
CONTINUOUS POSITIONAL PAYOFFS

ALEXANDER KOZACHINSKIY

Steklov Mathematical Institute of Russian Academy of Sciences, Moscow, Russia
e-mail address: akozachinskiy@mi-ras.ru

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.

1. INTRODUCTION

We study two-player turn-based games on finite directed graphs. In these games, two players called Max and Min travel over nodes of a given graph along its edges for infinitely many turns. In each turn, one of the players decides where to go next – which of the two depends on a predetermined partition of the nodes between the players.

After infinitely many turns, we get an infinite path in our graph. Each infinite path is mapped to a real number called its *reward*, according to some *payoff function* (or, for brevity, a payoff). The larger the reward is the more Max is happy; on the contrary, Min wants to minimize the reward.

We consider only payoffs that are defined through *edge labels*. Namely, we first fix some finite set A of labels. Then we label edges of our game graph by elements of A . After this, any function $\varphi: A^\omega \rightarrow \mathbb{R}$ can be viewed as a payoff in our graph. Namely, it takes an infinite path, considers an infinite word over A “written” on this path, and applies φ to this infinite word.

Fix a *strategy* of one of the players (that is, an instruction how to play in all possible developments of the game). If this is a strategy of Max, then its *value* is the infimum of the payoff over all infinite paths that can occur in a play with this strategy. Similarly, if this is a strategy of Min, then its value is the supremum of the payoff over all infinite paths that can occur in a play with this strategy. Values are a standard way of measuring the “worst-case” quality of a strategy.

Key words and phrases: Games on graphs, positional strategies, continuous payoffs.

* An extended abstract of this paper was presented at the CONCUR 2021 conference.

A strategy of Max is called *optimal* if its value is at least as large as the value of any other Max’s strategy. Similarly, a strategy of Min is called *optimal* if its value does not exceed the value of any other Min’s strategy.

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 so-called *determined* payoffs. These are payoffs that have an equilibrium in all game graphs.

Games that we are studying proceed for infinitely many turns, so strategies in them might be rather complicated. An area of *strategy complexity* classifies payoffs according to their “simplicity”. A “simple” payoff always admits an equilibrium of two “simple” strategies. This has been studied in various settings [2]. We study, perhaps, the most well-established one, where by “simple” strategies we understand *positional* strategies.

A strategy of one of the players is called positional if for every node there exists a single out-going edge such that this strategy always uses this edge from this node (of course, we only require this for nodes from where the corresponding player is the one to move). Essentially, a positional strategy is a strategy with no memory – at every location, it completely ignores the previous development of the game. Now, a payoff is called *positionally determined* if every game graph has an equilibrium of two positional strategies w.r.t. this payoff.

A lot of works are devoted to concrete positionally determined payoffs that are of interest in other areas of computer science. Classical examples of such payoffs are parity payoffs, mean payoffs and (multi-)discounted payoffs [5, 20, 19, 22]. Their applications range from logic, verification and finite automata theory [6, 12] to decision-making [21, 23] 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 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. Unfortunately, this property is far more technical than the fairly mixing property, and it is hard to use it in practice. Still, this property has a remarkable feature: if a payoff does not satisfy it, then this payoff is not positionally determined in some *one-player* game graph (i.e., in a game graph where one of the players owns all the nodes). As Gimbert and Zielonka indicate, this means that to establish the positional determinacy of a payoff, it is enough to do so only for one-player game graphs.

Unfortunately, these results do not quite answer how positionally determined payoffs are arranged in general. The goal of the present paper is to make progress in this direction – at least for payoffs that satisfy some natural additional properties. One such property 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 the input word. For instance, the parity and the mean payoffs are prefix-independent.

In [9], Gimbert and Zielonka briefly mention another interesting additional property, 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. A payoff is continuous if its range converges to just a single point as more and more initial letters of its input (which an

infinite word over the set of 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 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 payoffs, positional determinacy is equivalent to a simple property which we call *prefix-monotonicity*. A payoff φ is prefix-monotone if there are no two infinite words α and β and no two finite words x and y such that $\varphi(x\alpha) > \varphi(x\beta)$ and $\varphi(y\alpha) < \varphi(y\beta)$.

A proof of the fact that any continuous positionally determined payoff is prefix-monotone can be found in Section 3. We give three different proofs of the opposite direction of our main result, 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 Mean Payoff Games.

This argument can be found in Section 4.

- *A fixed point argument.* Then we give a proof which uses a fixed point approach due to Shapley [22]. 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.

This argument can be found in Section 5.

- *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.

This argument can be found in Section 7.

The simplest way to obtain our main result is via the inductive argument (at the cost of appealing without a proof to the sufficient condition of Gimbert and Zielonka). In turn, two other proofs give the following additional results:

- Using the fixed point approach, in Section 6 we give an explicit description of the set of continuous positionally determined payoffs. Namely, it turns out that all continuous positionally determined payoffs are, in a sense, non-affine multi-discounted payoffs. We use this to give an example of a positionally determined payoff which does not reduce to multi-discounted payoffs in an “algorithmic sense”.
- Using the strategy improvement approach, in Section 8 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.

We also 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 a “potential transformation lemma” [15, Lemma 3.6].

Finally, we study continuous payoffs that are positional in *stochastic games*. Namely, in Section 9 we show that any continuous payoff which is positional in Markov Decision Processes is multi-discounted. On the other hand, it is classical multi-discounted games are positional even in two-player stochastic games. Using it, we disprove the following conjecture of Gimbert [8]: “Any payoff function which is positional for the class of non-stochastic one-player games is positional for the class of Markov decision processes”.

2. PRELIMINARIES

2.1. Notation. We denote the function composition by \circ . For two sets A and B by B^A we denote the set of all functions from A to B . We write $C = A \sqcup B$ for three sets A, B, C if A and B are disjoint and $C = A \cup B$.

Take any set A . By A^* we denote the set of all finite words over the alphabet A . By A^+ we denote the set of all non-empty finite words over the alphabet A . Finally, by A^ω we denote the set of all infinite words over the alphabet A . For $w \in A^*$, we let $|w|$ be the length of w . For $\alpha \in A^\omega$ we define $|\alpha| = \infty$.

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 $u \in A^*$, by uA^ω we denote the set $\{u\alpha \mid \alpha \in A^\omega\}$. Alternatively, uA^ω is the set of all $\beta \in A^\omega$ such that u is a prefix of β .

For $u \in A^*$ and $k \in \mathbb{N}$ we define

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

In turn, if $u \in A^+$, we let $u^\omega \in A^\omega$ be a infinite word obtained by repeating u infinitely many times. We call $\alpha \in A^\omega$ ultimately periodic if $\alpha = uv^\omega$ for some $u \in A^*, v \in A^+$.

2.2. Deterministic infinite duration games on finite directed graphs.

Definition 2.1. Let A be a finite set. A tuple $G = \langle V, V_{\text{Max}}, V_{\text{Min}}, E \rangle$ is called an *A-labeled game graph* if the following conditions hold:

- $V, V_{\text{Max}}, V_{\text{Min}}, E$ are four finite sets with $V = V_{\text{Max}} \sqcup V_{\text{Min}}$, $E \subseteq V \times A \times V$;
- for every $s \in V$ there exist $a \in A$ and $t \in V$ such that $(s, a, t) \in E$.

Elements of V are called nodes of G . Nodes from V_{Max} (resp., V_{Min}) are called Max’s nodes (resp., Min’s nodes). Elements of E are called edges of E . For an edge $e = (s, a, t) \in E$ we define $\text{source}(e) = s$, $\text{lab}(e) = a$, $\text{target}(e) = t$. We imagine e as an arrow from $\text{source}(s)$ to $\text{target}(e)$ with the label $\text{lab}(a)$.

We will apply the function lab not only to individual edges, but also to arbitrary finite or infinite sequences of edges. Namely, given a sequence of edges, we first apply lab to its elements, and then concatenate the resulting letters from A in the same order as in the sequence. We will get a word over A of the same length as the initial sequence of edges.

The out-degree of a node $v \in V$ is the number of $e \in E$ with $\text{source}(e) = v$. The last requirement in the definition of an A -labeled game graph means that the out-degree of every node must be positive.

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

For technical convenience, we also consider 0-length paths. Namely, for every node $s \in V$ we consider a 0-length path λ_s , for which we define $\text{source}(\lambda_s) = \text{target}(\lambda_s) = s$. Hence, there are $|V|$ different 0-length paths.

If p and q are two paths of positive length and p is finite, then we can consider their concatenation pq . This will be a path if and only if $\text{target}(p) = \text{source}(q)$.

Now, if $p = \lambda_s$ is a 0-length path, then $\lambda_s q$ is a path if and only if $\text{source}(q) = s$. In this case, $\lambda_s q = q$. Similarly, if $q = \lambda_s$ is a 0-length path, then $p \lambda_s$ is a path if and only if $\text{target}(p) = s$. In this case, $p \lambda_s = p$.

Fix a finite set A and an A -labeled game graph $G = \langle V, V_{\text{Max}}, V_{\text{Min}}, E \rangle$. Consider the following infinite-duration game (IDG for short) which is played over G . Players are called Max and Min. Positions of the game are finite paths in G (informally, these are possible finite developments of the game). Possible starting positions are paths of length 0. Positions from where Max (resp., Min) is the one to move are finite paths with $\text{target}(p) \in V_{\text{Max}}$ (resp., $\text{target}(p) \in V_{\text{Min}}$).

The set of moves available at a position p is the set $\{e \in E \mid \text{source}(e) = \text{target}(p)\}$ of edges that come out of the endpoint of 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 position p with $\text{target}(p) \in V_{\text{Max}}$ some move available at p . Similarly, a Min's strategy τ in a game graph G is a mapping, assigning to every position p with $\text{target}(p) \in V_{\text{Min}}$ some move available at p .

Let $\mathcal{P} = e_1 e_2 e_3 \dots$ be an infinite path in G . We say that \mathcal{P} is *consistent* with a Max's strategy σ if for every finite prefix p of \mathcal{P} with $\text{target}(p) \in V_{\text{Max}}$ it holds that $\sigma(p)$ is the next edge of \mathcal{P} after p . For $s \in V$ and for a Max's strategy σ we let $\text{Cons}(s, \sigma)$ be a set of all infinite paths in G that start in s and are consistent with σ . We use a similar terminology and notation for strategies of Min.

Given a Max's strategy σ , a Min's strategy τ and $s \in V$, *the play of σ and τ from s* is an infinite path $\mathcal{P}_s^{\sigma, \tau}$ which can be obtained as follows. First, set $p = \lambda_s$. Then repeat the following infinitely many times. If $\text{target}(p) \in V_{\text{Max}}$, extend it by the edge $\sigma(p)$. Similarly, if $\text{target}(p) \in V_{\text{Min}}$, extend it by the edge $\tau(p)$. The resulting infinite path will be $\mathcal{P}_s^{\sigma, \tau}$. It is not hard to see that $\mathcal{P}_s^{\sigma, \tau}$ is a unique element of the intersection $\text{Cons}(s, \sigma) \cap \text{Cons}(s, \tau)$.

A Max's strategy σ in an A -labeled game graph $G = \langle V, V_{\text{Max}}, V_{\text{Min}}, E \rangle$ is called *positional* if $\sigma(p) = \sigma(q)$ for all finite paths p and q in G with $\text{target}(p) = \text{target}(q) \in V_{\text{Max}}$. For a positional strategy σ of Max and for $u \in V_{\text{Max}}$, we let $\sigma(u)$ be the move of σ from any position whose endpoint is u . That is, we can view a positional Max's strategy σ as a function $\sigma: V_{\text{Max}} \rightarrow E$. Obviously, this function satisfies $\text{source}(\sigma(u)) = u$ 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 $\text{source}(e) \in V_{\text{Min}}$ or $\text{source}(e) \in V_{\text{Max}}$, $e = \sigma(\text{source}(e))$. 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 $\text{source}(e) \in V_{\text{Max}}$ or $\text{source}(e) \in V_{\text{Min}}$, $e = \tau(\text{source}(e))$. The set of edges that are consistent with a Min's positional strategy τ is denoted by E_τ .

Fix a finite set A and an A -labeled game graph $G = \langle V, V_{\text{Max}}, V_{\text{Min}}, E \rangle$. Take any function $\varphi: A^\omega \rightarrow \mathbb{R}$ to which we will refer to as a *payoff*. Given a Max's strategy σ in G , its *value* in a node $s \in V$ (w.r.t. φ) is defined as follows:

$$\mathbf{Val}[\sigma](s) = \inf \varphi \circ \text{lab}(\text{Cons}(s, \sigma)).$$

That is, we first take all infinite paths from s that are consistent with σ . Then we consider all infinite words over A that are “written” on these paths. The set of these words is $\text{lab}(\text{Cons}(s, \sigma))$. Finally, we take the infimum of our payoff over this set.

Similarly, if τ is a Min's strategy in G , then the value of τ in a node $s \in V$ (w.r.t. φ) is the following quantity:

$$\mathbf{Val}[\tau](s) = \sup \varphi \circ \text{lab}(\text{Cons}(s, \tau)).$$

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

$$\mathbf{Val}[\sigma](s) \leq \varphi \circ \text{lab}(\mathcal{P}_s^{\sigma, \tau}) \leq \mathbf{Val}[\tau](s).$$

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

A pair (σ, τ) of a Max's strategy σ and a Min's strategy τ is called an *equilibrium* if $\mathbf{Val}[\sigma](s) = \mathbf{Val}[\tau](s)$ for every $s \in V$. It is easy to see that any strategy appearing in an equilibrium is optimal. On the other hand, if at least one equilibrium exists, then the following holds: the Cartesian product of the set of optimal Max's strategies and the set of optimal Min's strategies is 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 . We say that φ is *positionally determined* if every A -labeled game graph contains an equilibrium (w.r.t. φ) of two positional strategies.

Proposition 2.2. *If A is a finite set, $\varphi: A^\omega \rightarrow \mathbb{R}$ is a payoff and $g: \varphi(A^\omega) \rightarrow \mathbb{R}$ is a non-decreasing¹ function, then any equilibrium w.r.t. φ is also an equilibrium w.r.t. $g \circ \varphi$.*

Proof. Let (σ, τ) be an equilibrium w.r.t. φ , where σ is a Max's strategy and τ is a Min's strategy. Our goal is to show that (σ, τ) is also an equilibrium w.r.t. $g \circ \varphi$.

By definition, the values of σ and τ w.r.t. φ coincide. That, for every node s , we have:

$$\inf \varphi \circ \text{lab}(\text{Cons}(s, \sigma)) = \sup \varphi \circ \text{lab}(\text{Cons}(s, \tau)).$$

We have to derive from this that the values of σ and τ w.r.t. $g \circ \varphi$ also coincide. That is, we have to show that:

$$\inf g \circ \varphi \circ \text{lab}(\text{Cons}(s, \sigma)) = \sup g \circ \varphi \circ \text{lab}(\text{Cons}(s, \tau)) \quad (2.1)$$

for every node s . The sets $\varphi \circ \text{lab}(\text{Cons}(s, \sigma))$ and $\varphi \circ \text{lab}(\text{Cons}(s, \tau))$ have a common element $\varphi \circ \text{lab}(\mathcal{P}_s^{\sigma, \tau})$. Since the infimum of the first set equals the supremum of the second set, their common element $\varphi \circ \text{lab}(\mathcal{P}_s^{\sigma, \tau})$ must be the minimum of the first set and the maximum of the second set. Due to the fact that the function g is non-decreasing, we have that $g(\varphi \circ \text{lab}(\mathcal{P}_s^{\sigma, \tau}))$ is the minimum of $g \circ \varphi \circ \text{lab}(\text{Cons}(s, \sigma))$ and the maximum $g \circ \varphi \circ \text{lab}(\text{Cons}(s, \tau))$. This implies (2.1). \square

¹Throughout the paper, we call a function $f: S \rightarrow \mathbb{R}$, $S \subseteq \mathbb{R}$ non-decreasing if for all $x, y \in S$ we have $x \leq y \implies f(x) \leq f(y)$.

Corollary 2.3. *If A is a finite set, $\varphi: A^\omega \rightarrow \mathbb{R}$ is a positionally determined payoff and $g: \mathbb{R} \rightarrow \mathbb{R}$ is a non-decreasing function, then $g \circ \varphi$ is a positionally determined payoff.*

2.3. 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

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

where $S \subseteq A^*$. When we say that a payoff $\varphi: A^\omega \rightarrow \mathbb{R}$ 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 2.4. *Let A be a finite set. A payoff $\varphi: A^\omega \rightarrow \mathbb{R}$ 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 $n \geq 1$ it holds that α and β_n have the same prefixes of length n , then $\lim_{n \rightarrow \infty} \varphi(\beta_n)$ exists and equals $\varphi(\alpha)$.*

Proof. See Appendix A. □

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

2.4. MDPs. This subsection concerns stochastic games, but we deal with them only in Section 9. So for the rest of our results, one can skip this subsection.

In fact, we will need only one-player stochastic games, also known as Markov Decision Processes (MDPs). We will follow a formalization of Gimbert [8].

We use the following notation. Let A be a finite set. By \mathfrak{S}_A^{bor} we mean the Borel σ -algebra on A^ω (generated by the product topology from the previous subsection). By $\Delta(S)$ we denote the set of all probability distributions over a finite set S .

Definition 2.5. Let A be a finite set. An A -labeled MDP is a tuple $\mathcal{M} = \langle S, \text{Act}, \text{lab} \rangle$, where

- S is a finite set of *states* of \mathcal{M} ;
- $\text{Act} \subseteq S \times \Delta(S)$ is a finite set of *actions* of \mathcal{M} ;
- $\text{lab}: \text{Act} \times S \rightarrow A$ is the *labeling function* of \mathcal{M} ;

such that for every $s \in S$ there exists $P \in \Delta(S)$ such that $(s, P) \in \text{Act}$.

Given an A -labeled MDP $\mathcal{M} = \langle S, \text{Act}, \text{lab} \rangle$, we imagine that there is a single player called Max traveling over the states of \mathcal{M} . When Max is in a state $s \in S$, he considers the set of all actions of \mathcal{M} whose first coordinate is s (by definition, this set is non-empty for every $s \in S$). He chooses one such action (s, P) . Then Max samples his next location according to P . This continues for infinitely many turns.

For $e = (s, P) \in \text{Act}$, we define $\text{source}(e) = s$ and $\text{Dist}[e] = P$.

The set $T = \text{Act} \times S$, which is the domain of the function lab , is called the set of *transitions* of \mathcal{M} . Informally, transitions describe what happens in one turn. Namely, a

transition $(e, s) \in \text{Act} \times S$ means that in the beginning of a turn, Max was in the state $\text{source}(e)$, then he took the action e , and this led him to the state s .

Consistent sequences of transitions are called *histories*. Namely, a non-empty sequence $h = (e_1, s_1)(e_2, s_2)(e_3, s_3) \dots \in T^+ \cup T^\omega$ is called a history if for every $2 \leq i \leq |h|$, we have $s_{i-1} = \text{source}(e_i)$. We set $\text{source}(h) = \text{source}(e_1)$ and, if h is finite, $\text{target}(h) = s_{|h|}$. We also map each state $s \in S$ to a 0-length history λ_s with $\text{source}(\lambda_s) = \text{target}(\lambda_s) = s$. These histories correspond to $|S|$ possible starting positions of Max.

A strategy σ of Max is a mapping, which to every finite history h assigns an action $\sigma(h) \in \text{Act}$ such that $\text{target}(h) = \text{source}(\sigma(h))$. Informally, $\sigma(h)$ is the action which, according to σ , Max takes after h .

Given $s \in S$, a strategy σ defines a function $P_s^\sigma: T^* \rightarrow [0, 1]$. Informally, $P_s^\sigma(h)$ is the probability that we will see a history h if Max starts in s and plays according to σ . It can be defined inductively.

First, we set $P_s^\sigma(\text{empty word}) = 1$. The empty word here corresponds to the initial history λ_s . Now, given $(e, s_1) \in T$, we set $P_s^\sigma((e, s_1)) = 0$ if $e \neq \sigma(\lambda_s)$ and $P_s^\sigma((e, s_1)) = \text{Dist}[e](s_1)$ if $e = \sigma(\lambda_s)$. That is, if e is not the action played by Max according to σ in the starting position, then the probability of the transition (e, s_1) is 0. Otherwise, the probability of (e, s_1) is the probability that the action $e = \sigma(\lambda_s)$ brings us to s_1 .

More generally, assume that P_s^σ is already defined for all $h \in T^1 \cup T^2 \dots \cup T^n$. Take any $h = (e_1, s_1) \dots (e_n, s_n)(e_{n+1}, s_{n+1})$. If h is not a history, or if $\text{source}(h) \neq s$, we set $P_s^\sigma(h) = 0$. Similarly, if $e_{n+1} \neq \sigma((e_1, s_1) \dots (e_n, s_n))$, that is, if e_{n+1} is not the action played by σ after $(e_1, s_1) \dots (e_n, s_n)$, then we set $P_s^\sigma(h) = 0$. Finally, if h is a history with $\text{source}(h) = s$, and if $e_{n+1} = \sigma((e_1, s_1) \dots (e_n, s_n))$, then we set

$$P_s^\sigma(h) = P_s^\sigma((e_1, s_1) \dots (e_n, s_n)) \cdot \text{Dist}[e_{n+1}](s_{n+1}).$$

Obviously, we have $P_s^\sigma(h) = \sum_{t \in T} P_s^\sigma(ht)$ for every $h \in T^*$. Hence, by the Caratheodory's extension theorem, there is a unique probability measure \mathcal{P}_s^σ on $\mathfrak{S}_T^{\text{bor}}$ such that

$$\mathcal{P}_s^\sigma(hT^\omega) = P_s^\sigma(h)$$

for every $h \in T^*$. Intuitively, \mathcal{P}_s^σ is the probability distribution over infinite histories, generated by playing σ for infinitely many turns, starting from s .

Any sequence of transitions $h \in T^* \cup T^\omega$ can be mapped to a word $\text{lab}(h)$ over the set of labels by setting:

$$\text{lab}(h) = \text{lab}(h_1)\text{lab}(h_2)\text{lab}(h_3) \dots$$

Now, fix a payoff function $\varphi: A^\omega \rightarrow \mathbb{R}$. It maps any infinite history $\mathcal{H} \in T^\omega$ to its reward, defined as $\varphi(\text{lab}(\mathcal{H}))$. Max wants a strategy which maximizes the expected value of the reward. That is, Max wants to attain

$$\mathbb{E}_{\mathcal{H} \sim \mathcal{P}_s^\sigma} [\varphi \circ \text{lab}(\mathcal{H})] \rightarrow \max \tag{2.2}$$

over his strategies σ , for all $s \in S$. This expectation is well defined if $\varphi \circ \text{lab}$ is bounded and measurable with respect to $\mathfrak{S}_T^{\text{bor}}$. Since $\text{lab}: T^\omega \rightarrow A^\omega$ is continuous, it is well-defined if $\varphi: A^\omega \rightarrow \mathbb{R}$ is bounded and measurable with respect to $\mathfrak{S}_A^{\text{bor}}$.

For brevity, we will abbreviate the expectation in (2.2) by

$$\mathbb{E}\varphi \circ \text{lab}(\mathcal{P}_s^\sigma).$$

Definition 2.6. Let A be a finite set and $\varphi: A^\omega \rightarrow \mathbb{R}$ be a bounded measurable payoff.

We say that a strategy σ in an A -labeled MDP \mathcal{M} is *optimal* if for any strategy σ' and for any state s of \mathcal{M} we have:

$$\mathbb{E}\varphi \circ \text{lab}(\mathcal{P}_s^\sigma) \geq \mathbb{E}\varphi \circ \text{lab}(\mathcal{P}_s^{\sigma'}).$$

We say that a strategy σ in an A -labeled MDP \mathcal{M} is *positional* if for any two finite histories h_1 and h_2 in \mathcal{M} we have $\text{target}(h_1) = \text{target}(h_2) \implies \sigma(h_1) = \sigma(h_2)$.

We say that φ is *positionally determined in MDPs* if every A -labeled MDP has an optimal positional strategy w.r.t. φ .

In the paper, we will use this definition only for continuous φ – they all are, of course, bounded (by compactness of A^ω) and measurable.

3. STATEMENT OF THE MAIN RESULT AND ITS “ONLY IF” PART

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

Definition 3.1. Let A be a finite set. A payoff $\varphi: A^\omega \rightarrow \mathbb{R}$ 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 is continuous, unless it takes just 1 value.)

Theorem 3.2. *Let A be a finite set and $\varphi: A^\omega \rightarrow \mathbb{R}$ 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 is proved below in this section. Three different proofs of the “if” part of Theorem 3.2 are given in, respectively, Sections 4, 5 and 7. As an illustration of our result, we first give a formal definition of multi-discounted payoffs and show that they are continuous and prefix-monotone.

Definition 3.3. Let A be a finite set. Then a payoff $\varphi: A^\omega \rightarrow \mathbb{R}$ is **multi-discounted** if there are functions $\lambda: A \rightarrow [0, 1)$ and $w: A \rightarrow \mathbb{R}$ such that

$$\varphi(a_1 a_2 a_3 \dots) = \sum_{n=1}^{\infty} \lambda(a_1) \cdot \dots \cdot \lambda(a_{n-1}) \cdot w(a_n) \quad (3.1)$$

for all $a_1 a_2 a_3 \dots \in A^\omega$.

Proposition 3.4. *All multi-discounted payoffs are continuous and prefix-monotone.*

Proof. Let A be a finite set and $\varphi: A^\omega \rightarrow \mathbb{R}$ be a multi-discounted payoff, defined by $\lambda: A \rightarrow [0, 1)$ and $w: A \rightarrow \mathbb{R}$. Take any $W > 0$ such that $\lambda(a) < 1 - \frac{1}{W}$ and $|w(a)| < W$ for every $a \in A$.

Let us first show that φ is continuous. Take any $\alpha, \beta \in A^\omega$ that coincide in the first n elements. It is sufficient to bound the difference $|\varphi(\alpha) - \varphi(\beta)|$ by some quantity which depends only on n and tends to 0 as $n \rightarrow \infty$. First, observe that the value of φ never exceeds $W \cdot \frac{1}{1 - (1 - \frac{1}{W})} = W^2$. Now, let $u = a_1 a_2 \dots a_n \in A^n$ be the first n letters of α and β . Then $\alpha = u\alpha', \beta = u\beta'$ for some $\alpha', \beta' \in A^\omega$. It is not hard to derive from (3.1) that:

$$\varphi(\alpha) - \varphi(\beta) = \varphi(u\alpha') - \varphi(u\beta') = \lambda(a_1) \cdot \dots \cdot \lambda(a_n) \cdot (\varphi(\alpha') - \varphi(\beta')). \quad (3.2)$$

This means that the difference $|\varphi(\alpha) - \varphi(\beta)|$ is bounded by $(1 - \frac{1}{W})^n \cdot W^2$. This quantity tends to 0 as $n \rightarrow \infty$. Hence, φ is continuous.

Equation (3.2) also implies that φ is prefix-monotone. Indeed, it gives that for any $u \in A^*$ and $\beta, \gamma \in A^\omega$ there exists $\lambda \geq 0$ such that $\varphi(u\beta) - \varphi(u\gamma) = \lambda \cdot (\varphi(\beta) - \varphi(\gamma))$. This equality gives us that:

$$\varphi(u\beta) > \varphi(u\gamma) \implies \varphi(\beta) > \varphi(\gamma), \quad \varphi(u\beta) < \varphi(u\gamma) \implies \varphi(\beta) < \varphi(\gamma).$$

Hence, there are no $u, v \in A^*$ and $\beta, \gamma \in A^\omega$ such that $\varphi(u\beta) > \varphi(u\gamma)$ and $\varphi(v\beta) < \varphi(v\gamma)$. \square

Proof of the “only if” part of Theorem 3.2. Assume that φ is not prefix-monotone. Then for some $u, v \in A^*$ and $\alpha, \beta \in A^\omega$ we have

$$\varphi(u\alpha) > \varphi(u\beta) \text{ and } \varphi(v\alpha) < \varphi(v\beta). \quad (3.3)$$

First, notice that by the 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 infinite words over A such that α_n and α (respectively, β_n and β) have the same prefix of length n . Then, by continuity of φ , we have:

$$\begin{aligned} \lim_{n \rightarrow \infty} \varphi(u\alpha_n) &= \varphi(u\alpha), & \lim_{n \rightarrow \infty} \varphi(v\alpha_n) &= \varphi(v\alpha), \\ \lim_{n \rightarrow \infty} \varphi(u\beta_n) &= \varphi(u\beta), & \lim_{n \rightarrow \infty} \varphi(v\beta_n) &= \varphi(v\beta). \end{aligned}$$

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

Thus, we assume from now on that 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 its nodes are owned by Max).




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

There are two positional strategies of Max in this game graph, one which goes along p from c , and the other which goes along w from c . 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 (3.3)). So φ is not positionally determined in this game graph. \square

Remark 3.5. In this argument, it is crucial that our definition of positional determinacy is “uniform”. That is, 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 for continuous payoffs. At the same time, we are not even aware of a payoff which is positional in the non-uniform sense, but not in the uniform sense.

4. INDUCTIVE ARGUMENT

In this section, we show that any continuous prefix-monotone payoff is positionally determined, using the following sufficient condition due to Gimbert and Zielonka [9, Theorem 1]:

Proposition 4.1. *Let A be a finite set. Any payoff $\varphi: A^\omega \rightarrow \mathbb{R}$, satisfying the following three conditions:*

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

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

- **(c)** for any infinite sequence $\{x_n \in A^+\}_{n=0}^\infty$ it holds that:

$$\begin{aligned} \min\{\varphi(x_0x_2x_4\dots), \varphi(x_1x_3x_5\dots), \inf_{n \geq 0} \varphi(x_n^\omega)\} &\leq \varphi(x_0x_1x_2\dots) \\ &\leq \max\{\varphi(x_0x_2x_4\dots), \varphi(x_1x_3x_5\dots), \sup_{n \geq 0} \varphi(x_n^\omega)\} \end{aligned}$$

is positionally determined.

We observe that in case of continuous payoffs, one can get rid of the conditions **(b)** and **(c)** in this Proposition. A weaker version of this statement was proved in the on-line version of [9]. Namely, it was shown there that one can get rid of the condition **(c)** for continuous payoffs.

Proposition 4.2. *For continuous payoffs, the condition **(a)** of Proposition 4.1 implies the conditions **(b)** and **(c)** of Proposition 4.1.*

Proof. Take any finite set A and any continuous payoff $\varphi: A^\omega \rightarrow \mathbb{R}$ satisfying the condition **(a)** of Proposition 4.1. We first show that φ satisfies the condition **(b)** of this proposition. We will only show that $\varphi(u\alpha) \leq \max\{\varphi(u^\omega), \varphi(\alpha)\}$, the other inequality from this condition can be proved similarly. If $\varphi(u\alpha) \leq \varphi(\alpha)$, then we are done. Assume now that $\varphi(u\alpha) > \varphi(\alpha)$. By repeatedly applying **(a)**, we obtain $\varphi(u^{i+1}\alpha) \geq \varphi(u^i\alpha)$ for every $i \in \mathbb{N}$. In particular, for every $i \geq 1$ we get that $\varphi(u^i\alpha) \geq \varphi(u\alpha)$. By continuity of φ , we have that $\lim_{i \rightarrow \infty} \varphi(u^i\alpha) = \varphi(u^\omega)$. Hence, $\varphi(u^\omega) \geq \varphi(u\alpha)$.

Now we show that φ satisfies the condition **(c)** of Proposition 4.1. We will only show that $\varphi(x_0x_1x_2\dots) \leq \max\{\varphi(x_0x_2x_4\dots), \varphi(x_1x_3x_5\dots), \sup_{n \geq 0} \varphi(x_n^\omega)\}$, the other inequality from this condition has the same proof. Namely, we will show that if $\varphi(x_0x_1x_2\dots) > \sup_{n \geq 0} \varphi(x_n^\omega)$, then $\varphi(x_0x_1x_2\dots) \leq \varphi(x_0x_2x_4\dots)$. Note that this claim is stronger than we need.

First, we show that $\varphi(x_nx_{n+1}x_{n+2}\dots) \leq \varphi(x_{n+1}x_{n+2}x_{n+3}\dots)$ for every $n \geq 0$. This can be easily proved by induction on n . Let us start with the induction base. By the condition **(b)**, which is already established for φ , we have $\varphi(x_0x_1x_2\dots) \leq \max\{\varphi(x_0^\omega), \varphi(x_1x_2x_3\dots)\}$. Since, $\varphi(x_0x_1x_2\dots) > \varphi(x_0^\omega)$, we have $\varphi(x_0x_1x_2\dots) \leq \varphi(x_1x_2x_3\dots)$.

Let us now perform the induction step. Assume that it is already proved that $\varphi(x_nx_{n+1}x_{n+2}\dots) \leq \varphi(x_{n+1}x_{n+2}x_{n+3}\dots)$ for all $n \leq N$. In particular, this means that

$\varphi(x_{N+1}x_{N+2}x_{N+3}\dots) \geq \varphi(x_0x_1x_2\dots) \geq \varphi(x_{N+1}^\omega)$. Then, by the same argument as in the induction base, we get $\varphi(x_{N+1}x_{N+2}x_{N+3}\dots) \leq \varphi(x_{N+2}x_{N+3}x_{N+4}\dots)$.

We will now prove that $\varphi(x_0x_2\dots x_{2n}x_{2n+1}x_{2n+2}\dots) \geq \varphi(x_0x_1x_2\dots)$ for every $n \geq 0$. For $n = 0$, the left-hand side and the right-hand side coincide. Then we show that $\varphi(x_0x_2\dots x_{2n}x_{2n+1}x_{2n+2}\dots) \leq \varphi(x_0x_2\dots x_{2n+2}x_{2n+3}x_{2n+4}\dots)$ for every $n \geq 0$. Due to the condition **(b)**, we have $\varphi(x_{2n+1}x_{2n+2}x_{2n+3}\dots) \leq \max\{\varphi(x_{2n+1}^\omega), \varphi(x_{2n+2}x_{2n+3}x_{2n+4}\dots)\}$. On the other hand, as we have shown, $\varphi(x_{2n+1}x_{2n+2}x_{2n+3}\dots) \geq \varphi(x_0x_1x_2\dots) > \varphi(x_{2n+1}^\omega)$. Hence $\varphi(x_{2n+1}x_{2n+2}x_{2n+3}\dots) \leq \varphi(x_{2n+2}x_{2n+3}x_{2n+4}\dots)$. It remains to apply **(a)** by appending $x_0x_2\dots x_{2n}$ to both sides.

Thus, we have established that $\varphi(x_0x_2\dots x_{2n}x_{2n+1}x_{2n+2}\dots) \geq \varphi(x_0x_1x_2\dots)$ for every $n \geq 0$. By continuity of φ , the left-hand side of this inequality converges to $\varphi(x_0x_2x_4\dots)$ as $n \rightarrow \infty$. Hence, we get that $\varphi(x_0x_2x_4\dots) \geq \varphi(x_0x_1x_2\dots)$, as required. \square

Thus, to establish that some continuous payoff is positionally determined, it is enough to demonstrate that this payoff satisfies the condition **(a)** of Proposition 4.1. Let us now reformulate this condition using the following definition.

Definition 4.3. Let A be a finite set. A payoff $\varphi: A^\omega \rightarrow \mathbb{R}$ is called **shift-deterministic** if for all $a \in A, \beta, \gamma \in A^\omega$ we have $\varphi(\beta) = \varphi(\gamma) \implies \varphi(a\beta) = \varphi(a\gamma)$.

Claim 4.4. Let A be a finite set. A payoff $\varphi: A^\omega \rightarrow \mathbb{R}$ satisfies the condition **(a)** of Proposition 4.1 if and only if φ is prefix-monotone and shift-deterministic.

Proof. Assume first that φ satisfies the condition **(a)** of Proposition 4.1. It is shift-deterministic, because

$$\begin{aligned} \varphi(\beta) = \varphi(\gamma) &\implies \varphi(\beta) \leq \varphi(\gamma) \wedge \varphi(\gamma) \leq \varphi(\beta) \\ &\implies \varphi(a\beta) \leq \varphi(a\gamma) \wedge \varphi(a\gamma) \leq \varphi(a\beta) \implies \varphi(a\beta) = \varphi(a\gamma) \end{aligned}$$

for every $a \in A, \beta, \gamma \in A^\omega$. In turn, assume for contradiction that φ is not prefix-monotone. Then $\varphi(u\beta) > \varphi(u\gamma)$ and $\varphi(v\beta) < \varphi(v\gamma)$ for some $u, v \in A^*$ and $\beta, \gamma \in A^\omega$. Due to the contraposition to the condition **(a)** of Proposition 4.1, we have $\varphi(u\beta) > \varphi(u\gamma) \implies \varphi(\beta) > \varphi(\gamma)$ and $\varphi(v\gamma) > \varphi(v\beta) \implies \varphi(\gamma) > \varphi(\beta)$, contradiction.

Now, assume that φ is prefix-monotone and shift-deterministic. Take any $u \in A^*$ and $\alpha, \beta \in A^\omega$ such that $\varphi(\alpha) \leq \varphi(\beta)$. We have to derive that $\varphi(u\alpha) \leq \varphi(u\beta)$. If $\varphi(\alpha) = \varphi(\beta)$, then $\varphi(u\alpha) = \varphi(u\beta)$ because φ is shift-deterministic (we apply the definition of the shift-determinism to letters of u from right to left). If $\varphi(\alpha) < \varphi(\beta)$, then $\varphi(u\alpha) \leq \varphi(u\beta)$ because otherwise φ is not prefix-monotone. \square

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

Proposition 4.5. *Let A be a finite set. Any continuous prefix-monotone shift-deterministic payoff $\varphi: A^\omega \rightarrow \mathbb{R}$ 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 4.6. *Let A be a finite set and let $\varphi: A^\omega \rightarrow \mathbb{R}$ be a continuous prefix-monotone payoff. Then $\varphi = g \circ \psi$ for some continuous prefix-monotone shift-deterministic payoff $\psi: A^\omega \rightarrow \mathbb{R}$ and for some continuous² non-decreasing function $g: \psi(A^\omega) \rightarrow \mathbb{R}$.*

²Throughout the paper we call a function $f: S \rightarrow \mathbb{R}$, $S \subseteq \mathbb{R}^n$ continuous if f is continuous with respect to the restriction of the standard topology of \mathbb{R}^n to S .

Due to Corollary 2.3, this proposition means that all continuous prefix-monotone payoffs are positionally determined. In fact, we do not need continuity of g here, but it will be useful later. Thus, once we establish Proposition 3.2, our first proof of Theorem 3.2 will be finished.

Proof of Proposition 4.6. Define a payoff $\psi: A^\omega \rightarrow \mathbb{R}$ as follows:

$$\psi(\gamma) = \sum_{w \in A^*} \left(\frac{1}{|A| + 1} \right)^{|w|} \varphi(w\gamma), \quad \gamma \in A^\omega. \quad (4.1)$$

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 $W > 0$, which means that (4.1) 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 $g: \mathbb{R} \rightarrow \mathbb{R}$.

Why is ψ continuous? We will use Proposition 2.4. 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 words α and β_n have the same prefix of length n . We have to show that $\psi(\beta_n)$ converges to $\psi(\alpha)$ as $n \rightarrow \infty$. By definition:

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

The first series converges to the second one term-wise, because of the continuity of φ . In turn, the first series is bounded by the following absolutely converging series for every n :

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

Hence, the term-wise convergence implies that the sum of the first series converges to the sum of the second one.

Why is ψ prefix-monotone? Take any $\beta, \gamma \in A^\omega$. We have to show that either $\psi(u\beta) \geq \psi(u\gamma)$ for all $u \in A^*$ or $\psi(u\beta) \leq \psi(u\gamma)$ for all $u \in A^*$.

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

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

consists of non-negative terms. Hence $\psi(u\beta) \geq \psi(u\gamma)$ 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

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

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. That is, we have $\varphi(w\beta) - \varphi(w\gamma) = 0$ for every $w \in A^*$. Therefore,

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

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

$$\varphi(\alpha) > \varphi(\beta) \implies \psi(\alpha) > \psi(\beta) \text{ for all } \alpha, \beta \in A^\omega. \quad (4.2)$$

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,

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

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 (4.2) implies that $\varphi = g \circ \psi$ for some non-decreasing $g: \psi(A^\omega) \rightarrow \mathbb{R}$. 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 (4.2) we also have $\varphi(\alpha) = \varphi(\beta)$, 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 to (4.2), we get that $g(x) = \varphi(\gamma_x) \leq \varphi(\gamma_y) = g(y)$, as required.

Finally, we show that any $g: \psi(A^\omega) \rightarrow \mathbb{R}$ such that $\varphi = g \circ \psi$ 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

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

On the left here we have $x - y$, and on the right we have $\varphi(\alpha) - \varphi(\beta) = g \circ \psi(\alpha) - g \circ \psi(\beta) = g(x) - g(y)$. That is, we get $x - y \geq g(x) - g(y) \geq 0$, and it implies that $|g(x) - g(y)| \leq |x - y|$. \square

5. FIXED POINT ARGUMENT

Here we present a way of establishing positional determinacy of continuous prefix-monotone shift-deterministic payoffs (Proposition 4.5) via a fixed point argument. Together with Proposition 4.6, this constitutes our second proof of Theorem 3.2.

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

Remark 5.1. Sometimes, when φ is clear from the context, we will simply write $\mathbf{s}[a]$ instead of $\mathbf{s}[a, \varphi]$.

Claim 5.2. A shift-deterministic payoff $\varphi: A^\omega \rightarrow \mathbb{R}$ is prefix-monotone if and only if $\mathbf{s}[a, \varphi]$ is non-decreasing for every $a \in A$.

Proof. A statement that $\mathbf{s}[a, \varphi]$ is non-decreasing for every $a \in A$ is equivalent to the condition **(a)** of Proposition 4.1. In turn, by Claim 4.4, this condition is equivalent to a statement that φ is prefix-monotone and shift-deterministic. \square

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

Definition 5.3. Let A be a finite set, $\varphi: A^\omega \rightarrow \mathbb{R}$ be a shift-deterministic payoff and $G = \langle V, V_{\text{Max}}, V_{\text{Min}}, E \rangle$ 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 :

$$\mathbf{x}_u = \max_{e \in E, \text{source}(e)=u} \mathbf{s}[\text{lab}(e), \varphi](\mathbf{x}_{\text{target}(e)}), \quad \text{for } u \in V_{\text{Max}}, \quad (5.1)$$

$$\mathbf{x}_u = \min_{e \in E, \text{source}(e)=u} \mathbf{s}[\text{lab}(e), \varphi](\mathbf{x}_{\text{target}(e)}), \quad \text{for } u \in V_{\text{Min}}. \quad (5.2)$$

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

Proposition 5.4. *For any finite set A , for any continuous prefix-monotone shift-deterministic payoff $\varphi: A^\omega \rightarrow \mathbb{R}$ and for any A -labeled game graph G there exists a solution to Bellman's equations for φ in G .*

This proposition requires some additional work. We first discuss why does it imply that all continuous prefix-monotone shift-deterministic payoffs are positionally determined. Assume that we are given a solution \mathbf{x} to (5.1–5.2). 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. Here an edge e is called \mathbf{x} -tight if $\mathbf{x}_{\text{source}(e)} = \mathbf{s}[a, \varphi](\mathbf{x}_{\text{target}(e)})$. Note that each node must contain an out-going \mathbf{x} -tight edge (this will be any edge on which the maximum/minimum in (5.1–5.2) 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 5.5. *If A is a finite set, $\varphi: A^\omega \rightarrow \mathbb{R}$ is a continuous prefix-monotone shift-deterministic payoff, and $\mathbf{x} \in \varphi(A^\omega)^V$ is a solution to (5.1–5.2) for an A -labeled game graph $G = \langle V, V_{\text{Max}}, V_{\text{Min}}, E \rangle$, then the following holds. Let σ^* be a positional strategy of Max and τ^* be a positional strategy of Min such that $\sigma^*(V_{\text{Max}})$ and $\tau^*(V_{\text{Min}})$ consist only of \mathbf{x} -tight edges. Then (σ^*, τ^*) is an equilibrium in G .*

Proof. For brevity, we will omit φ in the notation $\mathbf{s}[a, \varphi]$. We will also use a notation

$$\mathbf{s}[a_1 a_2 \dots a_n] = \mathbf{s}[a_1] \circ \mathbf{s}[a_2] \circ \dots \circ \mathbf{s}[a_n]$$

for $n \in \mathbb{N}, a_1 a_2 \dots a_n \in A^n$. In particular, $\mathbf{s}[\text{empty string}]$ will denote the identity function.

It is enough to show that

- **(a)** for any $v \in V$ and for any $\mathcal{P} \in \text{Cons}(v, \sigma^*)$ we have

$$\varphi \circ \text{lab}(\mathcal{P}) \geq \mathbf{x}_v^*.$$

- **(b)** for any $v \in V$ and for any $\mathcal{P} \in \text{Cons}(v, \tau^*)$ we have

$$\varphi \circ \text{lab}(\mathcal{P}) \leq \mathbf{x}_v^*.$$

Indeed, from these two inequalities we obtain that $\mathbf{Val}[\sigma^*](v) \geq \mathbf{x}_v^* \geq \mathbf{Val}[\tau^*](v)$ for every $v \in V$. But on the other hand, $\mathbf{Val}[\sigma^*](v) \leq \varphi \circ \text{lab}(\mathcal{P}_v^{\sigma^*, \tau^*}) \leq \mathbf{Val}[\tau^*](v)$. That is, we get that $\mathbf{Val}[\sigma^*](v) = \mathbf{Val}[\tau^*](v)$ for every $v \in V$, and this by definition means that (σ^*, τ^*) is an equilibrium.

We only show the first item, the second one can be proved similarly. Let e_n be the n th edge of \mathcal{P} for $n \geq 1$. Define $v_n = \text{target}(e_1 e_2 \dots e_n) \in V$ and $T_n = \mathbf{s}[\text{lab}(e_1 \dots e_n)](\mathbf{x}_{v_n}^*) \in \varphi(A^\omega)$. We also set $v_0 = v$ and $T_0 = \mathbf{x}_v^*$. Note that due to the continuity of φ we have that $\lim_{n \rightarrow \infty} T_n = \varphi \circ \text{lab}(\mathcal{P})$. Indeed, $\mathbf{x}_{v_n}^* \in \varphi(A^\omega)$, so there exists $\beta_n \in A^\omega$ with $\mathbf{x}_{v_n}^* = \varphi(\beta_n)$. Hence, $\varphi(\text{lab}(e_1 \dots e_n)\beta_n) = \mathbf{s}[\text{lab}(e_1 \dots e_n)](\varphi(\beta_n)) = \mathbf{s}[\text{lab}(e_1 \dots e_n)](\mathbf{x}_{v_n}^*) = T_n$. On the other hand, the first n letters of $\text{lab}(e_1 \dots e_n)\beta_n$ and $\text{lab}(\mathcal{P})$ coincide. Hence, $\varphi(\text{lab}(e_1 \dots e_n)\beta_n) = T_n$ converges to $\varphi \circ \text{lab}(\mathcal{P})$ as $n \rightarrow \infty$, as required.

So **(a)** is equivalent to a statement that $\lim_{n \rightarrow \infty} T_n \geq T_0$. To show this statement, we demonstrate that $T_{n+1} \geq T_n$ for every n . Indeed, assume first that $v_n \in V_{\text{Max}}$. Then, since \mathcal{P} is consistent with σ^* , we have $e_{n+1} = \sigma^*(v_n)$. In particular, e_{n+1} is \mathbf{x} -tight, by the conditions of the lemma. This gives us that $\mathbf{s}[\text{lab}(e_{n+1})](\mathbf{x}_{v_{n+1}}^*) = \mathbf{x}_{v_n}^*$. After applying the function $\mathbf{s}[\text{lab}(e_1 e_2 \dots e_n)]$ to this equality, we obtain $T_{n+1} = T_n$.

Now, if $v_n \in V_{\text{Min}}$, then $\mathbf{s}[\text{lab}(e_{n+1})](\mathbf{x}_{v_{n+1}}^*) \geq \mathbf{x}_{v_n}^*$ by (5.2). The function $\mathbf{s}[\text{lab}(e_1 e_2 \dots e_n)]$ is composed of non-decreasing functions due to Claim 5.2. Hence, after applying this function to the left-hand and the right-hand sides of the inequality $\mathbf{s}[\text{lab}(e_{n+1})](\mathbf{x}_{v_{n+1}}^*) \geq \mathbf{x}_{v_n}^*$, we obtain $T_{n+1} \geq T_n$. \square

We now proceed to details of our proof of Proposition 5.4. Consider a function $T: \varphi(A^\omega)^V \rightarrow \varphi(A^\omega)^V$, mapping $\mathbf{x} \in \varphi(A^\omega)^V$ to the vector of the right-hand sides of (5.1–5.2). We should argue that T has a fixed point. For that, we will construct a continuous metric $D: \varphi(A^\omega)^V \times \varphi(A^\omega)^V \rightarrow [0, +\infty)$ 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 $d: \varphi(A^\omega) \times \varphi(A^\omega) \rightarrow [0, +\infty)$ such that all functions $\mathbf{s}[a, \varphi], a \in A$ are 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. The main technical challenge is to prove the existence of d . We do so via the following general fact about compositions of continuous functions.

Theorem 5.6. *Let $K \subseteq \mathbb{R}$ be a compact set, $m \geq 1$ be a natural number and $f_1, \dots, f_m: K \rightarrow K$ be m continuous functions. Then the following two conditions are equivalent:*

- **(a)** for every $a_1 a_2 a_3 \dots \in \{1, 2, \dots, m\}^\omega$ we have $\lim_{n \rightarrow \infty} \text{diam}(f_{a_1} \circ f_{a_2} \circ \dots \circ f_{a_n}(K)) = 0$ (by $\text{diam}(S)$ for $S \subseteq \mathbb{R}$ we mean $\sup_{x, y \in S} |x - y|$);
- **(b)** there exists a continuous metric $d: K \times K \rightarrow [0, +\infty)$ such that f_1, f_2, \dots, f_m are all d -contracting (a function $h: K \rightarrow K$ 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, \dots, f_m are non-decreasing, then these two conditions are equivalent to the following condition:

- **(c)** there exists a continuous metric $d: K \times K \rightarrow [0, +\infty)$ such that, first, f_1, f_2, \dots, f_m are all d -contracting, and second, for all $x, y, s, t \in K$ we have $x \leq s \leq t \leq y \implies d(s, t) \leq d(x, y)$.

Proof. See Appendix B. □

To derive Proposition 5.4 from this theorem, we first show that it is applicable to functions $\mathbf{s}[a, \varphi], a \in A$ for continuous shift-deterministic φ .

Proposition 5.7. *Let A be a finite set and $\varphi: A^\omega \rightarrow \mathbb{R}$ be a continuous shift-deterministic payoff. Then the functions $\mathbf{s}[a, \varphi], a \in A$ are continuous and satisfy the condition **(a)** of Theorem 5.6 for $K = \varphi(A^\omega)$.*

Proof. We use the same abbreviations with respect to the notation $\mathbf{s}[a, \varphi]$ as in the proof of Lemma 5.5.

Let us first demonstrate that $\mathbf{s}[a]$ is continuous for every $a \in A$. Consider any sequence $\{x_n\}_{n \in \mathbb{N}}$ of elements of $\varphi(A^\omega)$, converging to $x \in \varphi(A^\omega)$. We shall show that $\mathbf{s}[a](x_n)$ converges to $\mathbf{s}[a](x)$. Assume for contradiction that for some $\varepsilon > 0$ there are infinitely many n with $|\mathbf{s}[a](x_n) - \mathbf{s}[a](x)| > \varepsilon$. By restricting our sequence to such n , we may assume that all of them satisfy this inequality. Let $\beta_n \in A^\omega$ be such that $x_n = \varphi(\beta_n)$. Due to the compactness of A^ω , there exists $\beta \in A^\omega$ such that any open set $\mathcal{S} \subseteq A^\omega$ containing β also contains β_n for infinitely many n . Indeed, otherwise any point of A^ω is contained in an open set which covers only finitely many elements of the sequence $\{\beta_n\}_{n \in \mathbb{N}}$. A collection of such open sets would be an open cover of A^ω without a finite subcover.

We derive from this that for every $k \in \mathbb{N}$ there exists $n_k \geq k$ such that the first k letters of β_{n_k} and β coincide. Indeed, consider a word $u \in A^k$, consisting of the first k letters of β . An open set $\mathcal{S} = uA^\omega$ contains β . Hence, there are infinitely many n such that $\beta_n \in uA^\omega$, or, equivalently, such that β_n starts with u . In particular, there exists such n which is at least as large as k .

Due to continuity of φ , we have that $\lim_{k \rightarrow \infty} \varphi(\beta_{n_k}) = \varphi(\beta)$. On the other hand, $\lim_{k \rightarrow \infty} \varphi(\beta_{n_k}) = \lim_{k \rightarrow \infty} x_{n_k} = x$. Hence, $\varphi(\beta) = x$. Using the continuity of φ again, we get $\lim_{k \rightarrow \infty} \varphi(a\beta_{n_k}) = \varphi(a\beta)$. But $\varphi(a\beta_{n_k}) = \mathbf{s}[a](\varphi(\beta_{n_k})) = \mathbf{s}[a](x_{n_k})$, $\varphi(a\beta) = \mathbf{s}[a](\varphi(\beta)) = \mathbf{s}[a](x)$, and $|\mathbf{s}[a](x_n) - \mathbf{s}[a](x)| > \varepsilon$ for every n , contradiction.

Now, let us show that $\mathbf{s}[a], a \in A$ satisfy item **(a)** of Theorem 5.6 for $K = \varphi(A^\omega)$. By definition of $\mathbf{s}[a]$, we have that

$$\mathbf{s}[a_1 \dots a_n](\varphi(A^\omega)) = \varphi(a_1 a_2 a_3 \dots a_n A^\omega)$$

for every $n \in \mathbb{N}$ and $a_1 a_2 \dots a_n \in A^n$. Thus, it is enough to establish that

$$\lim_{n \rightarrow \infty} \text{diam}(\varphi(a_1 a_2 \dots a_n A^\omega)) = 0$$

for any $a_1 a_2 a_3 \dots \in A^\omega$. This is a simple consequence of the continuity of φ . Indeed, assume for contradiction that for some $a_1 a_2 a_3 \dots \in A^\omega$ we have $\text{diam}(\varphi(a_1 a_2 \dots a_n A^\omega)) > \varepsilon$

for infinitely many n . Then for infinitely many n there exist $\beta_n, \gamma_n \in a_1 a_2 \dots a_n A^\omega$ with $|\varphi(\beta_n) - \varphi(\gamma_n)| \geq \varepsilon$. At the same time, by continuity of φ , both $\varphi(\beta_n)$ and $\varphi(\gamma_n)$ must converge to $\varphi(a_1 a_2 a_3 \dots)$, contradiction. \square

We finally derive Proposition 5.4 from Theorem 5.6 and Proposition 5.7. This will finish our second proof of the fact that all continuous prefix-monotone payoffs are positionally determined.

Proof of Proposition 5.4. We use the same abbreviations with respect to the notation $\mathbf{s}[a, \varphi]$ as in the proof of Lemma 5.5.

Define a mapping $T: K^V \rightarrow K^V$, where $K = \varphi(A^\omega)$, as follows:

$$T(\mathbf{x})_u = \max_{e \in E, \text{source}(e)=u} \mathbf{s}[\text{lab}(e)](\mathbf{x}_{\text{target}(e)}), \quad \text{for } u \in V_{\text{Max}}, \quad (5.3)$$

$$T(\mathbf{x})_u = \min_{e \in E, \text{source}(e)=u} \mathbf{s}[\text{lab}(e)](\mathbf{x}_{\text{target}(e)}), \quad \text{for } u \in V_{\text{Min}}. \quad (5.4)$$

Recall that K is a compact set (because A^ω is compact and φ is continuous). It is enough to show that T has a fixed point. By Proposition 5.7, the functions $\mathbf{s}[a]$, $a \in A$ are continuous (which means that T is also continuous) and satisfy the item **(a)** of Theorem 5.6. By Claim 5.2, the functions $\mathbf{s}[a]$, $a \in A$ are non-decreasing. Hence, these functions satisfy the item **(c)** of Theorem 5.6. That is, there exists a continuous metric $d: K \times K \rightarrow [0, +\infty)$ such that, first, the function $\mathbf{s}[a]$ is d -contracting for every $a \in A$, and second, for every $x, s, t, y \in K$ we have $x \leq s \leq t \leq y \implies d(s, t) \leq d(x, y)$.

Define a metric $D: K^V \times K^V \rightarrow [0, +\infty)$ as follows:

$$D(\mathbf{x}, \mathbf{y}) = \max_{u \in V} d(\mathbf{x}_u, \mathbf{y}_u).$$

It is enough to show $D(T(\mathbf{x}), T(\mathbf{y})) < D(\mathbf{x}, \mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in K^V$, $\mathbf{x} \neq \mathbf{y}$. Indeed, assume that this inequality is already established. Consider a point $\mathbf{x}^* \in K^V$ minimizing $D(\mathbf{x}, T(\mathbf{x}))$. Such \mathbf{x}^* exists because $D(\mathbf{x}, T(\mathbf{x}))$ is continuous and $K^V \times K^V$ is a compact set. If $\mathbf{x}^* \neq T(\mathbf{x}^*)$, then $D(T(\mathbf{x}^*), T \circ T(\mathbf{x}^*)) < D(\mathbf{x}^*, T(\mathbf{x}^*))$, contradiction.

Now, take any $\mathbf{x}, \mathbf{y} \in K^V$, $\mathbf{x} \neq \mathbf{y}$. Let $u \in V$ be such that $D(T(\mathbf{x}), T(\mathbf{y})) = d(T(\mathbf{x})_u, T(\mathbf{y})_u)$. Assume w.l.o.g. that $u \in V_{\text{Max}}$. Also, up to swapping \mathbf{x} and \mathbf{y} , we may assume that $T(\mathbf{x})_u \leq T(\mathbf{y})_u$. Let e be an edge on which the maximum in (5.3) is attained for \mathbf{y} . That is, the source of e is u and $T(\mathbf{y})_u = \mathbf{s}[\text{lab}(e)](\mathbf{y}_w)$, where $w = \text{target}(e)$. On the other hands, by (5.3) applied to \mathbf{x} , we get $\mathbf{s}[\text{lab}(e)](\mathbf{x}_w) \leq T(\mathbf{x})_u$. Overall,

$$\mathbf{s}[\text{lab}(e)](\mathbf{x}_w) \leq T(\mathbf{x})_u \leq T(\mathbf{y})_u = \mathbf{s}[\text{lab}(e)](\mathbf{y}_w).$$

Since for any $x, s, t, y \in K$ it holds that $x \leq s \leq t \leq y \implies d(s, t) \leq d(x, y)$, we get:

$$d(T(\mathbf{x})_u, T(\mathbf{y})_u) \leq d(\mathbf{s}[\text{lab}(e)](\mathbf{x}_w), \mathbf{s}[\text{lab}(e)](\mathbf{y}_w)).$$

If $\mathbf{x}_w = \mathbf{y}_w$, then $0 = d(T(\mathbf{x})_u, T(\mathbf{y})_u) = D(T(\mathbf{x}), T(\mathbf{y})) < D(\mathbf{x}, \mathbf{y})$, because $\mathbf{x} \neq \mathbf{y}$. Now, if $\mathbf{x}_w \neq \mathbf{y}_w$, then due to the fact that the function $\mathbf{s}[\text{lab}(e)]$ is d -contracting, we have:

$$d(\mathbf{s}[\text{lab}(e)](\mathbf{x}_w), \mathbf{s}[\text{lab}(e)](\mathbf{y}_w)) < d(\mathbf{x}_w, \mathbf{y}_w) \leq D(\mathbf{x}, \mathbf{y}),$$

which gives us $D(T(\mathbf{x}), T(\mathbf{y})) = d(T(\mathbf{x})_u, T(\mathbf{y})_u) \leq d(\mathbf{s}[\text{lab}(e)](\mathbf{x}_w), \mathbf{s}[\text{lab}(e)](\mathbf{y}_w)) < D(\mathbf{x}, \mathbf{y})$, as required. \square

6. THE STRUCTURE OF CONTINUOUS POSITIONAL PAYOFFS

In this section, we give an explicit description of the set of continuous positionally determined payoffs, see Theorem 6.3 below. Then, in Proposition 6.5, we use our description to give an alternative definition of the class of multi-discounted payoffs. Finally, in Proposition 6.7, we give an example of a continuous positionally determined payoff which, in a quite strong sense, does not “reduce” to multi-discounted payoffs.

We start with some terminology. Let $K \subseteq \mathbb{R}$ be a compact set. We call a family of m continuous functions $f_1, \dots, f_m: K \rightarrow K$ a *contracting base* if there exists a continuous metric $d: K \times K \rightarrow [0, +\infty)$ such that f_i is d -contracting for every $i \in \{1, \dots, m\}$ (in other words, if f_1, \dots, f_m satisfy the item **(b)** of Theorem 5.6). If f_1, \dots, f_m are non-decreasing, then we call this family of functions a non-decreasing contracting base.

Claim 6.1. Let $K \subseteq \mathbb{R}$ be a compact set and $f_1, \dots, f_m: K \rightarrow K$ be m continuous functions forming a contracting base. Then for any $a_1 a_2 a_3 \dots \in \{1, 2, \dots, m\}^\omega$ we have

$$\left| \bigcap_{n=1}^{\infty} f_{a_1} \circ f_{a_2} \circ \dots \circ f_{a_n}(K) \right| = 1.$$

Moreover, for any $x \in K$, the quantity $f_{a_1} \circ f_{a_2} \circ \dots \circ f_{a_n}(x)$ converges to the unique element of this intersection as $n \rightarrow \infty$.

Proof. This intersection is non-empty due to Cantor’s intersection theorem. To show that this intersection contains just one point, observe that f_1, \dots, f_m satisfy the item **(b)** of Theorem 5.6 by definition. Hence, they also satisfy the item **(a)** of this theorem. This means that the diameter of this intersection is 0.

As for the second claim, note that the distance between the unique element of our intersection and the point $f_{a_1} \circ f_{a_2} \circ \dots \circ f_{a_n}(x)$ is at most $\text{diam} f_{a_1} \circ f_{a_2} \circ \dots \circ f_{a_n}(K)$, because both these points belong to $f_{a_1} \circ f_{a_2} \circ \dots \circ f_{a_n}(K)$. It remains to refer to the item **(a)** of Theorem 5.6 once again. \square

This claim means that any contracting base $f_1, \dots, f_m: K \rightarrow K$ induces a payoff $\psi[f_1, \dots, f_m]: \{1, \dots, m\}^\omega \rightarrow K \subseteq \mathbb{R}$, defined by

$$\{\psi[f_1, \dots, f_m](a_1 a_2 a_3 \dots)\} = \bigcap_{n=1}^{\infty} f_{a_1} \circ f_{a_2} \circ \dots \circ f_{a_n}(K),$$

for every $a_1 a_2 a_3 \dots \in \{1, 2, \dots, m\}^\omega$. By the second part of Claim 6.1, we have:

$$\psi[f_1, \dots, f_m](a_1 a_2 a_3 \dots) = \lim_{n \rightarrow \infty} f_{a_1} \circ \dots \circ f_{a_n}(x)$$

for every $x \in K$.

Claim 6.2. Let $K \subseteq \mathbb{R}$ be a compact set and $f_1, \dots, f_m: K \rightarrow K$ be m continuous functions forming a contracting base. Then the payoff $\psi = \psi[f_1, \dots, f_m]$, induced by f_1, \dots, f_m , is continuous and shift-deterministic. Moreover, for all $i \in \{1, 2, \dots, m\}$ and $\alpha = a_1 a_2 a_3 \dots \in \{1, 2, \dots, m\}^\omega$ we have:

$$\psi(i\alpha) = f_i(\psi(\alpha)). \tag{6.1}$$

Proof. Let us first establish (6.1). Take any $x \in K$. By Claim 6.1, we have $\psi(i\alpha) = \lim_{n \rightarrow \infty} f_i \circ f_{a_1} \circ \dots \circ f_{a_n}(x)$. By continuity of f_i , we get $\psi(i\alpha) = f_i(\lim_{n \rightarrow \infty} f_{a_1} \circ \dots \circ f_{a_n}(x)) = f_i(\psi(\alpha))$.

This immediately implies that ψ is shift-deterministic. To show that ψ is continuous, we use Proposition 2.4. Take any $\alpha = a_1 a_2 a_3 \dots \in \{1, 2, \dots, m\}^\omega$ and any infinite sequence $\{\beta_n\}_{n \geq 1}$ of elements of $\{1, 2, \dots, m\}^\omega$ such that α and β_n have the same prefixes of length n , for every $n \geq 1$. We have to show that $\psi(\beta_n) \rightarrow \psi(\alpha)$ as $n \rightarrow \infty$. By (6.1), both $\psi(\beta_n)$ and $\psi(\alpha)$ belong to the set $f_{a_1} \circ f_{a_2} \circ \dots \circ f_{a_n}(K)$. Hence, the difference between $\psi(\beta_n)$ and $\psi(\alpha)$ does not exceed the diameter of this set. But by the item **(a)** of Theorem 5.6, the diameter of this set converges to 0 as $n \rightarrow \infty$. \square

Theorem 6.3. *Let $m \geq 1$ be a natural number. Set³ $A = \{1, 2, \dots, m\}$. Then the set of continuous positionally determined payoffs from A^ω to \mathbb{R} coincides with the set of $\varphi: A^\omega \rightarrow \mathbb{R}$ 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 $d: K \times K \rightarrow [0, +\infty)$.
- **Step 3.** Take m continuous non-decreasing d -contracting functions $f_1, f_2, \dots, f_m: K \rightarrow K$. They will form a non-decreasing contracting base.
- **Step 4.** Consider the payoff $\psi = \psi[f_1, \dots, f_m]$ induced by f_1, \dots, f_m .
- **Step 5.** Choose a continuous non-decreasing function $g: \psi(A^\omega) \rightarrow \mathbb{R}$ and set $\varphi = g \circ \psi$.

Proof. Assume first that $\varphi: A^\omega \rightarrow \mathbb{R}$ is continuous and positionally determined. Then φ is prefix-monotone by Theorem 3.2. By Proposition 4.6, there is a continuous prefix-monotone shift-deterministic payoff $\psi: A^\omega \rightarrow \mathbb{R}$ and a continuous non-decreasing function $g: \psi(A^\omega) \rightarrow \mathbb{R}$ such that $\varphi = g \circ \psi$. Set $K = \psi(A^\omega)$. Note that K is compact due to the continuity of ψ . Define $f_i = \mathbf{s}[i, \psi]$. By Claim 5.2, the functions f_1, f_2, \dots, f_m are non-decreasing. By Proposition 5.7, the functions f_1, f_2, \dots, f_m are continuous and satisfy the item **(a)** of Theorem 5.6. Hence, they form a non-decreasing contracting base with respect to some continuous metric $d: K \times K \rightarrow [0, +\infty)$. It remains to show that ψ coincides with the payoff induced by f_1, \dots, f_m . For that, take any $x \in K = \psi(A^\omega)$. By the second part of Claim (6.1), it is sufficient to show that $\psi(a_1 a_2 a_3 \dots) = \lim_{n \rightarrow \infty} f_{a_1} \circ f_{a_2} \circ \dots \circ f_{a_n}(x)$ for every $a_1 a_2 a_3 \dots \in A^\omega$. There exists $\beta \in A^\omega$ such that $x = \psi(\beta)$. Observe that $f_{a_1} \circ f_{a_2} \circ \dots \circ f_{a_n}(x) = \mathbf{s}[a_1, \psi] \circ \mathbf{s}[a_2, \psi] \circ \dots \circ \mathbf{s}[a_n, \psi](\psi(\beta)) = \psi(a_1 a_2 \dots a_n \beta)$. This quantity converges to $\psi(a_1 a_2 a_3 \dots)$ as $n \rightarrow \infty$ due to the continuity of ψ .

In turn, assume that φ was obtained in these 5 steps. By Theorem 3.2, we only have to show that φ is continuous and prefix-monotone. First, by Claim 6.2, we have that ψ is continuous. Since $\varphi = g \circ \psi$ and g is continuous, we have that φ is continuous as well. In turn, since f_1, \dots, f_m are non-decreasing, from (6.1) we get that ψ is prefix-monotone. This easily implies that φ is also prefix-monotone. Indeed, if φ is not prefix-monotone, then $g \circ \psi(u\beta) > g \circ \psi(u\gamma)$ and $g \circ \psi(v\beta) > g \circ \psi(v\gamma)$ for some $u, v \in A^*$ and $\beta, \gamma \in A^\omega$. Since g is non-decreasing, this implies that $\psi(u\beta) > \psi(u\gamma)$ and $\psi(v\beta) < \psi(v\gamma)$, contradiction with the prefix-monotonicity of ψ . \square

Remark 6.4. Recall that we did not use the continuity of g from Proposition 4.6 in the inductive argument, but we use it in the proof of Theorem 6.3.

Next, we characterize the class of multi-discounted payoffs, using the language of Theorem 6.3.

³We assume that edge labels are natural numbers, for the notation simplicity.

Proposition 6.5. *Let $m \geq 1$ be a natural number and set $A = \{1, 2, \dots, m\}$. Then the set of multi-discounted payoffs from A^ω to \mathbb{R} coincides with the set of $\psi: A^\omega \rightarrow \mathbb{R}$ that can be obtained as in Theorem 6.3 with the following additional requirements:*

- $K = [-W, W]$ for some $W > 0$;
- d is a standard metric $d(x, y) = |x - y|$;
- f_1, \dots, f_m are affine functions with the slope from $[0, 1)$. That is, for each $a \in \{1, \dots, m\}$ there exists $\lambda(a) \in [0, 1)$ and $w(a) \in \mathbb{R}$ such that $f_a(x) = \lambda(a)x + w(a)$ (observe that we must have $\lambda(a)W + w(a) \leq W$ and $\lambda(a)(-W) + w(a) \geq -W$, because $f_a: [-W, W] \rightarrow [-W, W]$).
- g is the identity function.

Proof. First, assume that $\psi: A^\omega \rightarrow \mathbb{R}$ is a multi-discounted payoff. Take $\lambda: A \rightarrow [0, 1)$ and $w: A \rightarrow \mathbb{R}$ such that

$$\psi(a_1 a_2 a_3 \dots) = \sum_{n=1}^{\infty} \lambda(a_1) \cdot \dots \cdot \lambda(a_{n-1}) \cdot w(a_n)$$

for all $a_1 a_2 a_3 \dots \in A^\omega$. Obviously, there exists $W > 0$ such that $\lambda(a)W + w(a) \leq W$ and $\lambda(a)(-W) + w(a) \geq -W$ for all $a \in A$. Set $K = [-W, W]$ and $f_a(x) = \lambda(a)x + w(a)$ for every $a \in A$. Obviously, $f_1, \dots, f_m: [-W, W] \rightarrow [-W, W]$ form a non-decreasing contracting base with respect to a standard metric $d(x, y) = |x - y|$. It remains to show that ψ coincides with the payoff induced by f_1, \dots, f_m . Indeed,

$$\begin{aligned} \psi(a_1 a_2 a_3 \dots) &= \sum_{n=1}^{\infty} \lambda(a_1) \cdot \dots \cdot \lambda(a_{n-1}) \cdot w(a_n) \\ &= \lim_{n \rightarrow \infty} [w(a_1) + \lambda(a_1) \cdot w(a_2) + \dots + \lambda(a_1) \cdot \dots \cdot \lambda(a_{n-1}) w(a_n)] \\ &= \lim_{n \rightarrow \infty} [w(a_1) + \lambda(a_1)(w(a_2) + \dots + (w(a_{n-1}) + \lambda(a_{n-1})w(a_n)))] \\ &= \lim_{n \rightarrow \infty} [f_{a_1} \circ f_{a_2} \circ \dots \circ f_{a_n}(0)]. \end{aligned}$$

This computation also establishes the opposite direction of the proposition. Indeed, assume that ψ was obtained as in Theorem 6.3 with the requirements of our proposition. Then ψ must be equal to the payoff induced by f_1, \dots, f_m , where f_1, \dots, f_m are affine functions with the slope from $[0, 1)$. By writing down the same chain of equalities as above, we get that ψ is multi-discounted. \square

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

Definition 6.6. Let A be a finite set, $\varphi, \psi: A^\omega \rightarrow \mathbb{R}$ 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 [24] 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 multi-discounted payoffs in this sense.

Proposition 6.7. *There exist a finite set A , a continuous positionally determined payoff $\varphi: A^\omega \rightarrow \mathbb{R}$ and an A -labeled game graph G such that there exists no multi-discounted payoff to which φ reduces inside G .*

Proof. It is sufficient to establish the following lemma.

Lemma 6.8. *There exist a finite set A , a continuous positionally determined payoff $\varphi: A^\omega \rightarrow \mathbb{R}$ and three pairs $(\alpha_1, \beta_1), (\alpha_2, \beta_2), (\alpha_3, \beta_3) \in A^\omega \times A^\omega$ of ultimately periodic infinite words such that:*

- for all $i \in \{1, 2, 3\}$ we have $\varphi(\alpha_i) > \varphi(\beta_i)$
- for every multi-discounted payoff $\psi: A^\omega \rightarrow \mathbb{R}$ there exists $i \in \{1, 2, 3\}$ such that $\psi(\alpha_i) \leq \psi(\beta_i)$.

Indeed, assume that this lemma is proved. Consider a game graph from Figure 2, consisting of three pairs of “lassos”. The only optimal positional strategy of Max there w.r.t. φ is to go to the left from v_1, v_2 and v_3 . On the other hand, any multi-discounted payoff has an optimal positional strategy which for some $i \in \{1, 2, 3\}$ goes to the right from v_i . Hence, there is no multi-discounted payoff to which φ positionally reduces inside the game graph from Figure 2.




Figure 2: All nodes are owned by Max. For every $i = 1, 2, 3$, the node v_i has two lassos \mathcal{L}_i and \mathcal{R}_i starting at it, one going to the left, and the other going to the right. We label their edges in such a way that $\text{lab}(\mathcal{L}_i) = \alpha_i$ and $\text{lab}(\mathcal{R}_i) = \beta_i$. This is possible because $\alpha_1, \alpha_2, \alpha_3$ and $\beta_1, \beta_2, \beta_3$ are ultimately periodic.

To show Lemma 6.8, we observe that following property of the multi-discounted payoffs.

Claim 6.9. Let A be a finite set and $\psi: A^\omega \rightarrow \mathbb{R}$ be a multi-discounted payoff. Then there are no $a, b \in A, \gamma \in A^\omega$ such that

$$\begin{aligned} \psi(a\gamma) &> \psi(b\gamma), \\ \psi(aa\gamma) &< \psi(bb\gamma), \\ \psi(aaa\gamma) &> \psi(bbb\gamma). \end{aligned}$$

Proof. Assume for contradiction that such a, b, γ exist. Take $\lambda: A \rightarrow [0, 1)$ and $w: A \rightarrow [0, 1)$ defining ψ as in (3.1). Set $\lambda = \lambda(a)$, $\mu = \lambda(b)$, $u = w(a)$, $v = w(b)$ and $x = \psi(\gamma)$. Then

$\lambda, \mu \in [0, 1)$ and

$$\lambda x + u > \mu x + v, \quad (6.2)$$

$$\lambda^2 x + (1 + \lambda)u < \mu^2 x + (1 + \mu)v, \quad (6.3)$$

$$\lambda^3 x + (1 + \lambda + \lambda^2)u > \mu^3 x + (1 + \mu + \mu^2)v. \quad (6.4)$$

Multiply (6.2) by $\lambda + \mu + \lambda\mu$, multiply (6.3) by $-(1 + \lambda + \mu)$, multiply (6.4) by 1 and take the sum. This will give us $0 > 0$, contradiction. \square

To finish a proof of Lemma 6.8, we construct a continuous positionally determined payoff $\varphi: \{1, 2, 3\}^\omega \rightarrow \mathbb{R}$ such that:

$$\varphi(13^\omega) > \varphi(23^\omega),$$

$$\varphi(113^\omega) < \varphi(223^\omega),$$

$$\varphi(1113^\omega) > \varphi(2223^\omega).$$

For that, we use Theorem 6.3. Namely, we set $K = [0, 1]$ and $d(x, y) = |x - y|$. Next, we let $f_1 = \frac{x}{2}$, $f_3 = \frac{x}{2} + \frac{1}{2}$. These two functions are clearly d -contracting. Finally, we let $f_2: [0, 1] \rightarrow [0, 1]$ be a piece-wise linear function whose graph has the following break-points:

$$(0, 0), (0.26, 0.11), (0.49, 0.26), (1, 0.49).$$

Observe that its slope is always from $[0, 1)$, so f_2 is also d -contracting. So f_1, f_2, f_3 is a non-decreasing contracting base. Let $\varphi: \{1, 2, 3\}^\omega \rightarrow \mathbb{R}$ be the payoff induced by f_1, f_2, f_3 , that is,

$$\varphi(a_1 a_2 a_3 \dots) = \lim_{n \rightarrow \infty} f_{a_1} \circ f_{a_2} \circ \dots \circ f_{a_n}(1).$$

(Of course, 1 can here be changed to any point from $[0, 1]$, but 1 is the most convenient for computations below.) By Theorem 6.3, we have that φ is a continuous positionally determined payoff. Now, it is easy to see that $\varphi(3^\omega) = 1$ and

$$\varphi(13^\omega) = 0.5 > \varphi(23^\omega) = 0.49,$$

$$\varphi(113^\omega) = 0.25 < \varphi(223^\omega) = 0.26,$$

$$\varphi(1113^\omega) = 0.125 > \varphi(2223^\omega) = 0.11.$$

\square

7. STRATEGY IMPROVEMENT ARGUMENT

Here we establish the existence of a solution to Bellman's equations (Proposition 5.4) via the *strategy improvement*. This will yield our third proof of Theorem 3.2. We start with an observation that the vector of values of a positional strategy always gives a solution to a restriction of Bellman's equations to edges that are consistent with this strategy.

Lemma 7.1. *Let A be a finite set, $\varphi: A^\omega \rightarrow \mathbb{R}$ be a continuous prefix-monotone shift-deterministic payoff and $G = \langle V, V_{\text{Max}}, V_{\text{Min}}, E \rangle$ be an A -labeled game graph. Then for every positional strategy σ of Max in G we have for all $u \in V$:*

$$\mathbf{Val}[\sigma](u) = \min_{e \in E^\sigma, \text{source}(e)=u} \mathbf{s}[\text{lab}(e), \varphi](\mathbf{Val}[\sigma](\text{target}(e))) \text{ for } u \in V_{\text{Min}}.$$

(If $u \in V_{\text{Max}}$, the minimum is over a single edge $e = \sigma(u)$. If $u \in V_{\text{Min}}$, the minimum is over all edges that start at u).

Remark 7.2. Technically, Bellman's equations are over $\mathbf{x} \in \varphi(A^\omega)^V$. So we have to argue that $\mathbf{Val}[\sigma](u) \in \varphi(A^\omega)$ for every $u \in V$. This is because $\mathbf{Val}[\sigma](u)$ is the infimum of some subset of $\varphi(A^\omega)$. In turn, since φ is continuous, we have that $\varphi(A^\omega)$ is compact, and hence is closed.

Proof of Lemma 7.1. For brevity, we will denote $C_u = \text{Cons}(u)$. By definition, $\mathbf{Val}[\sigma](u)$ is the infimum of the image of $\varphi \circ \text{lab}$ on the set C_u . Now, the set C_u is exactly the set of infinite paths that start at u and consist only of edges from E^σ . So we can write:

$$C_u = \bigcup_{\substack{e \in E^\sigma \\ \text{source}(e)=u}} eC_{\text{target}(e)}.$$

The infimum of a union of finitely many sets is the minimum of the infimums of these sets. So we get:

$$\mathbf{Val}[\sigma](u) = \min_{\substack{e \in E^\sigma \\ \text{source}(e)=u}} \inf \varphi \circ \text{lab} (eC_{\text{target}(e)}).$$

It is sufficient to show that:

$$\inf \varphi \circ \text{lab} (eC_{\text{target}(e)}) = \mathbf{s}[\text{lab}(e)](\mathbf{Val}[\sigma](\text{target}(e))). \quad (7.1)$$

For any $a \in A, \mathcal{S} \subseteq A^\omega$, by definition of $\mathbf{s}[a]$, we can write:

$$\varphi(a\mathcal{S}) = \mathbf{s}[a](\varphi(\mathcal{S})).$$

After applying this to $a = \text{lab}(e), \mathcal{S} = \text{lab}(C_{\text{target}(e)})$, we obtain:

$$\varphi \circ \text{lab} (eC_{\text{target}(e)}) = \mathbf{s}[\text{lab}(e), \varphi](\varphi \circ \text{lab} (C_{\text{target}(e)})).$$

Now, since $\mathbf{s}[\text{lab}(e)]$ is non-decreasing (by Claim 5.2) and continuous (by Proposition 5.7), we can interchange \inf and $\mathbf{s}[\text{lab}(e)]$. This gives us:

$$\begin{aligned} \inf \mathbf{s}[\text{lab}(e)](\varphi \circ \text{lab} (C_{\text{target}(e)})) \\ &= \mathbf{s}[\text{lab}(e)](\inf \varphi \circ \text{lab} (C_{\text{target}(e)})) \\ &= \mathbf{s}[\text{lab}(e)](\mathbf{Val}[\sigma](\text{target}(e))). \end{aligned}$$

Hence, (7.1) is proved. \square

Next, take a positional strategy σ of Max. If the vector $\{\mathbf{Val}[\sigma](u)\}_{u \in V}$ happens to be a solution to the Bellman's equations, then we are done. Otherwise by Lemma 7.1 there must exist an edge $e \in E$ with $\text{source}(e) \in V_{\text{Max}}$ such that $\mathbf{Val}[\sigma](\text{source}(e)) < \mathbf{s}[\text{lab}(e), \varphi](\mathbf{Val}[\sigma](\text{target}(e)))$. 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 .

Lemma 7.3. *Let A be a finite set, $\varphi: A^\omega \rightarrow \mathbb{R}$ be a continuous prefix-monotone shift-deterministic payoff and $G = \langle V, V_{\text{Max}}, V_{\text{Min}}, E \rangle$ be an A -labeled game graph. Next, let σ be a positional strategy of Max in G . Assume that the vector $\mathbf{Val}[\sigma] = \{\mathbf{Val}[\sigma](u)\}_{u \in V}$ does not satisfy (5.1–5.2) and let $e' \in E$ be any σ -violating edge. Define a positional strategy σ' of Max as follows:*

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

Then $\sum_{u \in V} \mathbf{Val}[\sigma'](u) > \sum_{u \in V} \mathbf{Val}[\sigma](u)$.

Proof. For $\mathbf{x} \in \varphi(A^\omega)^V$, let the *modified cost* of an edge $e \in E$ with respect to \mathbf{x} be the following quantity:

$$R^{\mathbf{x}}(e) = \mathbf{s}[\text{lab}(e)](\mathbf{x}_{\text{target}(e)}) - \mathbf{x}_{\text{source}(e)}.$$

We need the following “potential transformation lemma” (its analog for discounted payoffs is well-known, see, e.g., [15, Lemma 3.6]).

Lemma 7.4. *Take any $\mathbf{x} \in \varphi(A^\omega)^V$. Let $\mathcal{P} = e_1 e_2 e_3 \dots$ be an infinite path in G . Then there exists an infinite sequence of non-negative real numbers $\lambda_1, \lambda_2, \lambda_3, \dots$ such that $\lambda_1 = 1$ and*

$$\varphi \circ \text{lab}(\mathcal{P}) - \mathbf{x}_{\text{source}(\mathcal{P})} = \sum_{n=1}^{\infty} \lambda_n \cdot R^{\mathbf{x}}(e_n).$$

Proof. For $u \in V$, let $\beta_u \in A^\omega$ be such that $\mathbf{x}_u = \varphi(\beta_u)$. Define $s_n = \text{target}(e_1 e_2 \dots e_n)$ for $n \geq 1$ and $s_0 = \text{source}(\mathcal{P})$. By the continuity of φ , we have

$$\varphi \circ \text{lab}(\mathcal{P}) = \lim_{n \rightarrow \infty} \varphi(\text{lab}(e_1 e_2 \dots e_n) \beta_{s_n}) = \lim_{n \rightarrow \infty} \mathbf{s}[\text{lab}(e_1 e_2 \dots e_n)](\mathbf{x}_{s_n}).$$

Hence we obtain

$$\begin{aligned} \varphi \circ \text{lab}(\mathcal{P}) - \mathbf{x}_{s_0} &= \lim_{n \rightarrow \infty} (\mathbf{s}[\text{lab}(e_1 e_2 \dots e_n)](\mathbf{x}_{s_n}) - \mathbf{x}_{s_0}) \\ &= \lim_{n \rightarrow \infty} \sum_{k=1}^n (\mathbf{s}[\text{lab}(e_1 e_2 \dots e_k)](\mathbf{x}_{s_k}) - \mathbf{s}[\text{lab}(e_1 e_2 \dots e_{k-1})](\mathbf{x}_{s_{k-1}})) \\ &= \sum_{n=1}^{\infty} (\mathbf{s}[\text{lab}(e_1 e_2 \dots e_n)](\mathbf{x}_{s_n}) - \mathbf{s}[\text{lab}(e_1 e_2 \dots e_{n-1})](\mathbf{x}_{s_{n-1}})). \end{aligned}$$

We can write each term in this series as:

$$\begin{aligned} &\mathbf{s}[\text{lab}(e_1 e_2 \dots e_n)](\mathbf{x}_{s_n}) - \mathbf{s}[\text{lab}(e_1 e_2 \dots e_{n-1})](\mathbf{x}_{s_{n-1}}) \\ &= \mathbf{s}[\text{lab}(e_1 \dots e_{n-1})](\mathbf{s}[\text{lab}(e_n)](\mathbf{x}_{s_n})) - \mathbf{s}[\text{lab}(e_1 e_2 \dots e_{n-1})](\mathbf{x}_{s_{n-1}}) \\ &= \lambda_n \cdot (\mathbf{s}[\text{lab}(e_n)](\mathbf{x}_{s_n}) - \mathbf{x}_{s_{n-1}}), \end{aligned}$$

for some $\lambda_n \in [0, +\infty)$, because $\mathbf{s}[\text{lab}(e_1 \dots e_{n-1})]$ is non-decreasing (for $n = 1$ we get $\lambda_1 = 1$ because $\mathbf{s}[\text{empty word}]$ is the identity function). It remains to notice that by definition:

$$\mathbf{s}[\text{lab}(e_n)](\mathbf{x}_{s_n}) - \mathbf{x}_{s_{n-1}} = R^{\mathbf{x}}(e_n)$$

(because $\text{source}(e_n) = s_{n-1}$, $\text{target}(e_n) = s_n$). \square

We apply this lemma to the vector $\mathbf{g} = \{\mathbf{Val}[\sigma](u)\}_{u \in V}$. Note that by Lemma 7.1 we have $R^{\mathbf{g}}(e) \geq 0$ for every $e \in E^\sigma$. In turn, since e' is σ -violating, we have $R^{\mathbf{g}}(e') > 0$.

Let us at first show that

$$\mathbf{Val}[\sigma'](u) \geq \mathbf{Val}[\sigma](u) = \mathbf{g}_u$$

for every $u \in V$. In other words, we will demonstrate that $\varphi \circ \text{lab}(\mathcal{P}) \geq \mathbf{g}_u$ for any infinite path $\mathcal{P} = e_1 e_2 e_3 \dots \in \text{Cons}(u, \sigma')$. Indeed, by Lemma 7.4 we can write:

$$\varphi \circ \text{lab}(\mathcal{P}) - \mathbf{g}_u = \sum_{n=1}^{\infty} \lambda_n R^{\mathbf{g}}(e_n) \tag{7.2}$$

for some $\lambda_n \in [0, +\infty)$, $\lambda_1 = 1$. All edges of \mathcal{P} are from $E^\sigma \cup \{e'\}$. Hence, all terms in this series are non-negative, and so is the left-hand side.

To establish that $\sum_{u \in V} \mathbf{Val}[\sigma'](u) > \sum_{u \in V} \mathbf{Val}[\sigma](u)$, it is now enough to show that $\mathbf{Val}[\sigma'](u) > \mathbf{Val}[\sigma](u) = \mathbf{g}_u$ for some $u \in V$. We will show this for $u = \text{source}(e')$. The first edge of any $\mathcal{P} \in \text{Cons}(u, \sigma')$ is e' . So the first term in (7.2) for any such \mathcal{P} equals $R^{\mathbf{g}}(e')$. All the other terms, as we have discussed, are non-negative. Hence, $\varphi \circ \text{lab}(\mathcal{P}) \geq R^{\mathbf{g}}(e') + \mathbf{Val}[\sigma](u)$ for any $\mathcal{P} \in \text{Cons}(u, \sigma')$. Since $R^{\mathbf{g}}(e')$ is strictly positive, we get that $\mathbf{Val}[\sigma'](u) > \mathbf{Val}[\sigma](u)$. \square

By this lemma, a Max's positional strategy σ^* maximizing the quantity $\sum_{u \in V} \mathbf{Val}[\sigma](u)$ (over positional strategies σ of Max) gives a solution to (5.1–5.2). Such σ^* exists just because there are only finitely many positional strategies of Max. This finishes our strategy improvement proof of Proposition 5.4. Let us note that the same argument can be carried out with positional strategies of Min (via analogues of Lemma 7.1 and Lemma 7.3 for Min).

8. SUBEXPONENTIAL-TIME ALGORITHM

In this subsection, we discuss implications of our strategy improvement argument to the *strategy synthesis problem*. The strategy synthesis for a positionally determined payoff φ is an algorithmic problem of finding an equilibrium (with respect to φ) of two positional strategies in a given game graph. It is classical that the strategy synthesis for parity, mean and multi-discounted 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. For that, 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 in a single node.

Proposition 8.1. *Let A be a finite set and $\varphi: A^\omega \rightarrow \mathbb{R}$ be a continuous positionally determined shift-deterministic payoff. Then for any A -labeled game graph $G = \langle V, V_{\text{Max}}, V_{\text{Min}}, E \rangle$ 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 $|\{u \in V_{\text{Max}} \mid \sigma(u) \neq \sigma'(u)\}| = 1$ and $\sum_{u \in V} \mathbf{Val}[\sigma'](u) > \sum_{u \in V} \mathbf{Val}[\sigma](u)$;
- 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 $|\{u \in V_{\text{Min}} \mid \tau(u) \neq \tau'(u)\}| = 1$ and $\sum_{u \in V} \mathbf{Val}[\tau'](u) < \sum_{u \in V} \mathbf{Val}[\tau](u)$.

Proof. Assume that σ is a non-optimal positional strategy of Max. First, let us show that the vector $\mathbf{Val}[\sigma]$ cannot be a solution to Bellman's equations. Indeed, by Lemma 7.1, it holds that σ uses only edges that are $\mathbf{Val}[\sigma]$ -tight. Hence, if $\mathbf{Val}[\sigma]$ were a solution to Bellman's equation, then, by Lemma 5.5, strategy σ would have been optimal.

Since $\mathbf{Val}[\sigma]$ is not a solution to Bellman's equation, we can take σ' as in Lemma 7.3, obtained by switching σ to some σ -violating edge. An argument for positional strategies of Min is similar. \square

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 $u \in V_{\text{Max}}$, will take values in the set of edges that start at u . Obviously, vertices of such a 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 8.1 means the following: any vertex σ , maximizing $\sum_{u \in V} \mathbf{Val}[\sigma](u)$ over its neighbors, also maximizes this quantity over the *whole* hypercube.

So an optimization problem of maximizing $\sum_{u \in V} \mathbf{Val}[\sigma](u)$ (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 8.1. 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 8.1 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 8.2. *Let A be a finite set and $\varphi: A^\omega \rightarrow \mathbb{R}$ 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 solves the strategy synthesis problem for φ with this oracle in expected $e^{O(\log m + \sqrt{n \log m})}$ 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 $e^{O(\log m + \sqrt{n \log m})}$.*

Proof. See Appendix C. □

So to deal with the issue of representation we assume a suitable oracle access to φ . Still, the oracle from Theorem 8.2 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 infinite words 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 8.2. Let us note that $[\exists w \in A^* \varphi(w\alpha) > \varphi(w\beta)] \iff \varphi(\alpha) > \varphi(\beta)$ if φ is additionally shift-deterministic.

9. MULTI-DISCOUNTED PAYOFFS AND MDPs

In this section, we establish the following result.

Theorem 9.1. *Let A be a finite set and $\varphi: A^\omega \rightarrow \mathbb{R}$ be a continuous payoff. Then φ is positionally determined in MDPs if and only if φ is multi-discounted.*

This theorem disproves the following conjecture of Gimbert [8]: “Any payoff function which is positional for the class of non-stochastic one-player games is positional for the class of Markov decision processes”. Indeed, by Proposition 6.7, there exists a continuous positionally determined payoff which is not multi-discounted. By Theorem 9, this payoff is not positionally determined in MDPs.

A fact that multi-discounted payoffs are positionally determined in MDPs (and in two-player stochastic games as well) is classical [21]. In the rest of this section, we show that any continuous payoff which is positionally determined in MDPs is multi-discounted. First, we establish the following two necessary conditions.

Proposition 9.2. *Let A be a finite set and $\varphi: A^\omega \rightarrow \mathbb{R}$ be a continuous payoff which is positionally determined in MDPs. Then there are no $a \in A$, $\beta, \gamma, \delta \in A^\omega$, $(p_1, p_2, p_3), (q_1, q_2, q_3) \in [0, +\infty)^3$ such that $p_1 + p_2 + p_3 = q_1 + q_2 + q_3 = 1$ and*

$$\begin{aligned} p_1\varphi(\beta) + p_2\varphi(\gamma) + p_3\varphi(\delta) &> q_1\varphi(\beta) + q_2\varphi(\gamma) + q_3\varphi(\delta), \\ p_1\varphi(a\beta) + p_2\varphi(a\gamma) + p_3\varphi(a\delta) &< q_1\varphi(a\beta) + q_2\varphi(a\gamma) + q_3\varphi(a\delta). \end{aligned}$$

Proposition 9.3. *If a continuous payoff is positionally determined in MDPs, then this payoff is prefix-monotone.*

We also show that these two necessary conditions imply that φ is multi-discounted.

Proposition 9.4. *Let A be a finite set and $\varphi: A^\omega \rightarrow \mathbb{R}$ be a continuous prefix-monotone payoff. Assume that there are no $a \in A$, $\beta, \gamma, \delta \in A^\omega$, $(p_1, p_2, p_3), (q_1, q_2, q_3) \in [0, +\infty)^3$ such that $p_1 + p_2 + p_3 = q_1 + q_2 + q_3 = 1$ and*

$$\begin{aligned} p_1\varphi(\beta) + p_2\varphi(\gamma) + p_3\varphi(\delta) &> q_1\varphi(\beta) + q_2\varphi(\gamma) + q_3\varphi(\delta), \\ p_1\varphi(a\beta) + p_2\varphi(a\gamma) + p_3\varphi(a\delta) &< q_1\varphi(a\beta) + q_2\varphi(a\gamma) + q_3\varphi(a\delta). \end{aligned}$$

Then φ is a multi-discounted payoff.

Note that Proposition 9.3 is already proved. Indeed, in Section 3 we have shown that for any continuous payoff which is not prefix-monotone, there exists a game graph where φ is not positional. This game graph had the following feature: all its nodes were controlled by Max. Thus, this game graph is a deterministic MDP, which means that any continuous payoff which is not prefix monotone is not positionally determined in MDPs.

To finish our proof of Theorem 9.1, it remains to prove Propositions 9.2 and 9.4.

9.1. Proof of Proposition 9.2. Assume for contradiction that such a, β, γ, δ , (p_1, p_2, p_3) and (q_1, q_2, q_3) exist. By the continuity of φ , we may assume that β, γ and δ are ultimately periodic. We construct an A -labeled MDP \mathcal{M} where φ has no optimal positional strategy. To define \mathcal{M} , consider an A -labeled game graph from Figure 3.

In this graph there are exactly 3 infinite paths (“lassos”) $\mathcal{P}_1, \mathcal{P}_2, \mathcal{P}_3$ that start at v . We label their edges in such a way that $\text{lab}(\mathcal{P}_1) = \beta, \text{lab}(\mathcal{P}_2) = \gamma, \text{lab}(\mathcal{P}_3) = \delta$. This is possible because β, γ and δ are ultimately periodic.

Next, we turn this graph into an MDP (formally, nodes of the graph will be states of the MDP). There will be two actions available at the node v . Both will be distributed on the set of successors of v . One gives a probability p_i to the successor which leads to the




Figure 3: A graph for an MDP where φ has no optimal positional strategy.

lasso \mathcal{P}_i , for $i = 1, 2, 3$. The other gives a probability q_i to the successor which leads to the lasso \mathcal{P}_i , for $i = 1, 2, 3$. For each node different from v there will be only one action with the source in this node, leading with probability 1 to its unique successor.

It remains to define the labeling function of \mathcal{M} . Fix a transition. It is over some edge of the graph from Figure 3. We define the label this transition as the label of this edge. This concludes a description of \mathcal{M} .

To show that φ is not positional in \mathcal{M} , note that in \mathcal{M} there are exactly 2 positional strategies, σ_p and σ_q , corresponding to two actions available at v . We show that none of these two strategies is optimal.

It is easy to see that:

$$\begin{aligned}\mathbb{E}\varphi \circ \text{lab}(\mathcal{P}_u^{\sigma_p}) &= p_1 \cdot \varphi(a\beta) + p_2 \cdot \varphi(a\gamma) + p_3 \cdot \varphi(a\delta), \\ \mathbb{E}\varphi \circ \text{lab}(\mathcal{P}_v^{\sigma_p}) &= p_1 \cdot \varphi(\beta) + p_2 \cdot \varphi(\gamma) + p_3 \cdot \varphi(\delta), \\ \mathbb{E}\varphi \circ \text{lab}(\mathcal{P}_u^{\sigma_q}) &= q_1 \cdot \varphi(a\beta) + q_2 \cdot \varphi(a\gamma) + q_3 \cdot \varphi(a\delta), \\ \mathbb{E}\varphi \circ \text{lab}(\mathcal{P}_v^{\sigma_q}) &= q_1 \cdot \varphi(\beta) + q_2 \cdot \varphi(\gamma) + q_3 \cdot \varphi(\delta).\end{aligned}$$

Due to our assumptions about $(p_1, p_2, p_3), (q_1, q_2, q_3)$, we obtain:

$$\mathbb{E}\varphi \circ \text{lab}(\mathcal{P}_u^{\sigma_p}) < \mathbb{E}\varphi \circ \text{lab}(\mathcal{P}_u^{\sigma_q}), \quad \mathbb{E}\varphi \circ \text{lab}(\mathcal{P}_v^{\sigma_p}) > \mathbb{E}\varphi \circ \text{lab}(\mathcal{P}_v^{\sigma_q}).$$

Therefore, neither σ_p nor σ_q is optimal.

9.2. Proof of Proposition 9.4. If $\varphi(\gamma) = \varphi(\delta)$ for all $\beta, \gamma \in A^\omega$, then clearly φ is multi-discounted (one can define $\lambda(a) = 0, w(a) = \varphi(\gamma)$ for all $a \in A$ and for an arbitrary $\gamma \in A^\omega$). So in what follows we fix any $\gamma, \delta \in A^\omega$ with $\varphi(\gamma) \neq \varphi(\delta)$. First we derive from the conditions of Proposition 9.4 the following:

Lemma 9.5. *For any $a \in A$ there exist $\lambda(a), w(a) \in \mathbb{R}$ such that for any $\beta \in A^\omega$ we have:*

$$\varphi(a\beta) = \lambda(a)\varphi(\beta) + w(a).$$

Proof. The following system in (λ, w) has a unique solution:

$$\begin{pmatrix} \varphi(a\gamma) \\ \varphi(a\delta) \end{pmatrix} = \begin{pmatrix} \varphi(\gamma) & 1 \\ \varphi(\delta) & 1 \end{pmatrix} \cdot \begin{pmatrix} \lambda \\ w \end{pmatrix}, \tag{9.1}$$

(because $\varphi(\gamma) \neq \varphi(\delta)$). Let its solution be $(\lambda(a), w(a))$. We show that $\varphi(a\beta) = \lambda(a)\varphi(\beta) + w(a)$ for all $\beta \in A^\omega$. Let us first show that

$$\det \begin{pmatrix} 1 & 1 & 1 \\ \varphi(\beta) & \varphi(\gamma) & \varphi(\delta) \\ \varphi(a\beta) & \varphi(a\gamma) & \varphi(a\delta) \end{pmatrix} = 0. \quad (9.2)$$

Indeed, otherwise there exists a vector $(x, y, z) \in \mathbb{R}^3$ such that

$$\begin{pmatrix} 1 & 1 & 1 \\ \varphi(\beta) & \varphi(\gamma) & \varphi(\delta) \\ \varphi(a\beta) & \varphi(a\gamma) & \varphi(a\delta) \end{pmatrix} \cdot \begin{pmatrix} x \\ y \\ z \end{pmatrix} = \begin{pmatrix} 0 \\ 1 \\ -1 \end{pmatrix}. \quad (9.3)$$

Let $P_1, P_2, P_3, Q_1, Q_2, Q_3$ be any positive real numbers such that $x = P_1 - Q_1, y = P_2 - Q_2, z = P_3 - Q_3$. From the first equality in (9.3) it follows that $P_1 + P_2 + P_3 = Q_1 + Q_2 + Q_3 = S > 0$. Define $p_i = P_i/S, q_i = Q_i/S$ for $i \in \{1, 2, 3\}$. Observe that $a, \beta, \gamma, \rho, (p_1, p_2, p_3), (q_1, q_2, q_3)$ violate the conditions of Proposition 9.4 (this can be seen from the second and the third equalities in (9.3)), contradiction. Therefore (9.2) is proved.

The first two rows of the matrix from (9.2) are linearly independent because $\varphi(\gamma) \neq \varphi(\delta)$. Hence, the third one must be a linear combination of the first two. I.e., there must exist $\lambda, w \in \mathbb{R}$ such that

$$(\varphi(a\beta), \varphi(a\gamma), \varphi(a\delta)) = \lambda(\varphi(\beta), \varphi(\gamma), \varphi(\delta)) + w(1, 1, 1).$$

From the second and the third coordinate we conclude that (λ, w) must be a solution to (9.1), so $\lambda = \lambda(a), w = w(a)$. Now, looking at the first coordinate, we obtain that $\varphi(a\beta) = \lambda(a)\varphi(\beta) + w(a)$, as required. \square

From now on let $\lambda(a), w(a)$ for $a \in A$ be as in Lemma 9.5. Let us show that $\lambda(a) \in [0, 1]$ for all $a \in A$.

Assume first that for some $a \in A$ we have $\lambda(a) < 0$. Without loss of generality, we may also assume that $\varphi(\gamma) < \varphi(\delta)$. Then $\varphi(a\gamma) = \lambda(a)\varphi(\gamma) + w(a) > \lambda(a)\varphi(\delta) + w(a) = \varphi(a\delta)$. But φ is prefix-monotone, so this is impossible.

Next, assume for contradiction that $\lambda(a) \geq 1$ for some $a \in A$. Consider the following two sequences $\{x_n\}_{n \in \mathbb{N}}$ and $\{y_n\}_{n \in \mathbb{N}}$ of real numbers:

$$x_n = \varphi(\underbrace{aa \dots a}_n \gamma), \quad y_n = \varphi(\underbrace{aa \dots a}_n \delta).$$

Note that by our choice of γ and δ , we have $x_0 = \varphi(\gamma) \neq \varphi(\delta) = y_0$. Next, since φ is continuous, we have:

$$\lim_{n \rightarrow \infty} x_n = \lim_{n \rightarrow \infty} y_n = \varphi(aaa \dots). \quad (9.4)$$

On the other hand, we can compute x_n and y_n through Lemma 9.5:

$$x_n = \lambda(a)^n x_0 + w(a)(1 + \lambda(a) + \dots + \lambda(a)^{n-1}), \quad (9.5)$$

$$y_n = \lambda(a)^n y_0 + w(a)(1 + \lambda(a) + \dots + \lambda(a)^{n-1}). \quad (9.6)$$

First, consider the case $\lambda(a) = 1$. Then x_n and y_n look as follows:

$$x_n = x_0 + nw(a), \quad y_n = y_0 + nw(a).$$

If $w(a) \neq 0$, then the sequences $\{x_n\}_{n \in \mathbb{N}}$ and $\{y_n\}_{n \in \mathbb{N}}$ are not convergent, contradiction with (9.4). If $w(a) = 0$, then one sequence converges to x_0 , and the other to y_0 . But $x_0 \neq y_0$, so this again gives a contradiction with (9.4).

Now, consider the case $\lambda(a) > 1$. Then we can rewrite (9.5–9.6) as follows:

$$x_n = \lambda(a)^n \left(x_0 + \frac{w(a)}{\lambda(a) - 1} \right) - \frac{w(a)}{\lambda(a) - 1}, \quad y_n = \lambda(a)^n \left(y_0 + \frac{w(a)}{\lambda(a) - 1} \right) - \frac{w(a)}{\lambda(a) - 1}.$$

Since $x_0 \neq y_0$, the coefficient before $\lambda(a)^n$ is non-zero in at least one of these expressions. Hence either $\{x_n\}_{n \in \mathbb{N}}$ or $\{y_n\}_{n \in \mathbb{N}}$ diverge. This contradicts (9.4).

We have established that $\lambda(a) \in [0, 1)$ for every $a \in A$. All that remains to do is to show that φ satisfies (3.1). For that, we again employ the continuity of φ . Take any $\beta \in A^\omega$. Note that by Lemma 9.5 we have:

$$\varphi(a_1 a_2 \dots a_n \beta) = \lambda(a_1) \cdot \dots \cdot \lambda(a_n) \varphi(\beta) + \sum_{i=1}^n \lambda(a_1) \cdot \dots \cdot \lambda(a_{i-1}) \cdot w(a_i). \quad (9.7)$$

We know that $\lambda(a_i)$ are all from $[0, 1)$. Since the set A is finite, all $\lambda(a_i)$ are bounded from above by some number smaller than 1. Hence, the first term in the right-hand side of (9.7) converges to 0 as $n \rightarrow \infty$. On the other hand, the second term in the right-hand side of (9.7) converges to the series from the right-hand side of (3.1). Finally, due to the continuity of φ , the left-hand side of (9.7) converges to $\varphi(a_1 a_2 a_3 \dots)$. Thus, φ is multi-discounted.

10. 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? Is it at least true for prefix-independent positionally determined 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. A fact that a non-optimal positional strategy can be improved by modifying it in a single node lies in the core of all subexponential-time algorithms for positionally determined payoffs [14, 1]. So if we could always prove positional determinacy by strategy improvement, maybe we can also solve any positionally determined payoff in subexponential time?

Finally, it would be interesting to obtain an explicit description of other classes of positionally determined payoffs – for example, of prefix-independent positionally determined payoffs.

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 Vandenholve. 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 Khossainov, Wei Li, and Frank Stephan. Deciding parity games in quasi-polynomial time. *SIAM Journal on Computing*, (0):STOC17–152, 2020.
- [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 Wiesław Zielonka. When can you play positionally? In *International Symposium on Mathematical Foundations of Computer Science*, pages 686–697. Springer, 2004.
- [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 Wiesław 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. 1960.
- [17] Marcin Jurdziński. Deciding the winner in parity games is in $\text{up}\cap\text{co-up}$. *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] Robert McNaughton. Infinite games played on finite graphs. *Annals of Pure and Applied Logic*, 65(2):149–184, 1993.
- [20] Andrzej W Mostowski. Games with forbidden positions. Technical Report 78, Uniwersytet Gdański, Instytut Matematyki, 1991.
- [21] Martin L Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014.
- [22] Lloyd S Shapley. Stochastic games. *Proceedings of the national academy of sciences*, 39(10):1095–1100, 1953.
- [23] Richard S Sutton, Andrew G Barto, et al. *Introduction to reinforcement learning*, volume 135. MIT press Cambridge, 1998.
- [24] Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343–359, 1996.

APPENDIX A. PROOF OF PROPOSITION 2.4

First, assume that φ is continuous. Take any $\varepsilon > 0$. We have to show that for some n_0 it holds that $\varphi(\beta_n) \in (\varphi(\alpha) - \varepsilon, \varphi(\alpha) + \varepsilon)$ for all $n \geq n_0$. The set $\varphi^{-1}((\varphi(\alpha) - \varepsilon, \varphi(\alpha) + \varepsilon))$ must be open. So for some $S \subseteq A^*$ we have:

$$\varphi^{-1}((\varphi(\alpha) - \varepsilon, \varphi(\alpha) + \varepsilon)) = \bigcup_{u \in S} uA^\omega.$$

Since obviously $\alpha \in \varphi^{-1}((\varphi(\alpha) - \varepsilon, \varphi(\alpha) + \varepsilon))$, there exists $u \in S$ such that $\alpha \in uA^\omega$. Hence for $n \geq |u|$ we have $\beta_n \in uA^\omega \subseteq \varphi^{-1}((\varphi(\alpha) - \varepsilon, \varphi(\alpha) + \varepsilon))$, as required.

Let us now establish the other direction of the proposition. It is enough to show that for any $x, y \in \mathbb{R}$ with $x < y$ the set $\varphi^{-1}((x, y))$ is open. Take any $\alpha \in \varphi^{-1}((x, y))$. Let us show that there exists $n(\alpha)$ such that all $\beta \in A^\omega$ that coincide with α in the first $n(\alpha)$ elements belong to $\varphi^{-1}((x, y))$. Indeed, otherwise for any n there exists β_n , coinciding with α in the first n elements, such that $\beta_n \notin \varphi^{-1}((x, y))$. Now, the limit $\lim_{n \rightarrow \infty} \varphi(\beta_n)$ must exist and must be equal to $\varphi(\alpha)$. But $\varphi(\alpha) \in (x, y)$ and all $\varphi(\beta_n)$ are not in this interval, contradiction.

Now, for $\alpha \in \varphi^{-1}((x, y))$ let $u_\alpha \in A^{n(\alpha)}$ be the $n(\alpha)$ -length prefix of α . Observe that

$$\varphi^{-1}((x, y)) = \bigcup_{\alpha \in \varphi^{-1}((x, y))} u_\alpha A^\omega.$$

So the set $\varphi^{-1}((x, y))$ is open, as required.

APPENDIX B. PROOF OF THEOREM 5.6

For the sake of readability, we will use the following notation. First, we will denote f_i by $f[i]$. Moreover, we will abbreviate

$$f[a_1 a_2 \dots a_n] = f[a_1] \circ f[a_2] \circ \dots \circ f[a_n]$$

for $n \in \mathbb{N}$, $a_1 a_2 \dots a_n \in \{1, 2, \dots, m\}^n$. In particular, $f[\text{empty word}]$ will denote the identity function).

Lemma B.1. *The condition (a) of Theorem 5.6 is equivalent to the following condition: for every $\varepsilon > 0$ there are only finitely many $w \in \{1, 2, \dots, m\}^*$ such that*

$$\text{diam}(f[w](K)) > \varepsilon.$$

Proof. Assume that the condition (a) of Theorem 5.6 holds. Take any $\varepsilon > 0$. Call $w \in \{1, 2, \dots, m\}^*$ bad if $\text{diam}(f[w](K)) > \varepsilon$. We have to show that the number of bad w is finite. Assume for contradiction that the number of bad w is infinite. Observe that any prefix of a bad w is also bad. Indeed, if $w = uv$ for some $u, v \in \{1, 2, \dots, m\}^*$, then $f[w](K) = f[u] \circ f[v](K) \subseteq f[u](K)$. Hence, by König's Lemma, there exists $\alpha = a_1 a_2 a_3 \dots \in \{1, 2, \dots, m\}^\omega$ such any finite prefix of α is bad. Observe that

$$\liminf_{n \rightarrow \infty} \text{diam}(f[a_1 a_2 \dots a_n](K)) \geq \varepsilon.$$

This is a contradiction with the condition (a) of Theorem 5.6.

The opposite direction of the lemma is obvious. □

The rest of the proof is organized as follows. We first show that (a) \implies (b). Then we observe that the same proof establishes (a) \implies (c) when $f[1], \dots, f[m]$ are non-decreasing. Finally, we show that (b) \implies (a). Since, obviously, (c) \implies (b), this will establish Theorem 5.6.

Proof of (a) \implies (b). Define

$$d: K \times K \rightarrow [0, +\infty), \quad d(x, y) = \sup_{w \in \{1, \dots, m\}^*} (2 - 2^{-|w|}) \cdot |f[w](x) - f[w](y)|. \quad (\text{B.1})$$

First, we obviously have $d(x, x) = 0$ and $d(x, y) = d(y, x)$. Notice also that $d(x, y) \geq |x - y|$, so $d(x, y) > 0$ for $x \neq y$. In turn, the triangle inequality $d(x, y) \leq d(x, z) + d(z, y)$ holds because, first, it holds for every $w \in \{1, 2, \dots, m\}^*$ in (B.1), and second, the supremum

of the sums is at most the sum of the supremums. These considerations show that d is a metric.

Note that the supremum in (B.1) is always attained on some $w \in \{1, 2, \dots, m\}^*$. Indeed, if $d(x, y) = 0$, then it is attained already on the empty word. Assume now that $d(x, y) > 0$. By Lemma B.1 all but finitely many $w \in \{1, 2, \dots, m\}^*$ satisfy $\text{diam}(f[w](K)) \leq d(x, y)/3$. So only finitely many terms in (B.1) are bigger than $2d(x, y)/3$, and hence the supremum (which is $d(x, y) > 2d(x, y)/3$) must be attained on one of them.

This already implies that $f[i]$ is d -contracting for every $i \in \{1, 2, \dots, m\}$. Indeed, take any $x, y \in K$. Then for some $w \in \{1, 2, \dots, m\}^*$ we have:

$$d(f[i](x), f[i](y)) = (2 - 2^{-|w|}) \cdot |f[w](f[i](x)) - f[w](f[i](y))|.$$

We have to show that if $x \neq y$, then $d(f[i](x), f[i](y)) < d(x, y)$. If $d(f[i](x), f[i](y)) = 0$, there is nothing to prove. Otherwise, the quantity

$$|f[w](f[i](x)) - f[w](f[i](y))|$$

is positive. Therefore, we can write:

$$\begin{aligned} d(f[i](x), f[i](y)) &= (2 - 2^{-|w|}) \cdot |f[w](f[i](x)) - f[w](f[i](y))| \\ &< (2 - 2^{-|w|-1}) \cdot |f[w](f[i](x)) - f[w](f[i](y))| \\ &= (2 - 2^{-|wi|}) \cdot |f[wi](x) - f[wi](y)| \leq d(x, y). \end{aligned}$$

It remains to show that d is continuous. Consider any $(x_0, y_0) \in K \times K$. We have to show that for any $\varepsilon > 0$ there exists $\delta > 0$ such that for all $(x, y) \in K \times K$ with $|x - x_0| + |y - y_0| \leq \delta$ we have $|d(x, y) - d(x_0, y_0)| \leq \varepsilon$.

By Lemma B.1, there exists $n \in \mathbb{N}$ such that for all $w \in \{1, 2, \dots, m\}^*$ with $|w| \geq n$ we have:

$$\text{diam}(f[w](K)) \leq \varepsilon/6.$$

In particular, this means that all terms in (B.1) corresponding to $w \in \{1, 2, \dots, m\}^*$ with $|w| \geq n$ are at most $\varepsilon/3$. Hence, for every $(x, y) \in K \times K$ we have that $d(x, y)$ is $(\varepsilon/3)$ -close to $d_n(x, y)$, where

$$d_n(x, y) = \max_{w \in A^*, |w| < n} (2 - 2^{-|w|}) \cdot |f[w](x) - f[w](y)|.$$

Now, notice that the function d_n is continuous (as a composition of finitely many continuous functions). Hence there exists $\delta > 0$ such that for all $(x, y) \in K \times K$ with $|x - x_0| + |y - y_0| \leq \delta$ we have $|d_n(x, y) - d_n(x_0, y_0)| \leq \varepsilon/3$. Obviously, for all such (x, y) we also have $|d(x, y) - d(x_0, y_0)| \leq \varepsilon$.

Remark B.2. Here we observe that if $f[1], \dots, f[m]$ are non-decreasing, then this construction establishes **(a)** \implies **(c)**. That is, we show that if $f[1], \dots, f[m]$ are non-decreasing, then $d(s, t) \leq d(x, y)$ for all $x, s, t, y \in K$ with $x \leq s \leq t \leq y$. Indeed, in this case the function $f[w]$ is non-decreasing for every $w \in \{1, 2, \dots, m\}^*$. Hence we have $f[w](x) \leq f[w](s) \leq f[w](t) \leq f[w](y)$ and $|f[w](s) - f[w](t)| \leq |f[w](x) - f[w](y)|$. By (B.1), this gives us $d(s, t) \leq d(x, y)$.

Proof of (b) \implies (a). We show that for every $\varepsilon > 0$ there exists $n \in \mathbb{N}$ such that for all $w \in \{1, 2, \dots, m\}^*$ with $|w| \geq n$ it holds that

$$\text{diam}(f[w](K)) \leq \varepsilon.$$

Obviously, this implies **(a)**.

Define $T = \{(x, y) \in K \times K \mid |x - y| \geq \varepsilon\}$. Note that T is a compact set. A function $d(x, y)/|x - y|$ is continuous on T . Hence, there exists

$$z = \min_{(x, y) \in T} d(x, y)/|x - y|.$$

Observe that $z > 0$. Indeed, for some $(x, y) \in T$ we have $z = d(x, y)/|x - y|$. By definition of T , we have $|x - y| \geq \varepsilon$. Hence $x \neq y$ and $d(x, y)$ is positive, as well as z .

Now, define $S = \{(x, y) \in K \times K \mid d(x, y) \geq z \cdot \varepsilon\}$. Again, S is a compact set. Consider a function:

$$h(x, y) = \max_{i \in \{1, \dots, m\}} \frac{d(f[i](x), f[i](y))}{d(x, y)}.$$

The function h is continuous on S (we never have 0 in its denominator on S). Hence there exists

$$\lambda = \max_{(x, y) \in S} h(x, y).$$

The function h is non-negative, so $\lambda \geq 0$. Let us show that $\lambda < 1$. Indeed, for some $(x, y) \in S$ we have $\lambda = h(x, y)$. By definition of h , for some $i \in \{1, 2, \dots, m\}$ we have:

$$\lambda = \frac{d(f[i](x), f[i](y))}{d(x, y)}.$$

Since $(x, y) \in S$, we have $d(x, y) \geq z \cdot \varepsilon > 0$. Hence $x \neq y$. Now, $f[i]$ is d -contracting. Therefore $d(f[i](x), f[i](y)) < d(x, y)$ and $\lambda < 1$.

Define $D = \sup_{x, y \in K} d(x, y)$. If $D = 0$, then K consists of a single point, which means that the condition **(a)** trivially holds. From now on we assume that $D > 0$. Take any $n \in \mathbb{N}$ such that

$$\lambda^n < \frac{z\varepsilon}{D}.$$

We claim that for any $w \in \{1, 2, \dots, m\}^*$ with $|w| \geq n$ we have $\text{diam}(f[w](K)) \leq \varepsilon$. We only have to show this for w of length exactly n . This is because if w' is of length at least n , then $f[w'](K)$ is contained in $f[w](K)$, where w is a prefix of w' of length n .

So take any $w = a_1 a_2 \dots a_n \in \{1, 2, \dots, m\}^n$. Let us first establish that:

$$\sup_{x, y \in K} d(f[w](x), f[w](y)) \leq z\varepsilon. \tag{B.2}$$

Assume for contradiction that $d(f[w](x), f[w](y)) > z\varepsilon$ for some $x, y \in K$. Define $w_{\geq i} = a_i a_{i+1} \dots a_n$ for $i = 1, \dots, n$, and let $w_{\geq n+1}$ be the empty string. Set

$$F_i = d(f[w_{\geq i}](x), f[w_{\geq i}](y)).$$

Note that $F_1 = d(f[w](x), f[w](y))$ and $F_{n+1} = d(x, y)$. Since $f[w_{\geq i}] = f[a_i] \circ f[w_{\geq i+1}]$, and since $f[a_i]$ is d -contracting, for every $i = 1, \dots, n$ we have

$$F_i = d(f[w_{\geq i}](x), f[w_{\geq i}](y)) \leq d(f[w_{\geq i+1}](x), f[w_{\geq i+1}](y)) = F_{i+1}.$$

In fact, if $F_{i+1} \geq z\varepsilon$, then, by definition of λ , it holds that $F_i \leq \lambda F_{i+1}$. Recall that $F_1 = d(f[w](x), f[w](y)) > z\varepsilon$. Since $F_1 \leq F_2 \leq \dots \leq F_{n+1}$, we have $F_i \geq z\varepsilon$ for every i . Therefore, $F_1 \leq \lambda F_2 \leq \dots \leq \lambda^n F_{n+1}$. On the other hand, $F_{n+1} = d(x, y) \leq D$. Hence, $z\varepsilon < F_1 \leq \lambda^n D$. But by definition of n we have $\lambda^n D < z\varepsilon$, contradiction.

It remains to show that (B.2) implies that $\text{diam}(f[w](K)) \leq \varepsilon$. We do so by showing for every $x, y \in K$ that $|f[w](x) - f[w](y)| > \varepsilon \implies d(f[w](x), f[w](y)) > z\varepsilon$. This means that if $\text{diam}(f[w](K)) > \varepsilon$, then (B.2) cannot hold.

Now, if $|f[w](x) - f[w](y)| > \varepsilon$, then $(f[w](x), f[w](y)) \in T$, so

$$\frac{d(f[w](x), f[w](y))}{|f[w](x) - f[w](y)|} \geq \min_{(x,y) \in T} \frac{d(x,y)}{|x-y|} = z \neq 0.$$

Therefore, $d(f[w](x), f[w](y)) \geq z \cdot |f[w](x) - f[w](y)| > z\varepsilon$.

APPENDIX C. PROOF OF THEOREM 8.2

First in Subsection C.1 it is demonstrated that w.l.o.g. we may assume that φ is shift-deterministic (so that we can use Proposition 8.1) and that we are given an oracle which simply compares values of φ on ultimately periodic infinite words. Then in Subsection C.2 we expose a framework of *recursively local-global functions* due to Björklund and Vorobyov. Finally, in Subsection C.3 we use this framework to show Theorem 8.2 in the assumptions of Subsection C.1.

C.1. Reducing to shift-deterministic payoffs. It is sufficient to establish Theorem 8.2 with the following assumptions.

Assumption C.1. *Payoff φ is continuous, positionally determined and **shift-deterministic**.*

Assumption C.2. *We are given an oracle which for $u, v, a, b \in A^*$ tells, whether $\varphi(u(v)^\omega) > \varphi(a(b)^\omega)$.*

To justify this, it is enough to show the following lemma.

Lemma C.3. *Let A be a finite set and let $\varphi: A^\omega \rightarrow \mathbb{R}$ be a continuous positionally determined payoff. Then there exist a continuous positionally determined shift-deterministic payoff $\psi: A^\omega \rightarrow \mathbb{R}$ and a non-decreasing function $g: \psi(A^\omega) \rightarrow \mathbb{R}$ such that $\varphi = g \circ \psi$ and*

$$[\exists w \in A^* \varphi(w\alpha) > \varphi(w\beta)] \iff \psi(\alpha) > \psi(\beta) \quad \text{for all } \alpha, \beta \in A^\omega.$$

Indeed, let φ be an arbitrary continuous positionally determined payoff and ψ be as in Lemma C.3. By Proposition 2.2, an equilibrium for ψ is also an equilibrium for $\varphi = g \circ \psi$. So to solve the strategy synthesis for φ , it is enough to do so for ψ . Clearly, ψ satisfies Assumption C.1. Finally, note that the oracle from Assumption C.2 for ψ simply coincides on every input with the oracle we are given for φ in Theorem 8.2.

Proof of Lemma C.3. We take ψ as in the proof of Proposition 4.6. It is established there that

- ψ is continuous, prefix-monotone (hence positionally determined) and shift-deterministic;
- $\varphi = g \circ \psi$ for some non-decreasing $g: \psi(A^\omega) \rightarrow \mathbb{R}$.

This information is sufficient to show that

$$[\exists w \in A^* \varphi(w\alpha) > \varphi(w\beta)] \implies \psi(\alpha) > \psi(\beta).$$

Indeed, if $g \circ \psi(w\alpha) = \varphi(w\alpha) > \varphi(w\beta) = g \circ \psi(w\beta)$ for some $w \in A^*$, then we also have $\psi(w\alpha) > \psi(w\beta)$, because g is non-decreasing. Due to prefix-monotonicity of ψ , we also have $\psi(\alpha) \geq \psi(\beta)$. It remains to demonstrate that $\psi(\alpha) \neq \psi(\beta)$. Indeed, $\psi(\alpha) = \psi(\beta) \implies \psi(w\alpha) = \psi(w\beta)$ because ψ is shift-deterministic.

To demonstrate that

$$\psi(\alpha) > \psi(\beta) \implies [\exists w \in A^* \varphi(w\alpha) > \varphi(w\beta)]$$

we have to recall the construction of ψ . By (4.1), we can write:

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

If $\varphi(w\alpha) \leq \varphi(w\beta)$ for all $w \in A^*$, then clearly $\psi(\alpha) \leq \psi(\beta)$. This is exactly the contraposition to the implication that we have to prove. \square

C.2. Recursively local-global functions. Fix $d \in \mathbb{N}$. A d -dimensional *structure* is a collection $\mathcal{S} = \{S_i\}_{i=1}^d$ of d non-empty finite sets S_1, S_2, \dots, S_d . *Vertices* of \mathcal{S} are elements of the Cartesian product $\prod_{i=1}^d S_i$. Two vertices $\sigma = (\sigma_1, \dots, \sigma_d), \sigma' = (\sigma'_1, \dots, \sigma'_d) \in \prod_{i=1}^d S_i$ of a structure \mathcal{S} are called *neighbors* if there is exactly one $i \in \{1, 2, \dots, d\}$ such that $\sigma_i \neq \sigma'_i$. A structure $\mathcal{S}' = \{S'_i\}_{i=1}^d$ is a *substructure* of a structure $\mathcal{S} = \{S_i\}_{i=1}^d$ if $S'_i \subseteq S_i$ for every $i \in \{1, 2, \dots, d\}$.

Let \mathcal{S} be a structure and f be a function from the set of vertices of \mathcal{S} to \mathbb{R} . A vertex σ of \mathcal{S} is called a *local* maximum of f if $f(\sigma) \geq f(\sigma')$ for every neighbor σ' of σ in \mathcal{S} . A vertex σ is called a *global* maximum of f if $f(\sigma) \geq f(\sigma')$ for every vertex σ' of \mathcal{S} . The function f is called *local-global* if all its local maxima are global. The function f is called *recursively local-global* if all its restrictions to substructures of \mathcal{S} are local-global.

Given a structure \mathcal{S} and a function f from the set of vertices of \mathcal{S} to \mathbb{R} , we are interested in finding a global maximum of f . In [1] Björklund and Vorobyov obtained the following result.

Theorem C.4 (Theorem 5.1 in [1]). *Let $\mathcal{S} = \{S_i\}_{i=1}^d$ be a d -dimensional structure and $f: \prod_{i=1}^d S_i \rightarrow \mathbb{R}$ be a recursively local-global function.*

Consider an oracle which, given two vertices σ^1 and σ^2 of \mathcal{S} that are neighbors of each other, compares $f(\sigma^1)$ and $f(\sigma^2)$. There is a randomized algorithm which find a global maximum of f with this oracle in expected

$$e^{O(\log m + \sqrt{d \log m})} \text{ time,}$$

where $m = \sum_{i=1}^d |S_i|$.

C.3. Deriving Theorem 8.2 with Assumptions C.1 and C.2. Let $G = \langle V, V_{\text{Max}}, V_{\text{Min}}, E \rangle$ be an A -labeled game graph in which we want to solve the strategy synthesis. We will only show how to find an optimal positional strategy of Max, an argument for Min is similar.

Let $d = |V_{\text{Max}}|$ and $V_{\text{Max}} = \{u_1, u_2, \dots, u_d\}$. Define $S_i = \{e \in E \mid \text{source}(e) = u_i\}$. Consider a structure $\mathcal{S} = \{S_i\}_{i=1}^d$. Obviously, we may identify vertices of \mathcal{S} with positional strategies of Max. Define

$$f: \prod_{i=1}^d S_i \rightarrow \mathbb{R}, \quad f(\sigma) = \sum_{u \in V} \mathbf{Val}[\sigma](u).$$

Lemma C.5. *Any global maximum of f is an optimal positional strategy of Max.*

Proof. Let σ be a global maximum of f and σ^* be any optimal positional strategy of Max. By optimality of σ^* , we have $\mathbf{Val}[\sigma^*](u) \geq \mathbf{Val}[\sigma](u)$ for every $u \in V$. On the other hand, σ maximizes the sum of the values (over all positional strategies of Max), so we must have $\mathbf{Val}[\sigma^*](u) = \mathbf{Val}[\sigma](u)$ for every $u \in V$. This means that σ is also optimal. \square

Lemma C.6. *The function f is recursively local-global.*

Proof. A fact that f is local-global is a simple consequence of Proposition 8.1 (note that by Assumption C.1, our payoff satisfies the requirements of this proposition). Indeed, a strategy σ which is not a global maximum of f cannot be optimal. Then take σ' as in Proposition 8.1. It is a neighbor of σ with $f(\sigma') > f(\sigma)$, so σ cannot be a local maximum as well.

To show that f is recursively local-global, it is sufficient to note that substructures of \mathcal{S} correspond to subgraphs of G , and for these subgraphs we also have Proposition 8.1. \square

Due to these two lemmas, if we run the algorithm from Theorem C.4, we get an optimal positional strategy of Max in expected

$$e^{O(\log m + \sqrt{d \log m})} \text{ time,}$$

where $m = \sum_{i=1}^d |S_i|$. Note that d does not exceed the number of nodes of G and m does not exceed the number of edges, so Theorem 8.2 follows.

Still, the algorithm from Theorem C.4 requires an oracle which, given any two vertices σ^1 and σ^2 of \mathcal{S} that are neighbors of each other, compares $f(\sigma^1)$ and $f(\sigma^2)$. In our case, this oracle, given two positional strategies σ^1, σ^2 of Max that differ from each other in exactly one node, compares the sums of their values:

$$\sum_{u \in V} \mathbf{Val}[\sigma^1](u), \quad \sum_{u \in V} \mathbf{Val}[\sigma^2](u).$$

We have to perform this comparison using the oracle from Assumption C.2.

Assume that the node where σ^1 and σ^2 differ is v . Let G^1 (respectively, G^2) be a game graph obtained from G by deleting all edges that are not consistent with σ^1 (resp., σ^2). Next, let $G^{1,2}$ be a game graph of all edges that appear either in G^1 or in G^2 .

Observe that in $G^{1,2}$, strategies σ^1, σ^2 are the only two positional strategies of Max (indeed, all nodes of Max except v have exactly one out-going edge in $G^{1,2}$, and v has exactly two). One of these strategies must be optimal in $G^{1,2}$. So either $\mathbf{Val}[\sigma^1](u) \geq \mathbf{Val}[\sigma^2](u)$ for all $u \in V$ or $\mathbf{Val}[\sigma^1](u) \leq \mathbf{Val}[\sigma^2](u)$ for all $u \in V$. This means that

$$\begin{aligned} \sum_{u \in V} \mathbf{Val}[\sigma^1](u) &> \sum_{u \in V} \mathbf{Val}[\sigma^2](u) \\ \iff \exists u \in V \mathbf{Val}[\sigma^1](u) &> \mathbf{Val}[\sigma^2](u). \end{aligned}$$

So our task reduces to a task of comparing $\mathbf{Val}[\sigma^1](u)$ and $\mathbf{Val}[\sigma^2](u)$ for $u \in V$.

Assume first that our game graph G is *one-player*. This means that for one of the players it holds that all nodes of this player have out-degree 1. In our case, this must be Min, because Max has two distinct positional strategies σ^1 and σ^2 . In particular, there is exactly one strategy τ of Min in G , and this strategy is positional (even if there are no nodes controlled by Min, we assume that Min has a unique empty strategy τ). Hence, $\mathbf{Val}[\sigma^1](u) = \varphi \circ \mathbf{lab}(\mathcal{P}_u^{\sigma^1, \tau})$ and $\mathbf{Val}[\sigma^2](u) = \varphi \circ \mathbf{lab}(\mathcal{P}_u^{\sigma^2, \tau})$. It remains to compare the value of φ on $\mathbf{lab}(\mathcal{P}_u^{\sigma^1, \tau})$ and on $\mathbf{lab}(\mathcal{P}_u^{\sigma^2, \tau})$ using the oracle from Assumption C.2. These two infinite words are written over

some lassos in G , so we can decompose them as $\text{lab}(\mathcal{P}_u^{\sigma^1, \tau}) = u(v)^\omega$ and $\text{lab}(\mathcal{P}_u^{\sigma^2, \tau}) = a(b)^\omega$ in polynomial time.

Theorem 8.2 is already proved for one-player game graphs. Hence, at the cost of increasing the expected running time by a factor of $e^{O(\log m + \sqrt{n \log m})}$, we may assume that we also have an oracle which can solve the strategy synthesis for φ in one-player game graphs. Then we can find an optimal Min's positional strategy τ^1 in G^1 and an optimal Min's positional strategy τ^2 in G^2 . Indeed, in these two graphs all nodes of Max have exactly one out-going edge. Observe that τ^1 is an optimal response to σ^1 and τ^2 is an optimal response to σ^2 , so we have:

$$\mathbf{Val}[\sigma^1](u) = \varphi \circ \text{lab}(\mathcal{P}_u^{\sigma^1, \tau^1}), \quad \mathbf{Val}[\sigma^2](u) = \varphi \circ \text{lab}(\mathcal{P}_u^{\sigma^2, \tau^2}).$$

It remains to compare the value of φ on $\text{lab}(\mathcal{P}_u^{\sigma^1, \tau^1})$ and $\text{lab}(\mathcal{P}_u^{\sigma^2, \tau^2})$. We can do this via the oracle from Assumption C.2.