_O \mathsf{col}(h_2)$.

We say that an objective O is *step-monotonic* if for any two finite words $w_1, w_2 \in C^*$ such that $|w_1| = |w_2|$, w_1 and w_2 are comparable for $\preceq_O$. In other words, for any two finite words that are read up to the same state of a step counter, one of the words must include at least the winning continuations of the other word. This is a specialisation of the $\mathcal{M}$*-strong-monotony* property [5] for the step-counter memory structure $\mathcal{M} = \mathcal{S}$.

▶ **Example 11.** Let $C = \{a, b\}$. The open objective $O = aaC^\omega \cup bbC^\omega$ is *not* step-monotonic, since for $w_1 = a$ and $w_2 = b$, we have that $|w_1| = |w_2|$, but w_1 and w_2 are not comparable for $\preceq_O$. Indeed, a^ω (resp. b^ω) is a winning continuation of w_1 but not w_2 (resp. w_2 but not w_1).

Now, let $C = \mathbb{Q}$ and $s \in \mathbb{N}$. The open objective $O_s = \{w \in C^\omega \mid \exists j \geq s, \mathsf{TP}(w_{\leq j}) \geq 0\}$ (containing all infinite words whose total payoff goes over 0 at some point after s steps) is step-monotonic. Indeed, consider two finite words $w_1, w_2 \in C^*$ such that $|w_1| = |w_2|$. If w_2 already satisfies O_s (i.e., $w_2 C^\omega \subseteq O_s$), then necessarily, $w_1 \preceq_{O_s} w_2$. Similarly, if w_1 already satisfies O_s, then $w_2 \preceq_{O_s} w_1$. When neither w_1 nor w_2 already satisfies O_s, they can be compared by their current total payoff: if $\mathsf{TP}(w_1) \leq \mathsf{TP}(w_2)$, then $w_1 \preceq_{O_s} w_2$. ⌟

▶ Remark 12. Variations of objective O_s are used as building blocks to define quantitative objectives (as can be seen in the descriptions in Table 1), and will be considered again later. An important remark is that $\preceq_{O_s}$ is not completely determined by the current total payoff of words. For instance, if $w_1 = -1, 0$ and $w_2 = 0, -100$, we have $w_1 \prec_{O_1} w_2$ even though $\mathsf{TP}(w_1) > \mathsf{TP}(w_2)$. The reason is that w_2 *already satisfies* O_1 after 1 step, and any continuation is therefore winning, despite the current total payoff being lower. ⌟

Step-counter strategies for open objectives. In general, the step-monotonicity property is necessary for the uniform sufficiency of step-counter strategies over finitely branching arenas (this is a specialisation of [5, Lemma 5.2] to the step-counter memory structure $\mathcal{S}$). However, the results of [5] do not yield a characterisation for open objectives in full generality. For the special case of the step-counter memory structure, we can actually show a converse: for open objectives, step-monotonicity implies that step-counter strategies suffice over finitely branching arenas. This is what we show over the next three lemmas (the missing proofs in this section are available in [3, Appendix B]).

First, a handy result about open objectives is that in a *finitely branching* arena, any winning strategy already satisfies the objective within a bounded number of steps.

▶ **Lemma 13.** *Let $O \subseteq C^\omega$ be an open objective, $\mathcal{A}$ be a finitely branching arena, and v_0 be an initial vertex in $\mathcal{A}$. If a strategy σ is winning from v_0 for O, then there is $s \in \mathbb{N}$ such that all histories h of length $\geq s$ consistent with σ already satisfy O, i.e., $\mathsf{col}(h)C^\omega \subseteq O$.*

Second, the following lemma shows that for step-monotonic objectives, step-counter strategies can be "locally not worse" than arbitrary strategies.

▶ **Lemma 14.** *Let $O \subseteq C^\omega$ be a step-monotonic objective. Let $\mathcal{A} = (V, V_1, V_2, E)$ be a finitely branching arena, $v_0 \in V$ be an initial vertex, and σ' be any strategy of Player 1 on $\mathcal{A}$. There is a step-counter strategy σ such that, for every history h from v_0 consistent with σ, there is a history h' from v_0 consistent with σ' such that $|h'| = |h|$, $\mathsf{to}(h') = \mathsf{to}(h)$, and $h' \preceq_O h$.*

The previous two lemmas imply that step-counter strategies suffice to win for open, step-monotonic objectives.

▶ **Corollary 15.** *Let $O \subseteq C^\omega$ be an open, step-monotonic objective. Step-counter strategies suffice for O over finitely branching arenas.*

Proof. Let $\mathcal{A}$ be a finitely branching arena. Let v_0 be a vertex from the winning region and σ' be an arbitrary winning strategy from v_0. By Lemma 13, using that O is open and $\mathcal{A}$ is finitely branching, for all histories h of length $\geq s$ consistent with σ', we have $\mathsf{col}(h)C^\omega \subseteq O$.

As O is step-monotonic, let σ be the step-counter strategy provided by Lemma 14. Every history h of length s from v_0 consistent with σ is at least as good (for $\preceq_O$) as a history h' of length s from v_0 consistent with σ'. Since h' only has winning continuations, so does h. Therefore, strategy σ is winning from v_0. ◀

5 Prefix-independent $\mathbf{\Pi}_2^0$ objectives

In this section, we show that step-counter strategies suffice for Player 1 for objectives $\overline{\mathsf{MP}}_{\geq 0}$ and $\overline{\mathsf{TP}}_{=+\infty}$. In fact, we give a sufficient condition for when step-counter strategies suffice for Player 1 in finitely branching games where the objectives are prefix-independent and in $\mathbf{\Pi}_2^0$.

Recall that an objective is in $\mathbf{\Pi}_2^0$ if it can be written as $\bigcap_{m\in\mathbb{N}} O_m$ for some open objectives O_m.

▶ **Theorem 16.** *Let $O = \bigcap_{m\in\mathbb{N}} O_m \subseteq C^\omega$ be a prefix-independent $\mathbf{\Pi}_2^0$ objective such that the objectives O_m are open and step-monotonic. Then, step-counter strategies suffice uniformly for O over finitely branching arenas.*

Proof. Let $\mathcal{A} = (V, V_1, V_2, E)$ be a finitely branching arena, and let $v_0 \in V$ be an initial vertex. Let $W_{\mathcal{A},1} \subseteq V$ be the winning region of $\mathcal{A}$ for O. We assume that v_0 is in the winning region $W_{\mathcal{A},1}$, and build a winning *step-counter* strategy from v_0.

We build a winning step-counter strategy $\sigma \colon V_1 \times \mathbb{N} \to E$ from v_0 by induction on parameter m used in the definition of $O = \bigcap_{m\in\mathbb{N}} O_m$. We consider the product arena $\mathcal{A} \otimes \mathcal{S}$, and fix a strategy for increasingly high step values. The inductive scheme is as follows: for every $m \in \mathbb{N}$, we fix σ on $V_1 \times \{0, \ldots, k_m - 1\}$ for some step bound $k_m \in \mathbb{N}$. We ensure that

- along all histories from v_0 consistent with σ of length at most k_m, the history does not leave $W_{\mathcal{A},1}$ (i.e., for all reachable (v, s), we have $v \in W_{\mathcal{A},1}$), and
- the open objectives $O_{m'}$ for $m' \leq m$ are already satisfied within k_m steps (i.e., any history of length k_m consistent with σ only has winning continuations for $O_{m'}$).

For the base case, we may assume that we start the induction at $m = -1$ with $k_{-1} = 0$ and $O_{-1} = C^\omega$. We indeed have that from $(v_0, 0)$, the winning region is not left within $k_{-1} = 0$ step and that the open objective O_{-1} is already satisfied.

Now, assume that for some $m \geq 0$, the above properties hold, so we have already fixed the moves of σ in $\mathcal{A} \otimes \mathcal{S}$ on $V_1 \times \{0, \ldots, k_m - 1\}$, yielding arena $(\mathcal{A} \otimes \mathcal{S})_m$. We first show that in arena $(\mathcal{A} \otimes \mathcal{S})_m$ the vertex $(v_0, 0)$ still belongs to the winning region. We have assumed by induction that the winning region $W_{\mathcal{A},1}$ is not left within k_m steps. This means that for all (v, k_m) reachable from $(v_0, 0)$ in $(\mathcal{A} \otimes \mathcal{S})_m$, v is in $W_{\mathcal{A},1}$. As O is prefix-independent, no matter the history from $(v_0, 0)$ to (v, k_m), there is still a winning strategy from (v, k_m) (recall that no choice for Player 1 has been fixed beyond step k_m). Hence, no matter how Player 2 plays in the first k_m steps, there is still a way to win for O from $(v_0, 0)$.

We therefore take an (arbitrary) winning strategy σ'_{m+1} of Player 1 from $(v_0, 0)$ in $(\mathcal{A} \otimes \mathcal{S})_m$. As σ'_{m+1} is winning for $O = \bigcap_{m \in \mathbb{N}} O_m$, σ'_{m+1} wins in particular for the open O_{m+1}. Since the arena is finitely branching and O_{m+1} is open, applying Lemma 13, there is $k'_{m+1} \in \mathbb{N}$ such that for all histories h' of length $\geq k'_{m+1}$ consistent with σ'_{m+1}, h' already satisfies O_{m+1} (i.e., $\mathsf{col}(h')C^\omega \subseteq O_{m+1}$). As O_{m+1} is step-monotonic, by Lemma 14, there is a step-counter strategy σ_{m+1} such that for every history h from $(v_0, 0)$ consistent with σ_{m+1}, there is a history h' from $(v_0, 0)$ consistent with σ'_{m+1} such that $|h'| = |h|$, $\mathsf{to}(h') = \mathsf{to}(h)$, and $h' \preceq_{O_{m+1}} h$.

To ensure that we fix at least one extra step of the strategy in the inductive step, let $k_{m+1} = \max\{k'_{m+1}, k_m + 1\}$. We extend the definition of σ to play the same moves as σ_{m+1} on $V_1 \times \{k_m, \ldots, k_{m+1} - 1\}$, which also defines $(\mathcal{A} \otimes \mathcal{S})_{m+1}$. We prove the two items of the inductive scheme.

First, σ still does not leave $W_{\mathcal{A},1}$ up to step k_{m+1}: indeed, for every history consistent with σ_{m+1}, there is a history consistent with σ'_{m+1} reaching the same vertex. Since σ'_{m+1} is winning and O is prefix-independent, no such vertex can be outside of the winning region.

Second, strategy σ then guarantees O_{m+1} within k_{m+1} steps: after k_{m+1} steps, every history consistent with σ_{m+1} is at least as good for $\preceq_{O_{m+1}}$ as a history of length k_{m+1} of σ'_{m+1}. But every history h' of length k_{m+1} consistent with σ' is such that $\mathsf{col}(h')C^\omega \subseteq O_{m+1}$, and therefore has only winning continuations.

This concludes the induction argument and shows the existence of a winning step-counter strategy from v_0 as we iterate this process for $m \to \infty$.

We now know that for any vertex from the winning region, there is a winning step-counter strategy. The existence of a *uniformly* winning step-counter strategy can be shown using prefix-independence of O; this part of the proof is standard and is detailed in [3, Appendix C]. ◀

This theorem applies to $\overline{\mathsf{MP}}_{\geq 0}$ and $\overline{\mathsf{TP}}_{=+\infty}$ (see [3, Appendix D] for a full proof).

▶ **Corollary 17.** *Step-counter strategies suffice uniformly for* $\overline{\mathsf{MP}}_{\geq 0}$ *and* $\overline{\mathsf{TP}}_{=+\infty}$.

To illustrate Theorem 16 further, we apply it to a non-quantitative objective.

▶ **Example 18.** Let C be at most countable and $O \subseteq C^\omega$ be the objective requiring that all colours are seen infinitely often (it is an intersection of *Büchi conditions*). Formally, $O = \bigcap_{c \in C} \bigcap_{i \geq 1} \bigcup_{j \geq i} \{w = c_1 c_2 \ldots \in C^\omega \mid c_j = c\}$. This objective is prefix-independent and in $\mathbf{\Pi}^0_2$: it is the countable intersection of the open, step-monotonic objectives $O_{c,i} = \bigcup_{j \geq i} \{w = c_1 c_2 \ldots \in C^\omega \mid c_j = c\}$. By Theorem 16, step-counter strategies suffice over finitely branching arenas for O. This result is relatively tight: finite-memory strategies do not suffice over finitely branching arenas when C is infinite, and step-counter strategies do not suffice over infinitely branching arenas when $|C| = 2$ (see [3, Appendix D] for details). ⌟

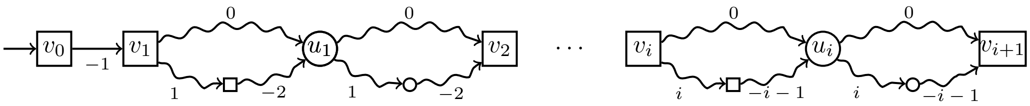

Figure 6 Arena $\mathcal{A}$ used in Example 19. Player 1 has a winning 1-bit strategy for $\overline{\mathsf{TP}}_{\geq 0}$, but no winning step-counter strategy.

6 A non-prefix-independent $\mathbf{\Pi}_2^0$ objective

In this section, we consider objective $\overline{\mathsf{TP}}_{\geq 0} = \bigcap_{m\geq 1}\bigcap_{i\geq 1}\bigcup_{j\geq i}\{w \in C^\omega \mid \mathsf{TP}(w_{\leq j}) \geq \frac{-1}{m}\}$ (in $\mathbf{\Pi}_2^0$). Its definition is very close to the one of $\overline{\mathsf{MP}}_{\geq 0}$ from the previous section, but one important difference is that it is not prefix-independent (for instance, $0^\omega \in \overline{\mathsf{TP}}_{\geq 0}$, but $-1, 0^\omega \notin \overline{\mathsf{TP}}_{\geq 0}$). Hence, Theorem 16 does not apply.

As argued in Lemma 5, it turns out that step-counter strategies do not suffice for $\overline{\mathsf{TP}}_{\geq 0}$, even over finitely branching arenas. We show a second example, only suited for this particular objective, illustrating more clearly the trade-off to consider to build simple winning strategies.

▶ **Example 19.** Consider the arena $\mathcal{A}$ in Figure 6. We assume that a play starts in v_0, hence reaching sum of weights -1 in v_1. We assume that a play is decomposed into rounds, where round i corresponds to the choice of Player 2 and Player 1 in v_i and u_i respectively. At each round i, Player 2 and then Player 1 choose either 0, or i followed by $-i-1$. As previously, we can assume that this arena only uses weights in $C = \{-1, 0, 1\}$, and that all histories from v_0 reaching the same vertex have the same length.

Player 1 has a winning strategy, consisting of playing "the opposite" of what Player 2 just played: if Player 2 played the sequence of 0 (resp. $i, -i-1$), then Player 1 replies with $i, -i-1$ (resp. the sequence of 0). This ensures that (i) the current sum of weights in v_i is exactly $-i$ (it starts at -1 in v_1 and decreases by 1 at each round), and (ii) the current sum of weights reaches exactly 0 once during each round, after i is played. This shows that this strategy is winning for $\overline{\mathsf{TP}}_{\geq 0}$. Such a strategy can be implemented with two memory states that simply remember the choice of Player 2 at each round.

As all histories leading to vertices u_i have the same length, a step-counter strategy cannot distinguish the choices of Player 2 (Lemma 1). Any step-counter strategy is losing:

- either Player 1 only plays 0, in which case Player 2 wins by only playing 0, thereby ensuring that the current sum of weights is -1 from v_1 onwards;
- or Player 1 plays $i, -i-1$ at some u_i. In this case, Player 2 wins by only playing $i, -i-1$. This means that the sum of weights decreases by at least 1 at every round, but decreases by 2 in round i. Hence, for $j \geq i$, the sum of weights at round j is at most $-j-1$. Such a sum can never go above 0 again when a player plays $j, -j-1$. ⌟

This example shows that in general, there is a trade-off between "obtaining a high value for a short time, to go above 0 temporarily" and "playing safe in order not to decrease the value too much". Two memory states sufficed: if the opponent just saw a high sum of weights (≥ 0), then we can play it safe temporarily; if the opponent played it safe, we may need to aim for a high value, even if the overall sum decreases. This reasoning generalises to all finitely branching arenas: in general, step-counter + 1-bit strategies suffice for $\overline{\mathsf{TP}}_{\geq 0}$.

▶ **Theorem 20.** *Step-counter + 1-bit strategies suffice for $\overline{\mathsf{TP}}_{\geq 0}$ over finitely branching arenas.*

We provide a proof sketch here (full proof in [3, Appendix E]). It follows the same scheme as the proof of Theorem 16, where we inductively fix choices for ever longer histories. However, we need to be more careful not to leave the winning region. As the objective is not prefix-independent, the winning region $W'_{\mathcal{A},1}$ is described not just by a set of vertices, but by pairs of a vertex and current total payoff (i.e., the current sum of weights), i.e, $W'_{\mathcal{A},1} \subseteq V \times \mathbb{Q}$.

We start with a lemma about the sufficiency of memoryless strategies to *stay* in this winning region. Staying in $W'_{\mathcal{A},1}$ is necessary but not sufficient to win for $\overline{\mathsf{TP}}_{\geq 0}$.

▶ **Lemma 21.** *Let $\mathcal{A} = (V, V_1, V_2, E)$ be a finitely branching arena. There exists a memoryless strategy σ_{safe} of Player 1 in $\mathcal{A}$ such that, for every $(v_0, r) \in W'_{\mathcal{A},1}$, σ_{safe} never leaves $W'_{\mathcal{A},1}$ from v_0 with initial weight value r.*

The following lemma is an analogue of Lemma 14, but ensures a stronger property with a more complex memory structure (using an extra bit). It says that locally, with a step-counter + 1-bit strategy, we can guarantee a high value temporarily while staying in the winning region $W'_{\mathcal{A},1}$, generalising the phenomenon of Example 19. The bit is used to aim for a high value (bit value 0) or stay in the winning region (bit value 1) by playing σ_{safe} from Lemma 21.

We use a rewriting of $\overline{\mathsf{TP}}_{\geq 0}$: observe that

$$\overline{\mathsf{TP}}_{\geq 0} = \bigcap_{m \geq 1} \bigcup_{j \geq m} \left\{ w \in C^\omega \mid \mathsf{TP}(w_{\leq j}) \geq \frac{-1}{m} \right\}, \tag{1}$$

where the variable m is used both for the $-\frac{1}{m}$ lower bound and for the m lower bound on the step count. Indeed, this also enforces that, for arbitrarily long prefixes, the current total payoff goes above values arbitrarily close to 0. For $m \geq 1$, let $O_m = \bigcup_{j \geq m} \{w \mid \mathsf{TP}(w_{\leq j}) \geq \frac{-1}{m}\}$ be the open set used in the definition of $\overline{\mathsf{TP}}_{\geq 0}$ in (1).

▶ **Lemma 22.** *Let $\mathcal{A} = (V, V_1, V_2, E)$ be an arena and $v_0 \in V$ be an initial vertex in the winning region of Player 1 for $\overline{\mathsf{TP}}_{\geq 0}$. For all $m \geq 1$, there exists a step-counter + 1-bit strategy σ_m such that σ_m is winning for O_m from v_0 and never leaves $W'_{\mathcal{A},1}$ (i.e., for all histories h from v_0 consistent with σ_m, $(\mathsf{to}(h), \mathsf{TP}(h)) \in W'_{\mathcal{A},1}$).*

The inductive scheme used in the proof of Theorem 20 is similar to that of Theorem 16, building a step-counter + 1-bit strategy $\sigma \colon V_1 \times \mathbb{N} \times \{0, 1\} \to E$.

For $\mathcal{M}$ a step-counter and 1-bit memory structure, consider the product arena $\mathcal{A}' = \mathcal{A} \otimes \mathcal{M}$ (in which the bit updates are not fixed yet, and will be fixed inductively). We have that $(v_0, (0, 0))$ is in the winning region of $\mathcal{A}'$. The inductive scheme is as follows: for infinitely many $m \in \mathbb{N}$, for some step bound $k_m \in \mathbb{N}$, we fix σ on $V_1 \times \{0, \dots, k_m - 1\} \times \{0, 1\}$, yielding arena $\mathcal{A}'_m$. Using Lemma 22, we ensure that

- along all histories h from v_0 consistent with σ of length at most k_m, $W'_{\mathcal{A},1}$ is not left, and
- the open objective O_m is already satisfied within k_m steps (i.e., any history of length k_m consistent with σ only has winning continuations for O_m).

Iterating this procedure defines a step-counter + 1-bit strategy σ that satisfies O_m for infinitely many $m \geq 1$. Since the sequence $(O_m)_{m \geq 1}$ is decreasing ($O_1 \supseteq O_2 \supseteq \dots$), we have that σ is winning for O_m for all $m \geq 1$. Hence, σ is winning for $\overline{\mathsf{TP}}_{\geq 0}$.

▶ Remark 23. Unlike for Theorem 16, the upper bound in this section does not apply *uniformly* in general (an arena illustrating this is in [3, Appendix E]). ⌟

▶ Remark 24. Over integer weights ($C \subseteq \mathbb{Z}$), $\overline{\mathsf{TP}}_{>0} = \overline{\mathsf{TP}}_{\geq 1} \in \mathbf{\Pi}^0_2$. As $\overline{\mathsf{TP}}_{\geq 1}$ behaves like $\overline{\mathsf{TP}}_{\geq 0}$ (Remark 2), the results from this section apply to $\overline{\mathsf{TP}}_{>0}$ over integer weights. Up to some scaling factor, this also applies to rational weights with bounded denominators. However, for general rational weights, $\overline{\mathsf{TP}}_{>0}$ can only be shown to be in $\mathbf{\Sigma}^0_3$, so the above does not apply. ⌟

7 Conclusion

We established whether step-counter strategies (possibly with finite memory) suffice for the objectives $\overline{\mathsf{MP}}_{\geq 0}$, $\underline{\mathsf{TP}}_{>0}$, $\underline{\mathsf{TP}}_{\geq 0}$, $\overline{\mathsf{TP}}_{=+\infty}$, and $\overline{\mathsf{TP}}_{\geq 0}$. We used the structure of these objectives as sets in the Borel hierarchy, and pinpointed the strategy complexity for all classical quantitative objectives on the second level of Borel hierarchy. This leaves open the cases of $\overline{\mathsf{MP}}_{>0}$, $\underline{\mathsf{MP}}_{\geq 0}$, $\overline{\mathsf{TP}}_{>0}$ (over $\mathbb{Q}$), and $\underline{\mathsf{TP}}_{=+\infty}$, all on the third level. The sufficiency of other less common infinite memory structures, such as *reward counters* [25], could also be investigated.

References

1 Roderick Bloem, Krishnendu Chatterjee, and Barbara Jobstmann. Graph games and reactive synthesis. In Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem, editors, *Handbook of Model Checking*, pages 921–962. Springer, 2018. doi:10.1007/978-3-319-10575-8_27.

2 Benjamin Bordais, Patricia Bouyer, and Stéphane Le Roux. Optimal strategies in concurrent reachability games. In Florin Manea and Alex Simpson, editors, *30th EACSL Annual Conference on Computer Science Logic, CSL 2022, February 14–19, 2022, Göttingen, Germany (Virtual Conference)*, volume 216 of *LIPIcs*, pages 7:1–7:17, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. doi:10.4230/LIPIcs.CSL.2022.7.

3 Sougata Bose, Rasmus Ibsen-Jensen, David Purser, Patrick Totzke, and Pierre Vandenhove. The power of counting steps in quantitative games. *CoRR*, abs/2406.17482, 2024. doi:10.48550/arXiv.2406.17482.

4 Sougata Bose, Rasmus Ibsen-Jensen, and Patrick Totzke. Bounded-memory strategies in partial-information games. In *ACM/IEEE Symposium on Logic in Computer Science*, LICS '24, New York, NY, USA, 2024. Association for Computing Machinery. doi:10.1145/3661814.3662096.

5 Patricia Bouyer, Nathanaël Fijalkow, Mickael Randour, and Pierre Vandenhove. How to play optimally for regular objectives? In Kousha Etessami, Uriel Feige, and Gabriele Puppis, editors, *50th International Colloquium on Automata, Languages, and Programming, ICALP 2023, July 10–14, 2023, Paderborn, Germany*, volume 261 of *LIPIcs*, pages 118:1–118:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2023. doi:10.4230/LIPIcs.ICALP.2023.118.

6 J. Richard Büchi and Lawrence H. Landweber. Definability in the monadic second-order theory of successor. *Journal of Symbolic Logic*, 34(2):166–170, 1969. doi:10.2307/2271090.

7 Antonio Casares. On the minimisation of transition-based Rabin automata and the chromatic memory requirements of Muller conditions. In Florin Manea and Alex Simpson, editors, *30th EACSL Annual Conference on Computer Science Logic, CSL 2022, February 14–19, 2022, Göttingen, Germany (Virtual Conference)*, volume 216 of *LIPIcs*, pages 12:1–12:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPICS.CSL.2022.12.

8 Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga. Resource interfaces. In Rajeev Alur and Insup Lee, editors, *Proceedings of the 3rd International Conference on Embedded Software, EMSOFT 2003, Philadelphia, PA, USA, October 13–15, 2003*, volume 2855 of *Lecture Notes in Computer Science*, pages 117–133. Springer, 2003. doi:10.1007/978-3-540-45212-6_9.

9 Thomas Colcombet, Nathanaël Fijalkow, and Florian Horn. Playing safe, ten years later. *Logical Methods in Computer Science*, 20(1), 2024. doi:10.46298/LMCS-20(1:10)2024.

10 Stefan Dziembowski, Marcin Jurdziński, and Igor Walukiewicz. How much memory is needed to win infinite games? In *Proceedings of the 12th Annual IEEE Symposium on Logic in Computer Science, LICS 1997, Warsaw, Poland, June 29 – July 2, 1997*, pages 99–110, 1997. doi:10.1109/.1997.614939.

11 Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, 1979. doi:10.1007/BF01768705.

12 E. Allen Emerson and Charanjit S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *Proceedings of the 32nd Annual Symposium on Foundations of Computer Science, FOCS 1991, San Juan, Puerto Rico, October, 1991*, pages 368–377. IEEE Computer Society, 1991. doi:10.1109/SFCS.1991.185392.

13 Nathanaël Fijalkow, Nathalie Bertrand, Patricia Bouyer-Decitre, Romain Brenguier, Arnaud Carayol, John Fearnley, Hugo Gimbert, Florian Horn, Rasmus Ibsen-Jensen, Nicolas Markey, Benjamin Monmege, Petr Novotný, Mickael Randour, Ocan Sankur, Sylvain Schmitz, Olivier Serre, and Mateusz Skomra. Games on graphs, 2023. arXiv:2305.10546.

14 Jerzy Filar and Koos Vrieze. *Competitive Markov Decision Processes*. Springer New York, 1996. URL: https://books.google.fr/books?id=21lcbnzDNwsC.

15 Dean Gillette. *Stochastic Games with Zero Stop Probabilities*, pages 179–188. Princeton University Press, Princeton, 1957. doi:10.1515/9781400882151-011.

16 Hugo Gimbert and Wiesław Zielonka. When can you play positionally? In Jiří Fiala, Václav Koubek, and Jan Kratochvíl, editors, *Proceedings of the 29th International Symposium on Mathematical Foundations of Computer Science, MFCS 2004, Prague, Czech Republic, August 22–27, 2004*, volume 3153 of *Lecture Notes in Computer Science*, pages 686–697. Springer, 2004. doi:10.1007/978-3-540-28629-5_53.

17 Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research [outcome of a Dagstuhl seminar, February 2001]*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002. doi:10.1007/3-540-36387-4.

18 Yuri Gurevich and Leo Harrington. Trees, automata, and games. In Harry R. Lewis, Barbara B. Simons, Walter A. Burkhard, and Lawrence H. Landweber, editors, *Proceedings of the 14th Annual ACM Symposium on Theory of Computing, STOC 1982, San Francisco, CA, USA, May 5–7, 1982*, pages 60–65. ACM, 1982. doi:10.1145/800070.802177.

19 Kristoffer Arnsfelt Hansen, Rasmus Ibsen-Jensen, and Abraham Neyman. The big match with a clock and a bit of memory. *Mathematics of Operations Research*, 48(1):419–432, 2023. doi:10.1287/moor.2022.1267.

20 Stefan Kiefer, Richard Mayr, Mahsa Shirmohammadi, and Patrick Totzke. Strategy complexity of parity objectives in countable MDPs. In Igor Konnov and Laura Kovács, editors, *Proceedings of the 31st International Conference on Concurrency Theory, CONCUR 2020, September 1–4, 2020, Vienna, Austria (Virtual Conference)*, volume 171 of *LIPIcs*, pages 39:1–39:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.CONCUR.2020.39.

21 Stefan Kiefer, Richard Mayr, Mahsa Shirmohammadi, and Patrick Totzke. Memoryless strategies in stochastic reachability games, 2024. To appear in Lecture Notes in Computer Science.

22 Stefan Kiefer, Richard Mayr, Mahsa Shirmohammadi, and Patrick Totzke. Strategy complexity of Büchi objectives in concurrent stochastic games, 2024. arXiv:2404.15483.

23 Ashok P. Maitra and William D. Sudderth. *Discrete Gambling and Stochastic Games*. Springer-Verlag, 1996.

24 Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975. URL: http://www.jstor.org/stable/1971035.

25 Richard Mayr and Eric Munday. Strategy complexity of point payoff, mean payoff and total payoff objectives in countable MDPs. *Logical Methods in Computer Science*, 19(1), 2023. doi:10.46298/LMCS-19(1:16)2023.

26 Pierre Ohlmann. *Monotonic graphs for parity and mean-payoff games*. PhD thesis, IRIF – Research Institute on the Foundations of Computer Science, 2021.

27 Pierre Ohlmann and Michał Skrzypczak. Positionality in Σ_2^0 and a completeness result. In Olaf Beyersdorff, Mamadou Moustapha Kanté, Orna Kupferman, and Daniel Lokshtanov, editors, *Proceedings of the 41st International Symposium on Theoretical Aspects of Computer Science, STACS 2024, March 12–14, 2024, Clermont-Ferrand, France*, volume 289 of *LIPIcs*, pages 54:1–54:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPICS.STACS.2024.54.

28 Anuj Puri. *Theory of Hybrid Systems and Discrete Event Systems.* PhD thesis, EECS Department, University of California, Berkeley, December 1995. URL: http://www2.eecs.berkeley.edu/Pubs/TechRpts/1995/2950.html.

29 Martin L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming.* Wiley Series in Probability and Statistics. Wiley, 1994. doi:10.1002/9780470316887.

30 Michael O. Rabin. Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1–35, 1969. doi:10.2307/1995086.

31 Lloyd S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences*, 39(10):1095–1100, 1953. doi:10.1073/pnas.39.10.1095.

32 Frank Thuijsman. *Optimality and Equilibria in Stochastic Games.* Number no. 82 in CWI Tract – Centrum voor Wiskunde en Informatica. Centrum voor Wiskunde en Informatica, 1992. URL: https://books.google.co.uk/books?id=sfzuAAAAMAAJ.

33 Wiesław Zielonka. Infinite games on finitely coloured graphs with applications to automata on infinite trees. *Theoretical Computer Science*, 200(1-2):135–183, 1998. doi:10.1016/S0304-3975(98)00009-7.