Continuous Positional Payoffs

Alexander Kozachinskiy
University of Warwick, Department of Computer Science, UK

Abstract

What payoffs are positionally determined for deterministic two-player antagonistic games on finite directed graphs? In this paper we study this question for payoffs that are continuous. The main reason why continuous positionally determined payoffs are interesting is that they include the multi-discounted payoffs.

We show that for continuous payoffs positional determinacy is equivalent to a simple property called prefix-monotonicity. We provide three proofs of it, using three major techniques of establishing positional determinacy – inductive technique, fixed point technique and strategy improvement technique. A combination of these approaches provides us with better understanding of the structure of continuous positionally determined payoffs as well as with some algorithmic results.

2012 ACM Subject Classification Theory of computation → Logic

Keywords and phrases Games on graphs, positional strategies, continuous payoffs

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2021.27

Related Version *Full Version*: arXiv:2012.09047 [19]

Acknowledgements Supported by the EPSRC grant EP/P020992/1 (Solving Parity Games in Theory and Practice).

1 Introduction

We study games of the following kind. A game takes place on a finite directed graph. There is a token, initially located in one of the nodes. Before each turn there is exactly one node containing the token. In each turn one of the two antagonistic players called Max and Min chooses an edge starting in a node containing the token. As a result the token moves to the endpoint of this edge, and then the next turn starts. To determine who makes a move in a turn we are given in advance a partition of the nodes into two sets. If the token is in a node from the first set, then Max makes a move, otherwise Min.

Players make infinitely many moves, and this yields an infinite trajectory of the token. Technically, we assume that each node of the graph has at least one out-going edge so that there is always at least one available move. To introduce competitiveness, we should somehow compare the trajectories of the token with each other. For that we first fix some finite set A and label the edges of the game graph by elements of A. We also fix a *payoff* φ which is a function from the set of infinite sequences of elements of A to $\mathbb{R}$. Each possible infinite trajectory of the token is then mapped to a real number called *the reward* of this trajectory as follows: we form an infinite sequence of elements of A by taking the labels of edges along the trajectory, and apply φ to this sequence. The larger the reward is the more Max is happy; on the contrary, Min wants to minimize the reward.

For both of the players we are interested in indicating *an optimal strategy*, i.e., an optimal instruction of how to play in all possible developments of the games. To point out among all the strategies the optimal ones we first introduce a notion of a *value of a strategy*. The value of a Max's strategy σ is the infimum of the payoff over all infinite trajectories, consistent with the strategy. The reward of a play against σ cannot be smaller than its value, but can be arbitrarily close to it. Now, a strategy of Max is called optimal if its value is maximal over all Max's strategies. Similarly, the value of a Min's strategy is the supremum of the payoff

32nd International Conference on Concurrency Theory (CONCUR 2021).
Editors: Serge Haddad and Daniele Varacca; Article No. 27; pp. 27:1–27:17
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

over all infinite trajectories, consistent with this Min's strategy. Min's strategies minimizing the value are called optimal.

Observe that the value of any Min's strategy is at least as large as the value of any Max's strategy. A pair (σ, τ) of a Max's strategy σ and a Min's strategy τ is called an *equilibrium* if the value of σ *equals* the value of τ. Both strategies appearing in an equilibrium must be optimal – one proves the optimality of the other. In this paper we only study the so-called *determined* payoffs – payoffs for which all games on finite directed graphs with this payoff have an equilibrium.

For general determined payoffs an optimal strategy might be rather complicated (since the game is infinite, it might even have no finite description). For what determined payoffs both players always have a "simple" optimal strategy? A word "simple" can be understood in different ways [2], and this leads to different classes of determined payoffs. Among these classes we study one for which "simple" is understood in, perhaps, the strongest sense possible. Namely, we study a class of *positionally determined* payoffs.

For a positionally determined payoff all game graphs must have a pair of *positional* strategies which is an equilibrium no matter in which node the game starts. Now, a positional strategy is a strategy which totally ignores the previous trajectory of the token[1] and only looks at its current location. Formally, a positional strategy of Max maps each Max's node to an edge which starts in this node (i.e., to a single edge which Max will use whenever this node contains the token). Min's positional strategies are defined similarly.

A lot of works are devoted to concrete positionally determined payoffs that are of particular interest in other areas of computer science. Classical examples of such payoffs are parity payoffs, mean payoffs and (multi-)discounted payoffs [5, 21, 20, 23]. Their applications range from logic, verification and finite automata theory [6, 12] to decision-making [22, 24] and algorithm design [3].

Along with this specialized research, in [9, 10] Gimbert and Zielonka undertook a thorough study of positionally determined payoffs in general. In [9] they showed that all the so-called *fairly mixing* payoffs are positionally determined. They also demonstrated that virtually all classical positionally determined payoffs are fairly mixing. Next, in [10] they established a property of payoffs which is *equivalent* to positional determinacy. Despite being rather technical, this property has a remarkable feature: if a payoff does not satisfy it, then this payoff violates positional determinacy in some *one-player* game graph (where one of the players owns all the nodes). As Gimbert and Zielonka indicate, this means that to establish positional determinacy of a payoff it is enough to do so only for one-player game graphs.

One could try to gain more understanding about positionally determined payoffs that satisfy certain additional requirements. Of course, this is interesting only if there are practically important positionally determined payoffs that satisfy these requirements. One such requirement studied in the literature is called *prefix-independence* [4, 8]. A payoff is prefix-independent if it is invariant under throwing away any finite prefix from an infinite sequence of edge labels. For instance, the parity and the mean payoffs are prefix-independent.

In [9] Gimbert and Zielonka briefly mention another interesting additional requirement, namely, *continuity*. They observe that the multi-discounted payoffs are continuous (they utilize this in showing that the multi-discounted payoffs are fairly mixing). In this paper we study continuous positionally determined payoffs in more detail. Continuity of a payoff, loosely speaking, means that its range converges to just a single point as more and more

[1] In particular, a node in which the game has started.

initial characters of an infinite sequence of edge labels are getting fixed. This contrasts with prefix-independent payoffs (such as the parity and the mean payoffs), for which any initial finite segment is irrelevant. Thus, continuity serves as a natural property which separates the multi-discounted payoffs from the other classical positionally determined payoffs. This is our main motivation to study continuous positionally determined payoffs in general, besides the general importance of the notion of continuity.

We show that for continuous payoff positional determinacy is equivalent to a simple property which we call *prefix-monotonicity*. Loosely speaking, prefix-monotonicity means the result of a comparison of the payoff on two infinite sequences of labels does not change after appending or deleting the same finite prefix. In fact, we prove this result in three different ways, using three major techniques of establishing positional determinacy:

- *An inductive argument.* Here we use a sufficient condition of Gimbert and Zielonka [9], which is proved by induction on the number of edges of a game graph. This type of argument goes back to a paper of Ehrenfeucht and Mycielski [5], where they provide an inductive proof of the positional determinacy of the Mean Payoff Games.
- *A fixed point argument.* Then we give a proof which uses a fixed point approach due to Shapley [23]. Shapley's technique is a standard way of establishing positional determinacy of Discounted Games. In this argument one derives positional determinacy from the existence of a solution to a certain system of equations (sometimes called *Bellman's equations*). In turn, to establish the existence of a solution one uses Banach's fixed point theorem.
- *A strategy improvement argument.* For Discounted Games the existence of a solution to Bellman's equations can also be proved by *strategy improvement*. This technique goes back to Howard [16]; for its thorough treatment (as well as for its applications to other payoffs) we refer the reader to [7]. We generalize it to arbitrary continuous positionally determined payoffs.

The simplest way to obtain our main result is via the inductive argument (at the cost of appealing without a proof to the results of Gimbert and Zielonka). We provide two other proofs for the following reasons.

First, they have applications (and it is unclear how to get these applications within the framework of the inductive approach). The fixed point approach provides a precise understanding of what do continuous positionally determined payoffs look like in general. In the full version of this paper [19] we use this to answer a question of Gimbert [8] regarding positional determinacy in more general *stochastic* games. In turn, the strategy improvement approach has algorithmic consequences. More specifically, we show that a problem of finding a pair of optimal positional strategies is solvable in randomized subexponential time for any continuous positionally determined payoff.

Second, as far as we know, these two approaches were never used in such an abstract setting before. Thus, we believe that our paper makes a useful addition to these approaches from a technical viewpoint. For example, the main problem for the fixed point approach is to identify a metric with which one can carry out the same "contracting argument" as in the case of multi-discounted payoffs. To solve it, we obtain a result of independent interest about compositions of continuous functions. As for the strategy improvement approach, our main contribution is a generalization of such well-established tools as "modified costs" and "potential transformation lemma" [15, Lemma 3.6].

Organization of the paper. In Section 2 we formalize the concepts discussed in the introduction. Then in Sections 3–6 we expose our results in more detail. In Section 7 we indicate some possible future directions. Most of the proofs are omitted due to space constraints. In this version we provide only one of the three proofs of our main result completely (namely, one by the induction argument). Missing proofs can be found in the full version of this paper [19].

2 Preliminaries

We denote the function composition by $\circ$.

Sets and sequences. For two sets A and B by A^B we denote the set of all functions from B to A (sometime we will interpret A^B as the set of vectors consisting of elements of A and with coordinates indexed by elements of B). We write $C = A \sqcup B$ for three sets A, B, C if A and B are disjoint and $C = A \cup B$.

For a set A by A^* we denote the set of all finite sequences of elements of A and by A^ω we denote the set of all infinite sequences of elements of A. For $w \in A^*$ we let $|w|$ be the length of w. For $\alpha \in A^\omega$ we let $|\alpha| = \infty$.

For $u \in A^*$ and $v \in A^* \cup A^\omega$ we let uv denote the concatenation of u and v. We call $u \in A^*$ a prefix of $v \in A^* \cup A^\omega$ if for some $w \in A^* \cup A^\omega$ we have $v = uw$. For $w \in A^*$ by wA^ω we denote the set $\{w\alpha \mid \alpha \in A^\omega\}$. Alternatively, wA^ω is the set of all $\beta \in A^\omega$ such that w is a prefix of β.

For $u \in A^*$ and $k \in \mathbb{N}$ we use a notation

$$u^k = \underbrace{uu \ldots u}_{k \text{ times}}.$$

In turn, we let $u^\omega \in A^\omega$ be a unique element of A^ω such that u^k is a prefix of u^ω for every $k \in \mathbb{N}$. We call $\alpha \in A^\omega$ ultimately periodic if α is a concatenation of u and v^ω for some $u, v \in A^*$.

Graphs notation. By a finite directed graph G we mean a pair $G = (V, E)$ of two finite sets V and E equipped with two functions $\mathsf{source}, \mathsf{target} \colon E \to V$. Elements of V are called nodes of G and elements of E are called edges of G. For an edge $e \in E$ we understand $\mathsf{source}(e)$ (respectively, $\mathsf{target}(e)$) as the node in which e starts (respectively, ends). We allow parallel edges; i.e., there might be two distinct edges $e, e' \in E$ with $\mathsf{source}(e) = \mathsf{source}(e')$, $\mathsf{target}(e) = \mathsf{target}(e')$. We allow self-loops as well (i.e., edges with $\mathsf{source}(e) = \mathsf{target}(e)$).

The out-degree of a node $a \in V$ is $|\{e \in E \mid \mathsf{source}(e) = a\}|$. A node $a \in V$ is called a sink if its out-degree is 0. We call a graph G sinkless if there are no sinks in G.

A path in G is a non-empty (finite or infinite) sequence of edges of G with a property that $\mathsf{target}(e) = \mathsf{source}(e')$ for any two consecutive edges e and e' from the sequence. For a path p we define $\mathsf{source}(p) = \mathsf{source}(e)$, where e is the first edge of p. For a finite path p we define $\mathsf{target}(p) = \mathsf{target}(e')$, where e' is the last edge of p.

For technical convenience we also consider 0-length paths. Each 0-length path is associated with some node of G (so that there are $|V|$ different 0-length paths). For a 0-length path p, associated with $a \in V$, we define $\mathsf{source}(p) = \mathsf{target}(p) = a$.

When we write pq for two paths p and q we mean the concatenation of p and q (viewed as sequences of edges). Of course, this is well-defined only if p is finite. Note that pq is not necessarily a path. Namely, pq is a path if and only if $\mathsf{target}(p) = \mathsf{source}(q)$.

2.1 Deterministic infinite duration games on finite directed graphs

Mechanics of the game. By a *game graph* we mean a sinkless finite directed graph $G = \langle V, E, \mathsf{source}, \mathsf{target}\rangle$, equipped with two sets V_{Max} and V_{Min} such that $V = V_{\text{Max}} \sqcup V_{\text{Min}}$.

A game graph $G = \langle V = V_{\text{Max}} \sqcup V_{\text{Min}}, E, \mathsf{source}, \mathsf{target}\rangle$ induces a so-called *infinite duration game* (IDG for short) on G. The game is always between two players called Max and Min. Positions of the game are finite paths in G (informally, these are possible finite trajectories of the token). We call a finite path p a Max's (a Min's) position if $\mathsf{target}(p) \in V_{\text{Max}}$ (if $\mathsf{target}(p) \in V_{\text{Min}}$). Max makes moves in Max's positions and Min makes moves in Min's positions. We do not indicate any position as the starting one – it can be any node of G.

The set of moves available at a position p is the set $\{e \in E \mid \mathsf{source}(e) = \mathsf{target}(p)\}$. A move e from a position p leads to a position pe.

A Max's strategy σ in a game graph G is a mapping assigning to every Max's position p a move available at p. Similarly, a Min's strategy τ in a game graph G is a mapping assigning to every Min's position p a move available at p.

Let $\mathcal{P} = e_1e_2e_3\dots$ be an infinite path in G. We say that $\mathcal{P}$ is *consistent* with a Max's strategy σ if the following conditions hold:

- if $s = \mathsf{source}(\mathcal{P}) \in V_{\text{Max}}$, then $\sigma(s) = e_1$;
- for every $i \geq 1$ it holds that $\mathsf{target}(e_1e_2\dots e_i) \in V_{\text{Max}} \implies e_{i+1} = \sigma(e_1e_2\dots e_i)$.

For $a \in V$ and for a Max's strategy σ we let $\mathsf{Cons}(a,\sigma)$ be a set of all infinite paths in G that start in a and are consistent with σ. We use similar terminology and notation for strategies of Min.

Given a Max's strategy σ, a Min's strategy τ and $a \in V$, we let *the play of σ and τ from a* be a unique element of the intersection $\mathsf{Cons}(a,\sigma) \cap \mathsf{Cons}(a,\tau)$. The play of σ and τ from a is denoted by $\mathcal{P}_a^{\sigma,\tau}$.

Positional strategies. A Max's strategy σ in a game graph $G = \langle V = V_{\text{Max}} \sqcup V_{\text{Min}}, E, \mathsf{source}, \mathsf{target}\rangle$ is called *positional* if $\sigma(p) = \sigma(q)$ for all finite paths p and q in G with $\mathsf{target}(p) = \mathsf{target}(q) \in V_{\text{Max}}$. Clearly, a Max's positional strategy σ can be represented as a mapping $\sigma\colon V_{\text{Max}} \to E$ satisfying $\mathsf{source}(\sigma(u)) = u$ for all $u \in V_{\text{Max}}$. We define Min's positional strategies analogously.

We call an edge $e \in E$ *consistent* with a Max's positional strategy σ if either $\mathsf{source}(e) \in V_{\text{Min}}$ or $\mathsf{source}(e) \in V_{\text{Max}}, e = \sigma(\mathsf{source}(e))$. We denote the set of edges that are consistent with σ by E^σ. If τ is a Min's positional strategy, then we say that an edge $e \in E$ is consistent with τ if either $\mathsf{source}(e) \in V_{\text{Max}}$ or $\mathsf{source}(e) \in V_{\text{Min}}, e = \tau(\mathsf{source}(e))$. The set of edges that are consistent with a Min's positional strategy τ is denoted by E_τ.

Labels and payoffs. Let A be a finite set. A game graph $G = \langle, V = V_{\text{Max}} \sqcup V_{\text{Min}}, E, \mathsf{source}, \mathsf{target}\rangle$ equipped with a function $\mathsf{lab}\colon E \to A$ is called an *A-labeled game graph*. If $p = e_1e_2e_3\dots$ is a (finite or infinite) path in an A-labeled game graph $G = \langle V = V_{\text{Max}} \sqcup V_{\text{Min}}, E, \mathsf{source}, \mathsf{target}, \mathsf{lab}\rangle$, we define $\mathsf{lab}(p) = \mathsf{lab}(e_1)\mathsf{lab}(e_2)\mathsf{lab}(e_3)\dots \in A^* \cup A^\omega$. A *payoff* is a bounded function from A^ω to $\mathbb{R}$. Some papers allow A to be infinite and consider only infinite sequences that contain finitely many elements of A (as any game graph contains only finitely many labels). So basically we just have to deal with finite subsets of A, and this can be done with our approach.

Values, optimal strategies and equilibria. Let A be a finite set, $\varphi\colon A^\omega \to \mathbb{R}$ be a payoff and $G = \langle V = V_{\text{Max}} \sqcup V_{\text{Min}}, E, \mathsf{source}, \mathsf{target}, \mathsf{lab}\rangle$ be an A-labeled game graph. Take a Max's strategy σ in G. *The value* of σ in a node $a \in V$ is the following quantity:

$$\mathsf{Val}[\sigma](a) = \inf \varphi \circ \mathsf{lab}\big(\mathsf{Cons}(a,\sigma)\big).$$

Similarly, if τ is a Min's strategy in G, then the value of τ in a node $a \in V$ is the following quantity:

$$\mathsf{Val}[\tau](a) = \sup \varphi \circ \mathsf{lab}\big(\mathsf{Cons}(a, \tau)\big).$$

A Max's strategy σ is called *optimal* if $\mathsf{Val}[\sigma](a) \geq \mathsf{Val}[\sigma'](a)$ for any $a \in V$ and for any Max's strategy σ'. Similarly, A Min's strategy τ is called *optimal* if $\mathsf{Val}[\tau](a) \leq \mathsf{Val}[\tau'](a)$ for any $a \in V$ and for any Min's strategy τ'.

Observe that for any Max's strategy σ, for any Min's strategy τ and for any $a \in V$ we have:

$$\mathsf{Val}[\sigma](a) \leq \varphi \circ \mathsf{lab}\big(\mathcal{P}_a^{\sigma,\tau}\big) \leq \mathsf{Val}[\tau](a).$$

In particular, this inequality gives us the following. If a pair (σ, τ) of a Max's strategy σ and a Min's strategy τ is such that $\mathsf{Val}[\sigma](a) = \mathsf{Val}[\tau](a)$ for every $a \in V$, then both σ and τ are optimal for their players. We call any pair (σ, τ) with $\mathsf{Val}[\sigma](a) = \mathsf{Val}[\tau](a)$ for every $a \in V$ an *equilibrium*[2]. In fact, if at least one equilibrium exists, then the following holds: the Cartesian product of the set of the optimal strategies of Max and the set of the optimal strategies of Min is exactly the set of equilibria. We say that φ is *determined* if in every A-labeled game graph there exists an equilibrium (with respect to φ).

Positionally determined payoffs. Let A be a finite set and $\varphi\colon A^\omega \to \mathbb{R}$ be a payoff. We call φ *positionally determined* if all A-labeled game graphs have (with respect to φ) an equilibrium consisting of two positional strategies.

▶ **Proposition 1.** *If A is a finite set, $\varphi\colon A^\omega \to \mathbb{R}$ is a positionally determined payoff and $g\colon \varphi(A^\omega) \to \mathbb{R}$ is a non-decreasing[3] function, then $g \circ \varphi$ is a positionally determined payoff.*

2.2 Continuous payoffs

For a finite set A, we consider the set A^ω as a topological space. Namely, we take the discrete topology on A and the corresponding product topology on A^ω. In this product topology open sets are sets of the form

$$\mathcal{S} = \bigcup_{u \in S} uA^\omega,$$

where $S \subseteq A^*$. When we say that a payoff $\varphi\colon A^\omega \to \mathbb{R}$ is *continuous* we always mean continuity with respect to this product topology (and with respect to the standard topology on $\mathbb{R}$). The following proposition gives a convenient way to establish continuity of payoffs.

▶ **Proposition 2.** *Let A be a finite set. A payoff $\varphi\colon A^\omega \to \mathbb{R}$ is continuous if and only if for any $\alpha \in A^\omega$ and for any infinite sequence $\{\beta_n\}_{n=1}^\infty$ of elements of A^ω the following holds. If for all $n \geq 1$ the sequences α and β_n coincide in the first n elements, then $\lim\limits_{n\to\infty} \varphi(\beta_n)$ exists and equals $\varphi(\alpha)$.*

For a finite A by Tychonoff's theorem the space A^ω is compact (because any finite set A with the discrete topology is compact). This has the following consequence which is important for this paper: if $\varphi\colon A^\omega \to \mathbb{R}$ is a continuous payoff, then $\varphi(A^\omega)$ is a compact subset of $\mathbb{R}$.

---

[2] This definition is equivalent to a more standard one: (σ, τ) is an equilibrium if and only if σ is a "best response" to τ in every node, and vice versa.

[3] Throughout the paper we call a function $f\colon S \to \mathbb{R}, S \subseteq \mathbb{R}$ non-decreasing if for all $x, y \in S$ with $x \leq y$ we have $f(x) \leq f(y)$.

3 Statement of the Main Result and Preliminary Discussion

Our main result establishes a simple property which is equivalent to positional determinacy for continuous payoffs.

▶ **Definition 3.** *Let A be a finite set. A payoff $\varphi\colon A^\omega \to \mathbb{R}$ is called* ***prefix-monotone*** *if there are no $u, v \in A^*$, $\beta, \gamma \in A^\omega$ such that $\varphi(u\beta) > \varphi(u\gamma)$ and $\varphi(v\beta) < \varphi(v\gamma)$.*

(One can note that prefix-independence trivially implies prefix-monotonicity. On the other hand, no prefix-independent payoff which takes at least 2 values is continuous.)

▶ **Theorem 4.** *Let A be a finite set and $\varphi\colon A^\omega \to \mathbb{R}$ be a continuous payoff. Then φ is positionally determined if and only if φ is prefix-monotone.*

The fact that any continuous positionally determined payoff must be prefix-monotone[4] is proved in Appendix A. Three different proofs of the "if" part of Theorem 4 are discussed in, respectively, Sections 4, 5 and 6. Before going into the proofs, let us discuss the notions of continuity and prefix-monotonicity by means of the multi-discounted payoffs.

▶ **Definition 5.** *A payoff $\varphi\colon A^\omega \to \mathbb{R}$ for a finite set A is* ***multi-discounted*** *if there are functions $\lambda\colon A \to [0, 1)$ and $w\colon A \to \mathbb{R}$ such that*

$$\varphi(a_1a_2a_3\ldots) = \sum_{n=1}^{\infty} \lambda(a_1)\cdot\ldots\cdot\lambda(a_{n-1})\cdot w(a_n) \qquad (1)$$

for all $a_1a_2a_3\ldots \in A^\omega$.

A few technical remarks: since the set A is finite, the coefficients $\lambda(a)$ are bounded away from 1 uniformly over $a \in A$. This ensures that the series (1) converges. In fact, this means that a tail of this series converges to 0 uniformly over $a_1a_2a_3\ldots \in A^\omega$. Thus, the multi-discounted payoffs are continuous. As the multi-discounted payoffs are positionally determined, by Theorem 4 they also must be prefix-monotone. Of course, prefix-monotonicity of the multi-discounted payoffs can be established without Theorem 4. Indeed, from (1) it is easy to derive that $\varphi(a\beta) - \varphi(a\gamma) = \lambda(a)\cdot(\varphi(\beta) - \varphi(\gamma))$ for all $a \in A, \beta, \gamma \in A^\omega$. Due to the condition $\lambda(a) \geq 0$, we have that $\varphi(a\beta) > \varphi(a\gamma)$ implies that $\varphi(\beta) > \varphi(\gamma)$. Moreover, the same holds if we append more than one character to β and γ. Hence it is impossible to simultaneously have $\varphi(u\beta) > \varphi(u\gamma)$ and $\varphi(v\beta) < \varphi(v\gamma)$ for $u, v \in A^*$, as required in the definition of prefix-monotonicity.

4 Inductive Argument

Here we show that any continuous prefix-monotone payoff is positionally determined using a sufficient condition of Gimbert and Zielonka [9, Theorem 1], which, in turn, is proved by an inductive argument. As Gimbert and Zielonka indicate [9, Lemma 2], their sufficient condition takes the following form for continuous payoffs[5].

[4] Here it is crucial that in our definition of positional determinacy we require that some positional strategy is optimal for all the nodes. Allowing each starting node to have its own optimal positional strategy gives us a weaker, "non-uniform" version of positional determinacy. It is not clear whether non-uniform positional determinacy implies prefix-monotonicity. At the same time, we are not even aware of a payoff which is positional only "non-uniformly".

[5] Lemma 2 can only be found in the HAL version of their paper.

▶ **Proposition 6.** *Let A be a finite set. Any continuous payoff $\varphi\colon A^\omega \to \mathbb{R}$, satisfying the following two conditions:*

- ***(a)** for all $u \in A^*$ and $\alpha, \beta \in A^\omega$ we have that $\varphi(\alpha) \le \varphi(\beta) \implies \varphi(u\alpha) \le \varphi(u\beta)$;*
- ***(b)** for all non-empty $u \in A^*$ and for all $\alpha \in A^\omega$ we have that*

$$\min\{\varphi(u^\omega), \varphi(\alpha)\} \le \varphi(u\alpha) \le \max\{\varphi(u^\omega), \varphi(\alpha)\};$$

is positionally determined.

We observe that one can get rid of the condition ***(b)*** in this Proposition.

▶ **Proposition 7.** *For continuous payoffs the condition **(a)** of Proposition 6 implies the condition **(b)** of Proposition 6.*

Proof. See Appendix B. ◀

So to establish positional determinacy of a continuous payoff it is enough to demonstrate that this payoff satisfies the condition ***(a)*** of Proposition 6. Let us now reformulate this condition using the following definition.

▶ **Definition 8.** *Let A be a finite set. A payoff $\varphi\colon A^\omega \to \mathbb{R}$ is called **shift-deterministic** if for all $a \in A, \beta, \gamma \in A^\omega$ we have $\varphi(\beta) = \varphi(\gamma) \implies \varphi(a\beta) = \varphi(a\gamma)$.*

▶ **Observation 9.** *Let A be a finite set. A payoff $\varphi\colon A^\omega \to \mathbb{R}$ satisfies the condition **(a)** of Proposition 6 if and only if φ is prefix-monotone and shift-deterministic.*

The above discussion gives the following sufficient condition for positional determinacy.

▶ **Proposition 10.** *Let A be a finite set. Any continuous prefix-monotone shift-deterministic payoff $\varphi\colon A^\omega \to \mathbb{R}$ is positionally determined.*

Still, some argument is needed for continuous prefix-monotone payoffs that are not shift-deterministic. To tie up loose ends we prove the following:

▶ **Proposition 11.** *Let A be a finite set and let $\varphi\colon A^\omega \to \mathbb{R}$ be a continuous prefix-monotone payoff. Then $\varphi = g \circ \psi$ for some continuous prefix-monotone shift-deterministic payoff $\psi\colon A^\omega \to \mathbb{R}$ and for some continuous[6] non-decreasing function $g\colon \psi(A^\omega) \to \mathbb{R}$.*

Proof. See Appendix C. ◀

Due to Proposition 1 this finishes our first proof of Theorem 4. In fact, we do not need continuity of g here, but it will be useful later.

5 Fixed point argument

Here we present a way of establishing positional determinacy of continuous prefix-monotone shift-deterministic payoffs (Proposition 10) via a fixed point argument. Together with Proposition 11 this constitutes our second proof of Theorem 4.

Obviously, for any shift-deterministic payoff $\varphi\colon A^\omega \to \mathbb{R}$ and for any $a \in A$ there is a unique function $\mathsf{shift}[a,\varphi]\colon \varphi(A^\omega) \to \varphi(A^\omega)$ such that $\mathsf{shift}[a,\varphi]\big(\varphi(\beta)\big) = \varphi(a\beta)$ for all $\beta \in A^\omega$.

[6] Throughout the paper we call a function $f\colon S \to \mathbb{R}$, $S \subseteq \mathbb{R}^n$ continuous if f is continuous with respect to a restriction of the standard topology of $\mathbb{R}^n$ to S.

▶ **Observation 12.** *A shift-deterministic payoff* $\varphi\colon A^\omega \to \mathbb{R}$ *is prefix-monotone if and only if* $\mathsf{shift}[a,\varphi]$ *is non-decreasing for every* $a \in A$.

We use this notation to introduce the so-called *Bellman's equations*, playing a key role in our fixed point argument.

▶ **Definition 13.** *Let* A *be a finite set,* $\varphi\colon A^\omega \to \mathbb{R}$ *be a shift-deterministic payoff and* $G = \langle V = V_{\mathrm{Max}} \sqcup V_{\mathrm{Min}}, E, \mathsf{source}, \mathsf{target}, \mathsf{lab}\rangle$ *be an* A*-labeled game graph.*

The following equations in $\mathbf{x} \in \varphi(A^\omega)^V$ *are called* ***Bellman's equations*** *for* φ *in* G*:*

$$\mathbf{x}_u = \max_{e\in E, \mathsf{source}(e)=u} \mathsf{shift}[\mathsf{lab}(e),\varphi]\big(\mathbf{x}_{\mathsf{target}(e)}\big), \qquad \textit{for } u \in V_{\mathrm{Max}}, \tag{2}$$

$$\mathbf{x}_u = \min_{e\in E, \mathsf{source}(e)=u} \mathsf{shift}[\mathsf{lab}(e),\varphi]\big(\mathbf{x}_{\mathsf{target}(e)}\big), \qquad \textit{for } u \in V_{\mathrm{Min}}. \tag{3}$$

The most important step of our argument is to show the existence of a solution to Bellman's equations.

▶ **Proposition 14.** *For any finite set* A*, for any continuous prefix-monotone shift-deterministic payoff* $\varphi\colon A^\omega \to \mathbb{R}$ *and for any* A*-labeled game graph* G *there exists a solution to Bellman's equations for* φ *in* G*.*

(One can also show the uniqueness of a solution, but we do not need this for the argument).

This proposition requires some additional work, and we first discuss how to derive positional determinacy of continuous prefix-monotone shift-deterministic payoffs from it. Assume that we are give a solution $\mathbf{x}$ to (2–3). How can one extract an equilibrium of positional strategies from it? For that we take any pair of positional strategies that use only $\mathbf{x}$-*tight* edges. Now, an edge e is $\mathbf{x}$-tight if $\mathbf{x}_{\mathsf{source}(e)} = \mathsf{shift}[a,\varphi](\mathbf{x}_{\mathsf{target}(e)})$. Note that each node must contain an out-going $\mathbf{x}$-tight edge (this will be any edge on which the maximum/minimum in (2–3) is attained for this node). So clearly each player has at least one positional strategy which only uses $\mathbf{x}$-tight edges. It remains to show that for continuous prefix-monotone shift-deterministic φ any two such strategies of the players form an equilibrium.

▶ **Lemma 15.** *If* A *is a finite set,* $\varphi\colon A^\omega \to \mathbb{R}$ *is a continuous prefix-monotone shift-deterministic payoff, and* $\mathbf{x} \in \varphi(A^\omega)^V$ *is a solution to (2–3) for an* A*-labeled game graph* $G = \langle V = V_{\mathrm{Max}} \sqcup V_{\mathrm{Min}}, E, \mathsf{source}, \mathsf{target}, \mathsf{lab}\rangle$*, then the following holds. Let* σ^* *be a positional strategy of Max and* τ^* *be a positional strategy of Min such that* $\sigma^*(V_{\mathrm{Max}})$ *and* $\tau^*(V_{\mathrm{Min}})$ *consist only of* $\mathbf{x}$*-tight edges. Then* (σ^*, τ^*) *is an equilibrium in* G*.*

We now proceed to details of our proof of Proposition 14. Consider a function $T\colon \varphi(A^\omega)^V \to \varphi(A^\omega)^V$, mapping $\mathbf{x} \in \varphi(A^\omega)^V$ to the vector of the right-hand sides of (2–3). We should argue that T has a fixed point. For that we will construct a continuous metric $D\colon \varphi(A^\omega)^V \times \varphi(A^\omega)^V \to [0,+\infty)$ with respect to which T is *contracting.* More precisely, $D(T\mathbf{x}, T\mathbf{y})$ will always be smaller than $D(\mathbf{x},\mathbf{y})$ as long as $\mathbf{x}$ and $\mathbf{y}$ are distinct. Due to the compactness of the domain of T this will prove that T has a fixed point.

Now, to construct such D we show that for continuous shift-deterministic φ there must be a continuous metric $d\colon \varphi(A^\omega) \times \varphi(A^\omega) \to [0,+\infty)$ such that for all $a \in A$ the function $\mathsf{shift}[a,\varphi]$ is d-contracting. Once we have such d, we let $D(\mathbf{x},\mathbf{y})$ be the maximum of $d(\mathbf{x}_a, \mathbf{y}_a)$ over $a \in V$. Checking that T is contracting with respect to such D will be rather straightforward (technically, we will need an additional property of d which can be derived from the prefix-monotonicity of φ).

The main technical challenge is to prove the existence of d. In the full version of this paper we do so via the following general fact about compositions of continuous functions.

▶ **Theorem 16.** *Let $K \subseteq \mathbb{R}$ be a compact set, $m \geq 1$ be a natural number and $f_1, \ldots, f_m \colon K \to K$ be m continuous functions. Then the following two conditions are equivalent:*

- ***(a)*** *for any $a_1 a_2 a_3 \ldots \in \{1, 2, \ldots, m\}^\omega$ we have $\lim_{n\to\infty} \mathsf{diam}\big(f_{a_1} \circ f_{a_2} \circ \ldots \circ f_{a_n}(K)\big) = 0$ (by $\mathsf{diam}(S)$ for $S \subseteq \mathbb{R}$ we mean $\sup_{x,y\in S} |x - y|$);*
- ***(b)*** *there exists a continuous metric $d \colon K \times K \to [0, +\infty)$ such that $f_1, f_2, \ldots, f_m$ are all d-contracting (a function $h \colon K \to K$ is called d-contracting if for all $x, y \in K$ with $x \neq y$ we have $d(h(x), h(y)) < d(x, y)$).*

*If $f_1, \ldots, f_m$ are non-decreasing, then one can strengthen item **(b)** by demanding that d satisfies the following property: for all $x, y, s, t \in K$ with $x \leq s \leq t \leq y$ we have $d(s, t) \leq d(x, y)$.*

Namely, we apply this theorem to the functions $\mathsf{shift}[a, \varphi]$ for $a \in A$ (for that we first show that the continuity of φ implies that these functions satisfy item ***(a)*** of Theorem 16).

5.1 Applications of the fixed point technique

Theorem 16 additionally provides an exhaustive method of generating continuous positionally determined payoffs.

▶ **Theorem 17.** *Let m be a natural number. The set of continuous positionally determined payoffs from[7] $\{1, 2, \ldots, m\}^\omega$ to $\mathbb{R}$ coincides with the set of φ that can be obtained in the following 5 steps.*

- ***Step 1.*** *Take a compact set $K \subseteq \mathbb{R}$.*
- ***Step 2.*** *Take a continuous metric $d \colon K \times K \to [0, +\infty)$.*
- ***Step 3.*** *Take m non-decreasing d-contracting functions $f_1, f_2, \ldots, f_m \colon K \to K$ (they will automatically be continuous due to continuity of d).*
- ***Step 4.*** *Define $\psi \colon \{1, \ldots, m\}^\omega \to K$ so that*

$$\{\psi(a_1 a_2 a_3 \ldots)\} = \bigcap_{n=1}^{\infty} f_{a_1} \circ f_{a_2} \circ \ldots \circ f_{a_n}(K)$$

  for every [8] $a_1 a_2 a_3 \ldots \in \{1, 2, \ldots, m\}^\omega$.
- ***Step 5.*** *Choose a continuous non-decreasing function $g \colon \psi(\{1, 2, \ldots, m\}^\omega) \to \mathbb{R}$ and set $\varphi = g \circ \psi$.*

▶ Remark 18. Recall that we did not use continuity of g from Proposition 11 in the inductive argument. It becomes important for Theorem 17 – otherwise we could not argue that all continuous positionally payoffs can be obtained in these 5 steps.

We get the multi-discounted payoffs when the functions $f_1, f_2, \ldots, f_m$ are affine, each with the slope from $[0, 1)$. In this case they will be contracting with respect to a standard metric $d(x, y) = |x - y|$. We get the whole set of continuous positionally determined payoffs by relaxing the multi-discounted payoffs in the following three regards: ***(a)*** functions $f_1, f_2, \ldots, f_m$ do not have to be affine; ***(b)*** d can be an arbitrary continuous metric; ***(c)*** any continuous non-decreasing function g can be applied to a payoff.

We use Theorem 17 to construct a continuous positionally determined payoff which does not “reduce” to the multi-discounted ones, in a sense of the following definition.

---

[7] Of course, in this theorem a set of labels can be any finite set, we let it be $\{1, 2, \ldots, m\}$ for some $m \in \mathbb{N}$ just to simplify the notation.

[8] Note that this intersection always consists of a single point due to Cantor's intersection theorem and item ***(a)*** of Theorem 16. This will also be $\lim_{n\to\infty} f_{a_1} \circ f_{a_2} \circ \ldots \circ f_{a_n}(x)$ for any $x \in K$.

▶ **Definition 19.** *Let A be a finite set, $\varphi, \psi \colon A^\omega \to \mathbb{R}$ be two payoffs, and G be an A-labeled game graph. We say that φ* ***positionally reduces*** *to ψ* ***inside*** *G if any pair of positional strategies in G which is an equilibrium for ψ is also an equilibrium for φ.*

This definition has an algorithmic motivation. Namely, note that finding a positional equilibrium for ψ in G is at least as hard as for φ, provided that φ reduces to ψ inside G. There are classical reductions from Parity to Mean Payoff games [17] and from Mean Payoff to Discounted games [25] that work in exactly this way. See also [11] for a reduction from *Priority* Mean Payoff games to Multi-Discounted games. As far as we know, our next proposition provides the first example of a positionally determined payoff which does not reduce to the multi-discounted ones in this sense.

▶ **Proposition 20.** *There exist a finite set A, a continuous positionally determined payoff $\varphi \colon A^\omega \to \mathbb{R}$ and an A-labeled game graph G such that there exists no multi-discounted payoff to which φ reduces inside G.*

Proposition 20 means, in particular, that there exists a continuous positionally determined payoff which differs from all the multi-discounted ones (as was stated in Section 3). This fact alone can be used to disprove a conjecture of Gimbert [8]. Namely, Gimbert conjectured the following: "Any payoff function which is positional for the class of non-stochastic one-player games is positional for the class of Markov decision processes". To show that this is not the case, in the full version of this paper [19] we establish that all continuous payoffs that are positionally determined in Markov decision processes are multi-discounted.

6 Strategy improvement argument

Here we establish the existence of a solution to Bellman's equations (Proposition 14) via the *strategy improvement*. This will yield our third proof of Theorem 4. We start with an observation that a vector of values of a positional strategy always gives a solution[9] to a restriction of Bellman's equations to edges that are consistent with this strategy.

▶ **Lemma 21.** *Let A be a finite set, $\varphi \colon A^\omega \to \mathbb{R}$ be a continuous prefix-monotone shift-deterministic payoff and $G = \langle V = V_{\mathrm{Max}} \sqcup V_{\mathrm{Min}}, E, \mathsf{source}, \mathsf{target}, \mathsf{lab}\rangle$ be an A-labeled game graph. Then for every positional strategy σ of Max in G we have:*

$$\mathsf{Val}[\sigma](u) = \mathsf{shift}[\mathsf{lab}(\sigma(u)), \varphi]\Big(\mathsf{Val}[\sigma]\big(\mathsf{target}(\sigma(u))\big)\Big) \text{ for } u \in V_{\mathrm{Max}},$$

$$\mathsf{Val}[\sigma](u) = \min_{e \in E, \mathsf{source}(e) = u} \mathsf{shift}[\mathsf{lab}(e), \varphi]\Big(\mathsf{Val}[\sigma](\mathsf{target}(e))\Big) \text{ for } u \in V_{\mathrm{Min}}.$$

Next, take a positional strategy σ of Max. If the vector $\{\mathsf{Val}[\sigma](u)\}_{u \in V}$ happens to be a solution to the Bellman's equations, then we are done. Otherwise by Lemma 21 there must exist an edge $e \in E$ with $\mathsf{source}(e) \in V_{\mathrm{Max}}$ such that $\mathsf{Val}[\sigma](\mathsf{source}(e)) < \mathsf{shift}[\mathsf{lab}(e), \varphi]\big(\mathsf{Val}[\sigma](\mathsf{target}(e))\big)$. We call edges satisfying this property *σ-violating*. We show that *switching* σ to any σ-violating edge gives us a positional strategy which *improves* σ.

[9] Bellman's equations involve the functions $\mathsf{shift}[a, \varphi]$ for $a \in A$, and these functions are defined on $\varphi(A^\omega)$. So formally we should argue that the values of any strategy belong to $\varphi(A^\omega)$. Indeed, for continuous φ the set $\varphi(A^\omega)$ is compact and hence is closed, and all values are the infimums/supremums of some subsets of $\varphi(A^\omega)$.

▶ **Lemma 22.** *Let A be a finite set, $\varphi\colon A^\omega \to \mathbb{R}$ be a continuous prefix-monotone shift-deterministic payoff and $G = \langle V = V_{\mathrm{Max}} \sqcup V_{\mathrm{Min}}, E, \mathsf{source}, \mathsf{target}, \mathsf{lab}\rangle$ be an A-labeled game graph. Next, let σ be a positional strategy of Max in G. Assume that the vector $\mathsf{Val}[\sigma] = \{\mathsf{Val}[\sigma](u)\}_{u\in V}$ does not satisfy (2–3) and let $e' \in E$ be any σ-violating edge. Define a positional strategy σ' of Max as follows:*

$$\sigma'(u) = \begin{cases} e' & u = \mathsf{source}(e'), \\ \sigma(u) & \textit{otherwise.} \end{cases}$$

Then $\sum\limits_{u\in V} \mathsf{Val}[\sigma'](u) > \sum\limits_{u\in V} \mathsf{Val}[\sigma](u)$.

By this lemma, a Max's positional strategy σ^* maximizing the quantity $\sum_{u\in V} \mathsf{Val}[\sigma](u)$ (over positional strategies σ of Max) gives a solution to (2–3). Such σ^* exists just because there are only finitely many positional strategies of Max. This finishes our strategy improvement proof of Proposition 14. Let us note that the same argument can be carried out with positional strategies of Min (via analogues of Lemma 21 and Lemma 22 for Min).

6.1 Applications of the strategy improvement technique

In this subsection we discuss implications of our strategy improvement argument to the *strategy synthesis problem*. Strategy synthesis for a positionally determined payoff φ is an algorithmic problem of finding an equilibrium (with respect to φ) of two positional strategies for a given game graph. It is classical that strategy synthesis for classical positionally determined payoffs admits a randomized algorithm which is subexponential in the number of nodes [14, 1]. We obtain the same subexponential bound for all continuous positionally determined payoffs. From a technical viewpoint, we just observe that a technique which was used for classical positionally determined payoffs is applicable in a more general setting. Specifically, we use a framework of recursively local-global functions due to Björklund and Vorobyov [1].

Let us start with an observation that for continuous positionally determined shift-deterministic payoffs a non-optimal positional strategy can always be improved by changing it just in a single node.

▶ **Proposition 23.** *Let A be a finite set and $\varphi\colon A^\omega \to \mathbb{R}$ be a continuous positionally determined shift-deterministic payoff. Then for any A-labeled game graph $G = \langle V = V_{\mathrm{Max}} \sqcup V_{\mathrm{Min}}, E, \mathsf{source}, \mathsf{target}, \mathsf{lab}\rangle$ the following two conditions hold:*

- *if σ is a non-optimal positional strategy of Max in G, then in G there exists a Max's positional strategy σ' such that $|\{u \in V_{\mathrm{Max}} \mid \sigma(u) \neq \sigma'(u)\}| = 1$ and $\sum_{u\in V} \mathsf{Val}[\sigma'](u) > \sum_{u\in V} \mathsf{Val}[\sigma](u)$;*
- *if τ is a non-optimal positional strategy of Min in G, then in G there exists a Min's positional strategy τ' such that $|\{u \in V_{\mathrm{Min}} \mid \tau(u) \neq \tau'(u)\}| = 1$ and $\sum_{u\in V} \mathsf{Val}[\tau'](u) < \sum_{u\in V} \mathsf{Val}[\tau](u)$.*

It is instructive to visualize this proposition by imagining the set of positional strategies of one of the players (say, Max) as a *hypercube*. Namely, in this hypercube there will be as many dimensions as there are nodes of Max. A coordinate corresponding to a node $u \in V_{\mathrm{Max}}$ will take values in the set of edges that start at u. Obviously, vertices of such hypercube are in a one-to-one correspondence with positional strategies of Max. Let us call two vertices *neighbors* of each other if they differ in exactly one coordinate. Now, Proposition 23 means in this language the following: any vertex σ, maximizing $\sum_{u\in V} \mathsf{Val}[\sigma](u)$ over its neighbors, also maximizes this quantity over the *whole* hypercube.

So an optimization problem of maximizing $\sum_{u\in V} \mathsf{Val}[\sigma](u)$ (equivalently, finding an optimal positional strategy of Max) has the following remarkable feature: all its *local* maxima are also *global*. For positional strategies of Min the same holds for the minima. Optimization problems with this feature are in a focus of numerous works, starting from a classical area of convex optimization.

Observe that in our case this local-global property is *recursive*; i.e., it holds for any restriction to a *subcube* of our hypercube. Indeed, subcubes correspond to subgraphs of our initial game graph, and for any subgraph we still have Proposition 23. Björklund and Vorobyov [1] noticed that a similar phenomenon occurs for all classical positionally determined payoffs. In turn, they showed that any optimization problem on a hypercube with this recursive local-global property admits a randomized algorithm which is subexponential in the dimension of a hypercube. In our case this yields a randomized algorithm for the strategy synthesis problem which is subexponential in the number of nodes of a game graph.

Still, this only applies to continuous payoffs that are shift-deterministic (as we have Proposition 23 only for shift-deterministic payoffs). One more issue is that we did not specify how our payoffs are represented. We overcome these difficulties in the following result.

▶ **Theorem 24.** *Let A be a finite set and $\varphi\colon A^\omega \to \mathbb{R}$ be a continuous positionally determined payoff. Consider an oracle which for given $u, v, a, b \in A^*$ tells, whether there exists $w \in A^*$ such that $\varphi(wu(v)^\omega) > \varphi(wa(b)^\omega)$. There exists a randomized algorithm which with this oracle solves the strategy synthesis problem for φ in expected $e^{O\left(\log m + \sqrt{n \log m}\right)}$ time for game graphs with n nodes and m edges. In particular, every call to the oracle in the algorithm is for $u, v, a, b \in A^*$ that are of length $O(n)$, and the expected number of the calls is $e^{O\left(\log m + \sqrt{n \log m}\right)}$.*

So to deal with the issue of representation we assume a suitable oracle access to φ. Still, the oracle from Theorem 24 might look unmotivated. Here it is instructive to recall that all continuous positionally determined φ must be prefix-monotone. For prefix-monotone φ a formula $\exists w \in A^*\ \varphi(w\alpha) > \varphi(w\beta)$ defines a total preorder on A^ω, and our oracle just compares ultimately periodic sequences according to this preorder. In fact, it is easy to see that the formula $\exists w \in A^*\ \varphi(w\alpha) > \varphi(w\beta)$ defines a total preorder on A^ω if and *only if* φ is prefix-monotone. This indicates a fundamental role of this preorder for prefix-monotone φ and justifies a use of the corresponding oracle in Theorem 24. Let us note that $\left[\exists w \in A^*\ \varphi(w\alpha) > \varphi(w\beta)\right] \iff \varphi(\alpha) > \varphi(\beta)$ if φ is additionally shift-deterministic.

7 Discussion

As Gimbert and Zielonka show by their characterization of the class of positionally determined payoffs [10], positional determinacy can always be proved by an inductive argument. Does the same hold for two other techniques that we have considered in the paper – the fixed point technique and the strategy improvement technique? The answer is positive in the continuous case, so this suggests that the answer might also be positive at least in some other special cases, for instance, for prefix-independent payoffs. E.g., for the mean payoff, a major example of a prefix-independent positionally determined payoff, both the strategy improvement and the fixed point arguments are applicable [13, 18].

These questions are specifically interesting for the strategy improvement argument. Indeed, strategy improvement usually leads to subexponential-time (randomized) algorithms for the strategy synthesis. So this resonates with a question of how hard strategy synthesis for a positionally determined payoff can be. Loosely speaking, do we have this subexponential bound for all positionally determined payoffs (as we do, by Theorem 24, for all such payoffs that are additionally continuous)?

Finally, is it possible to characterize positionally determined payoffs more explicitly (say, as in Theorem 17)? This question sounds more approachable in special cases, and a natural special case to start is again the prefix-independent case.

References

1. Henrik Björklund and Sergei Vorobyov. Combinatorial structure and randomized subexponential algorithms for infinite games. *Theoretical Computer Science*, 349(3):347–360, 2005.
2. Patricia Bouyer, Stéphane Le Roux, Youssouf Oualhadj, Mickael Randour, and Pierre Vandenhove. Games where you can play optimally with arena-independent finite memory. In *31st International Conference on Concurrency Theory (CONCUR 2020)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2020.
3. Cristian S Calude, Sanjay Jain, Bakhadyr Khoussainov, Wei Li, and Frank Stephan. Deciding parity games in quasipolynomial time. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing*, pages 252–263, 2017.
4. Thomas Colcombet and Damian Niwiński. On the positional determinacy of edge-labeled games. *Theoretical Computer Science*, 352(1-3):190–196, 2006.
5. Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, 1979.
6. E Allen Emerson and Charanjit S Jutla. Tree automata, mu-calculus and determinacy. In *FoCS*, volume 91, pages 368–377. Citeseer, 1991.
7. John Fearnley. *Strategy iteration algorithms for games and Markov decision processes*. PhD thesis, University of Warwick, 2010.
8. Hugo Gimbert. Pure stationary optimal strategies in markov decision processes. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 200–211. Springer, 2007.
9. Hugo Gimbert and Wieslaw Zielonka. When can you play positionnaly? In I. Fiala, V. Koubek, and J. Kratochvil, editors, *Mathematical Foundations of Computer Science 2004*, Lecture Notes in Comp. Sci. 3153, pages 686–697, Prague, Czech Republic, 2004. Springer. URL: https://hal.archives-ouvertes.fr/hal-00160436.
10. Hugo Gimbert and Wiesław Zielonka. Games where you can play optimally without any memory. In *International Conference on Concurrency Theory*, pages 428–442. Springer, 2005.
11. Hugo Gimbert and Wieslaw Zielonka. Applying blackwell optimality: priority mean-payoff games as limits of multi-discounted games. In *Logic and automata*, pages 331–356, 2008.
12. Erich Gradel and Wolfgang Thomas. *Automata, logics, and infinite games: a guide to current research*, volume 2500. Springer Science & Business Media, 2002.
13. Vladimir A Gurvich, Alexander V Karzanov, and LG Khachivan. Cyclic games and an algorithm to find minimax cycle means in directed graphs. *USSR Computational Mathematics and Mathematical Physics*, 28(5):85–91, 1988.
14. Nir Halman. Simple stochastic games, parity games, mean payoff games and discounted payoff games are all lp-type problems. *Algorithmica*, 49(1):37–50, 2007.
15. Thomas Dueholm Hansen, Peter Bro Miltersen, and Uri Zwick. Strategy iteration is strongly polynomial for 2-player turn-based stochastic games with a constant discount factor. *Journal of the ACM (JACM)*, 60(1):1–16, 2013.
16. Ronald A Howard. *Dynamic programming and markov processes*. The M.I.T. Press, 1960.
17. Marcin Jurdziński. Deciding the winner in parity games is in up∩co-up. *Information Processing Letters*, 68(3):119–124, 1998.
18. Elon Kohlberg. Invariant half-lines of nonexpansive piecewise-linear transformations. *Mathematics of Operations Research*, 5(3):366–372, 1980.
19. Alexander Kozachinskiy. Continuous positional payoffs. *arXiv preprint arXiv:2012.09047*, 2020.

20 Robert McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65(2):149–184, 1993.

21 Andrzej W Mostowski. Games with forbidden positions. Technical Report 78, Uniwersytet Gdánski, Instytut Matematyki, 1991.

22 Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.

23 Lloyd S Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 39(10):1095–1100, 1953.

24 Richard S Sutton, Andrew G Barto, et al. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.

25 Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343–359, 1996.

A The "Only If" Part of Theorem 4

Assume that φ is not prefix-monotone. Then for some $u, v \in A^*$ and $\alpha, \beta \in A^\omega$ we have

$$\varphi(u\alpha) > \varphi(u\beta) \text{ and } \varphi(v\alpha) < \varphi(v\beta). \tag{4}$$

First, notice that by continuity of φ we may assume that α and β are ultimately periodic. Indeed, consider any two sequences $\{\alpha_n\}_{n\in\mathbb{N}}$ and $\{\beta_n\}_{n\in\mathbb{N}}$ of ultimately periodic sequences from A^ω such that α_n and α (respectively, β_n and β) have the same prefix of length n. Then from continuity of φ (by Proposition 2) we have:

$$\lim_{n\to\infty} \varphi(u\alpha_n) = \varphi(u\alpha), \qquad \lim_{n\to\infty} \varphi(v\alpha_n) = \varphi(v\alpha),$$
$$\lim_{n\to\infty} \varphi(u\beta_n) = \varphi(u\beta), \qquad \lim_{n\to\infty} \varphi(v\beta_n) = \varphi(v\beta).$$

So if u, v, α, β violate prefix-monotonicity, then so do u, v, α_n, β_n for some $n \in \mathbb{N}$.

Now, if α, β are ultimately periodic, then $\alpha = p(q)^\omega$ and $\beta = w(r)^\omega$ for some $p, q, w, r \in A^*$. Consider an A-labeled game graph from Figure 1 (all nodes there are owned by Max).

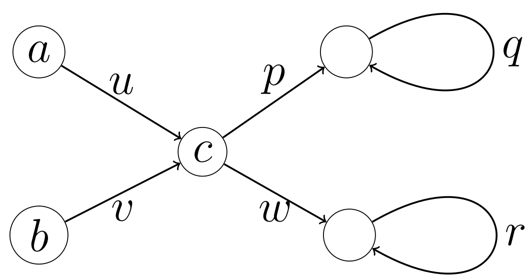

Figure 1 A game graph where φ is not positionally determined.

In this game graph there are two positional strategies of Max, one which from c goes by p and the other which goes from c by w. The first one is not optimal when the game starts in b, and the second one is not optimal when the game starts in a (because of (4)). So φ is not positionally determined in this game graph.

B Proof of Proposition 7

We only show that $\varphi(u\alpha) \le \max\{\varphi(u^\omega), \varphi(\alpha)\}$, the other inequality can be proved similarly. If $\varphi(u\alpha) \le \varphi(\alpha)$, then we are done. Assume now that $\varphi(u\alpha) > \varphi(\alpha)$. By repeatedly applying *(a)* we obtain $\varphi(u^{i+1}\alpha) \ge \varphi(u^i\alpha)$ for every $i \in \mathbb{N}$. In particular, for every $i \ge 1$ we get that $\varphi(u^i\alpha) \ge \varphi(u\alpha)$. By continuity of φ the limit of $\varphi(u^i\alpha)$ as $i \to \infty$ exists and equals $\varphi(u^\omega)$. Hence $\varphi(u^\omega) \ge \varphi(u\alpha)$.

C Proof of Proposition 11

Define a payoff $\psi\colon A^\omega \to \mathbb{R}$ as follows:

$$\psi(\gamma) = \sum_{w \in A^*} \left(\frac{1}{|A|+1}\right)^{|w|} \varphi(w\gamma), \qquad \gamma \in A^\omega. \tag{5}$$

First, why is ψ well-defined, i.e., why does this series converge? Since A^ω is compact, so is $\varphi(A^\omega) \subseteq \mathbb{R}$, because φ is continuous. Hence $\varphi(A^\omega) \subseteq [-W, W]$ for some $W > 0$ and (5) is bounded by the following absolutely converging series:

$$\sum_{w \in A^*} W \cdot \left(\frac{1}{|A|+1}\right)^{|w|}.$$

We shall show that ψ is continuous, prefix-monotone and shift-deterministic, and that $\varphi = g \circ \psi$ for some continuous non-decreasing $g\colon \psi(A^\omega) \to \mathbb{R}$.

Why is ψ continuous? Consider any $\alpha \in A^\omega$ and any infinite sequence $\{\beta_n\}_{n\in\mathbb{N}}$ of elements of A^ω such that for all n the sequences α and β_n coincide in the first n elements. We have to show that $\psi(\beta_n)$ converges to $\psi(\alpha)$ as $n \to \infty$. By definition:

$$\psi(\beta_n) = \sum_{w \in A^*} \left(\frac{1}{|A|+1}\right)^{|w|} \varphi(w\beta_n), \qquad \psi(\alpha) = \sum_{w \in A^*} \left(\frac{1}{|A|+1}\right)^{|w|} \varphi(w\alpha).$$

The first series, as we have seen, is bounded uniformly (in n) by an absolutely converging series. So it remains to note that the first series converges to the second one term-wise, by continuity of φ.

Why is ψ prefix-monotone? Let $\alpha, \beta \in A^\omega$. We have to show that either $\psi(u\alpha) \ge \psi(u\beta)$ for all $u \in A^*$ or $\psi(u\alpha) \le \psi(u\beta)$ for all $u \in A^*$.

Since φ is prefix-monotone, then either $\varphi(w\alpha) \ge \varphi(w\beta)$ for all $w \in A^*$ or $\varphi(w\alpha) \le \varphi(w\beta)$ for all $w \in A^*$. Up to swapping α and β we may assume that $\varphi(w\alpha) \ge \varphi(w\beta)$ for all $w \in A^*$. Then for any $u \in A^*$ the difference

$$\psi(u\alpha) - \psi(u\beta) = \sum_{w \in A^*} \left(\frac{1}{|A|+1}\right)^{|w|} \big[\varphi(wu\alpha) - \varphi(wu\beta)\big]$$

consists of non-negative terms. Hence $\psi(u\alpha) \ge \psi(u\beta)$ for all $u \in A^*$, as required.

Why is ψ shift-deterministic? Take any $a \in A$ and $\beta, \gamma \in A^\omega$ with $\psi(\beta) = \psi(\gamma)$. We have to show that $\psi(a\beta) = \psi(a\gamma)$. Indeed, assume that

$$0 = \psi(\beta) - \psi(\gamma) = \sum_{w \in A^*} \left(\frac{1}{|A|+1}\right)^{|w|} \big[\varphi(w\beta) - \varphi(w\gamma)\big].$$

If this series contains a non-zero term, then it must contain a positive term and a negative term. But this contradicts prefix-monotonicity of φ. So all the terms in this series must be 0. The same then must hold for a series:

$$\psi(a\beta) - \psi(a\gamma) = \sum_{w\in A^*} \left(\frac{1}{|A|+1}\right)^{|w|} \big[\varphi(wa\beta) - \varphi(wa\gamma)\big]$$

(all the terms in this series also appear in the series for $\psi(\beta) - \psi(\gamma)$). So we must have $\psi(a\beta) = \psi(a\gamma)$.

Why $\varphi = g \circ \psi$ for some continuous non-decreasing $g\colon \psi(A^\omega) \to \mathbb{R}$? Let us first show that

$$\varphi(\alpha) > \varphi(\beta) \implies \psi(\alpha) > \psi(\beta) \text{ for all } \alpha, \beta \in A^\omega. \tag{6}$$

Indeed, if $\varphi(\alpha) > \varphi(\beta)$, then we also have $\varphi(w\alpha) \geq \varphi(w\beta)$ for every $w \in A^*$, by prefix-monotonicity of φ. Now, by definition,

$$\psi(\alpha) - \psi(\beta) = \sum_{w\in A^*} \left(\frac{1}{|A|+1}\right)^{|w|} \big[\varphi(w\alpha) - \varphi(w\beta)\big].$$

All the terms in this series are non-negative, and the term corresponding to the empty w is strictly positive. So we have $\psi(\alpha) > \psi(\beta)$, as required.

Now, let us demonstrate that (6) implies that $\varphi = g \circ \psi$ for some non-decreasing $g\colon \psi(A^\omega) \to \mathbb{R}$. Namely, define g as follows. For $x \in \psi(A^\omega)$ take an arbitrary $\gamma \in \psi^{-1}(x)$ and set $g(x) = \varphi(\gamma)$. First, why do we have $\varphi = g \circ \psi$? By definition, $g(\psi(\alpha)) = \varphi(\gamma)$ for some $\gamma \in A^\omega$ with $\psi(\alpha) = \psi(\gamma)$. By (6) we also have $\varphi(\alpha) = \varphi(\gamma)$, so $g(\psi(\alpha)) = \varphi(\gamma) = \varphi(\alpha)$, as required. Now, why is g non-decreasing? I.e., why for all $x, y \in \psi(A^\omega)$ we have $x \leq y \implies g(x) \leq g(y)$? Indeed, $g(x) = \varphi(\gamma_x), g(y) = \varphi(\gamma_y)$ for some $\gamma_x \in \psi^{-1}(x)$ and $\gamma_y \in \psi^{-1}(y)$. Now, since $x \leq y$, we have $x = \psi(\gamma_x) \leq \psi(\gamma_y) = y$. By taking the contraposition of (6) we get that $g(x) = \varphi(\gamma_x) \leq \varphi(\gamma_y) = g(y)$, as required.

Finally, we show that any $g\colon \psi(A^\omega) \to \mathbb{R}$ with $\varphi = g \circ \psi$ must be continuous. For that we show that $|g(x) - g(y)| \leq |x - y|$ for all $x, y \in \psi(A^\omega)$. Take any $\alpha, \beta \in A^\omega$ with $x = \psi(\alpha)$ and $y = \psi(\beta)$. By prefix-monotonicity of φ we have that either $\varphi(w\alpha) \geq \varphi(w\beta)$ for all $w \in A^*$ or $\varphi(w\alpha) \leq \varphi(w\beta)$ for all $w \in A^*$. Up to swapping x and y we may assume that the first option holds. Then

$$\psi(\alpha) - \psi(\beta) = \sum_{w\in A^*} \left(\frac{1}{|A|+1}\right)^{|w|} \big[\varphi(w\alpha) - \varphi(w\beta)\big] \geq \varphi(\alpha) - \varphi(\beta) \geq 0.$$

On the left here we have $x - y$, and on the right we have $\varphi(\alpha) - \varphi(\beta) = g \circ \psi(\alpha) - g \circ \psi(\beta) = g(x) - g(y)$.