

Perfect Information Stochastic Priority Games

Hugo Gimbert¹ and Wiesław Zielonka^{2*}

¹ LIX, École Polytechnique, Palaiseau, France
gimbert@lix.polytechnique.fr

² LIAFA, Université Paris 7 and CNRS, Paris, France
zielonka@liafa.jussieu.fr

Abstract. We introduce stochastic priority games — a new class of perfect information stochastic games. These games can take two different, but equivalent, forms. In stopping priority games a play can be stopped by the environment after a finite number of stages, however, infinite plays are also possible. In discounted priority games only infinite plays are possible and the payoff is a linear combination of the classical discount payoff and of a limit payoff evaluating the performance at infinity. Shapley games [1] and parity games [2] are special extreme cases of priority games.

1 Introduction

Recently de Alfaro, Henzinger and Majumdar[3] introduced a new variant of μ -calculus: discounted μ -calculus. As it is known since the seminal paper [2] of Emerson and Jutla μ -calculus is strongly related to parity games and this relationship is preserved even for stochastic games, [4]. In this context it is natural to ask if there is a class of games that corresponds to discounted μ -calculus of [3]. A partial answer to this question was given in [5], where an appropriate class of infinite discounted games was introduced. However, in [5], only deterministic systems were considered and much more challenging problem of stochastic games was left open. In the present paper we return to the problem but in the context of perfect information stochastic games. The most basic and usually non-trivial question is if the games that we consider admit “simple” optimal strategies for both players. We give a positive answer, for all games presented in this paper both players have pure stationary optimal strategies. Since our games contain parity games as a very special case, our paper extends the result known for perfect information parity games [6–9].

However, we have an objective which is larger than just transferring to stochastic games the results known for deterministic systems. Parity games are used (directly or through an associated logic) in verification. Conditions that are verified often do not depend on any finite prefix of the play (take as a typical example a simple condition like “A wins if we visit infinitely often some set X of states”). However, certainly all real systems have a finite life span thus we can

* Partially supported by the french ANR-SETI project AVERISS

ask what is the meaning of infinite games when they are used to examine such systems. Notice that the same question arises in classical game theory [10]. The obvious answer is that the life span is finite but unknown or sufficiently long and thus infinite games are a convenient approximation of finite games. However, what finite games are approximated by parity games? Notice that for the games like mean-payoff games that are used in economics the answer is simple: infinite mean-payoff games approximate finite mean-payoff games of long or unknown duration. But we do not see any obvious candidate for “finite parity games”. Suppose that C is a parity condition and f_C a payoff mapping associated with C , i.e. f_C maps to 1 (win) all infinite sequence of states that satisfy C and to 0 all “loosing” sequences. Now we can look for a sequence $f_n, n \in \mathbb{N}$, of payoff functions, such that each f_n , defined for state sequences of length n , gives a payoff for games of length n and such that for each infinite sequence $s_0 s_1 \dots$ of states $f_n(s_0 \dots s_{n-1}) \xrightarrow{n \rightarrow \infty} f_C(s_0 s_1 \dots)$. However, except for very special parity conditions C , such payoff mappings f_n do not exist, thus parity games cannot approximate finite games in the same way as infinite mean-payoff games approximate finite mean-payoff games.

Nevertheless, it turns out that parity games approximate finite games, however “finite” does not mean here that the number of steps is fixed, instead these games are finite in the sense that they stop with probability 1. In Section 4 we present a class of priority stopping games. In the simplest case, when the stopping probabilities are positive for all states, stopping games are stochastic games defined by Shapley [1]. However, we examine also stopping games for which stopping probabilities are positive only for some states. One of the results of this paper can be interpreted in the following way: parity games are a limit of stopping games when the stopping probabilities tend to 0 but all at the same time but rather one after another, in the order determined by priorities.

2 Arenas and Perfect Information Games

Perfect information stochastic games are played by two players, that we call player 1 and player 2. We assume that player $i \in \{1, 2\}$ controls a finite set S_i of states, S_1 and S_2 are disjoint and $S = S_1 \cup S_2$ is the set of all states.

With each state $s \in S$ is associated a finite non-empty set A_s of actions that are available at s and we set $A = \cup_{s \in S} A_s$ to be the set of all actions.

If the current state is $s \in S_i$ then player i controlling this state chooses an available action $a \in A_s$ and, with a probability $p(s'|s, a)$, the systems changes its state to $s' \in S$. Thus $p(\cdot|s, a), s \in S, a \in A_s$, are transition probabilities satisfying the usual conditions: $0 \leq p(s'|s, a) \leq 1$ and $\sum_{s' \in S} p(s'|s, a) = 1$.

Let \mathcal{H}^ω be the set of *histories*, i.e. the set of all infinite sequences $s_0 a_0 s_1 a_1 s_2 \dots$ alternating states and actions. Assuming that the sets S and A are equipped with the discrete topology, we equip \mathcal{H}^ω with the product topology, i.e. the smallest topology for which the mappings

$$\mathbf{S}_i : \mathcal{H}^\omega \rightarrow S, \quad \mathbf{S}_i : \mathcal{H}^\omega \ni s_0 a_0 \dots s_i a_i \dots \mapsto s_i$$

and

$$\mathbf{A}_i : \mathcal{H}^\omega \rightarrow A, \quad \mathbf{A}_i : \mathcal{H}^\omega \ni s_0 a_0 \dots s_i a_i \dots \mapsto a_i$$

are continuous. Thus $(\mathbf{S}_i)_{i \in \mathbb{N}}$ and $(\mathbf{A}_i)_{i \in \mathbb{N}}$, are stochastic processes on the probability space $(\mathcal{H}^\omega, \mathcal{B})$, where \mathcal{B} is Borel σ -algebra generated by open subsets of \mathcal{H}^ω .

The data consisting of the state sets S_1, S_2 , available actions $(A_s)_{s \in S}$ and transition probabilities $p(\cdot, s, a)$ is an *arena* \mathcal{A} .

Let $u : \mathcal{H}^\omega \rightarrow \mathbb{R}$ be a bounded Borel measurable mapping. We interpret $u(h), h \in \mathcal{H}^\omega$, as the payoff obtained by player 1 from player 2 after an infinite play h .

A couple (\mathcal{A}, u) consisting of an arena and a payoff mapping is a *perfect information stochastic game*.

Let $\mathcal{H}_i^+ = (SA)^* S_i, i \in \{1, 2\}$, be the set of finite non-empty histories terminating at a state controlled by player i . A *strategy* for player i is a family of conditional probabilities $\sigma(a|h_n)$ for all $h_n = s_0 a_0 \dots s_n \in \mathcal{H}_i^+$ and $a \in A_{a_n}$. Intuitively, $\sigma(a|s_0 a_0 \dots s_n)$ gives the probability that player i controlling the last state s_n chooses an (available) action a , while the sequence h_n describes the first n steps of the game. As usual $0 \leq \sigma(a|s_0 a_0 \dots s_n) \leq 1$ and $\sum_{a \in A_{s_n}} \sigma(a|s_0 a_0 \dots s_n) = 1$.

A strategy σ is said to be *pure* if for each finite history $h_n = s_0 a_0 \dots s_n \in \mathcal{H}_1^+$ there is an action $a \in A_{s_n}$ such that $\sigma(a|h_n) = 1$, i.e. no randomization is used to choose an action to execute. A strategy σ is *stationary* if for each finite history $h_n = s_0 a_0 \dots s_n \in \mathcal{H}_1^+$, $\sigma(\cdot|h_n) = \sigma(\cdot|s_n)$, i.e. the probability distribution used to choose actions depends only on the last state.

Notice that pure stationary strategies for player i can be identified with mappings $\sigma : S_i \rightarrow A$ such that $\sigma(s) \in A_s$ for $s \in S_i$.

In the sequel we shall use σ , possibly with subscripts or superscripts, to denote a strategy of player 1. On the other hand, τ will always denote a strategy of player 2.

Given and initial state s , strategies σ, τ of both players determine a unique probability measure $\mathbb{P}_{\sigma, \tau}^s$ on $(\mathcal{H}^\omega, \mathcal{B})$, [11].

The expectation corresponding to the probability measure $\mathbb{P}_{\sigma, \tau}^s$ is denoted $\mathbb{E}_{\sigma, \tau}^s$. Thus $\mathbb{E}_{\sigma, \tau}^s(u)$ gives the expected payoff obtained by player 1 from player 2 in the game (\mathcal{A}, u) starting at state s when the players use strategies σ, τ respectively. If $\sup_\sigma \inf_\tau \mathbb{E}_{\sigma, \tau}^s(u) = \inf_\tau \sup_\sigma \mathbb{E}_{\sigma, \tau}^s(u)$ for each state s then the quantity appearing on both side of this equality is *the value of the game* (for initial state s) and is denoted $\text{val}_s(\mathcal{A}, u)$.

Strategies σ^\sharp and τ^\sharp of players 1, 2 are *optimal* in the game (\mathcal{A}, u) if for each state $s \in S$ and for all strategies $\sigma \in \Sigma, \tau \in \mathcal{T}$

$$\mathbb{E}_{\sigma, \tau^\sharp}^s[u] \leq \mathbb{E}_{\sigma^\sharp, \tau^\sharp}^s[u] \leq \mathbb{E}_{\sigma^\sharp, \tau}^s[u] .$$

If σ^\sharp and τ^\sharp are optimal strategies then $\text{val}_s(\mathcal{A}, u) = \mathbb{E}_{\sigma^\sharp, \tau^\sharp}^s[u]$, i.e. the expected payoff obtained when both players use optimal strategies is equal to the value of the game.

3 Priority Games

Starting from this moment we assume that each arena \mathcal{A} is equipped with a *priority mapping*

$$\varphi : S \rightarrow \{1, \dots, k\} \quad (1)$$

from the set S of states to the set $\{1, \dots, k\}$ of (positive integer) *priorities*. The composition

$$\varphi_n = \varphi \circ \mathbf{S}_n, \quad , n \in \mathbb{N} , \quad (2)$$

$\varphi_n : \mathcal{H}^\omega \rightarrow \{1, \dots, k\}$, gives therefore a stochastic process with values in $\{1, \dots, k\}$. Then $\liminf_i \varphi_i$ is a random variable

$$\mathcal{H}^\omega \ni h \mapsto \liminf_i \varphi_i(h)$$

giving for each infinite history $h \in \mathcal{H}^\omega$ its priority which the smallest priority visited infinitely often in h (we assume that $\{1, \dots, k\}$ is equipped with the usual order on integers and \liminf is taken for this order). From this moment onward, we assume that there is a fixed a *reward mapping*

$$r : \{1, \dots, k\} \rightarrow [0, 1] \quad (3)$$

from priorities to the interval $[0, 1]$.

The *priority payoff mapping* $u : \mathcal{H}^\omega \rightarrow [0, 1]$ is defined as

$$u(h) = r(\liminf_i \varphi_i(h)), \quad h \in \mathcal{H}^\omega . \quad (4)$$

Thus, in the priority game (\mathcal{A}, u) , the payoff received by player 1 from player 2 is the reward corresponding to the minimal priority visited infinitely often. If r maps odd priorities to 1 and even priorities to 0 then we get a parity game.

One of the referees drew our attention to the paper of McIver and Morgan [12] where a new stochastic game is introduced and it is proved that this game admits pure stationary optimal strategies. The framework used in [12] is so different from the one used in our paper that the direct comparison is difficult. However, it seems that, after an appropriate translation, the game of [12] corresponds to the priority game defined above thus [12] gives a direct proof that priority games admit pure stationary optimal strategies through a reduction to parity games. In our paper we define, in Sections 4 and 5, a much larger class of games, that includes priority games as a special case and not only we prove that all these games admit pure stationary optimal strategies but we show also that all these games can be seen as appropriate limits of classical discounted games of Shapley [1].

4 Stopping Priority Games

In the sequel we assume that besides the priority and reward mappings (1) and (3) we have also a mapping

$$\lambda : \{1, \dots, k\} \rightarrow [0, 1] \quad (5)$$

from priorities to the interval $[0, 1]$.

We modify the rules of the priority game of Section 3 in the following way.

Every time a state s is visited the game can stop with probability $1 - \lambda(\varphi(s))$, where $\varphi(s)$ is the priority of s . If the game stops at s then player 1 receives from player 2 the payoff $r(\varphi(s))$. If the game does not stop then the player controlling s chooses an action $a \in A_s$ and we go to a state t with probability $p(t|s, a)$. (Thus $p(t|s, a)$ should now be interpreted as the probability to go to t *under the condition that the game does not stop*.) The rules above determine the payoff in the case when the game stops at some state s . However, λ can be 1 for some states (priorities) and then it is possible to have also infinite plays with a positive probability. For such infinite plays the payoff is calculated as in priority games of the preceding section.

Let us note that if $\lambda(p) = 1$ for all priorities $p \in \{1, \dots, k\}$ then actually we never stop and the game described above is the same as the priority game of the preceding section.

On the other hand, if $\lambda(p) < 1$ for all priorities p , i.e. the stopping probabilities are positive for all states, then the game will stop with probability 1. Shapley [1] proved that for such games both players have optimal stationary strategies. In fact Shapley considered general stochastic games while we limit ourselves to perfect information stochastic games and for such games the optimal strategies constructed in [1] are not only stationary but also pure.

Theorem 1 (Shapley 1953). *If, for all priorities i , $\lambda(i) < 1$ then both players have pure stationary optimal strategies in the priority stopping game.*

Stopping games have an appealing intuitive interpretation but they are not consistent with the framework fixed in Section 2, where the probability space consisted of infinite histories only. This obstacle can be removed in the following way. For each priority $i \in \{1, \dots, k\}$ we create a new “stopping” state i^\sharp that we add to the arena \mathcal{A} . The priority of i^\sharp is set to i , $\varphi(i^\sharp) = i$. The set of newly created states is denoted S^\sharp . There is only one action available at each i^\sharp and executing this action we return immediately to i^\sharp with probability 1, it is impossible to leave a stopping state. Note also that since there is only one action available at i^\sharp it does not matter which of the two players controls “stopping” states. For each non-stopping state $s \in S$ we modify the transition probabilities. Formally we define new transition probabilities $p^\sharp(\cdot|\cdot, \cdot)$ by setting, for $s, t \in S$, $a \in A_s$,

$$p^\sharp(t|s, a) = \lambda(\varphi(s)) \cdot p(t|s, a)$$

and

$$p^\sharp(i^\sharp|s, a) = \begin{cases} 1 - \lambda(\varphi(s)) & \text{if } i = \varphi(s), \\ 0 & \text{otherwise.} \end{cases}$$

Let us note by $\mathcal{A}_\lambda^\sharp$ the arena obtained from \mathcal{A} in this way. It is worth noticing that, even if the set of finite histories of $\mathcal{A}_\lambda^\sharp$ strictly contains the set of finite histories of \mathcal{A} , we can identify the strategies in both arenas. In fact, given a strategy for arena \mathcal{A} there is only one possible way to complete it to a strategy

in $\mathcal{A}_\lambda^\sharp$ since for finite histories in $\mathcal{A}_\lambda^\sharp$ that end in a stopping state i^\sharp any strategy chooses always the unique action available at i^\sharp . Clearly, playing a stopping priority game on \mathcal{A} is the same as playing priority game on $\mathcal{A}_\lambda^\sharp$: stopping at state s in \mathcal{A} yields the same payoff as an infinite history in $\mathcal{A}_\lambda^\sharp$ that loops at i^\sharp , where $i = \varphi(s)$.

5 Discounted Priority Games

The aim of this section is to introduce a new class of infinite games that are equivalent to stopping priority games.

As previously, we suppose that arenas are equipped with a priority mapping (1) and that a reward mapping (3) is fixed.

On the other hand, the mapping λ of (5), although also present, has now another interpretation, it does not define stopping probabilities but it provides discount factors applied to one-step rewards.

Let

$$\mathbf{r}_i = r \circ \varphi_i \quad \text{and} \quad \boldsymbol{\lambda}_i = \lambda \circ \varphi_i, \quad i \in \mathbb{N}, \quad (6)$$

be stochastic processes giving respectively the reward and the discount factor at stage i . Then the payoff mapping $u_\lambda : \mathcal{H}^\omega \rightarrow \mathbb{R}$ of *discounted priority games* is defined as

$$u_\lambda = \sum_{i=0}^{\infty} \boldsymbol{\lambda}_0 \cdots \boldsymbol{\lambda}_{i-1} (1 - \boldsymbol{\lambda}_i) \mathbf{r}_i + \left(\prod_{i=0}^{\infty} \boldsymbol{\lambda}_i \right) \cdot r(\liminf_n \varphi_n). \quad (7)$$

Thus u_λ is composed of two parts, the *discount part*

$$u_\lambda^{\text{disc}} = \sum_{i=0}^{\infty} \boldsymbol{\lambda}_0 \cdots \boldsymbol{\lambda}_{i-1} (1 - \boldsymbol{\lambda}_i) \mathbf{r}_i \quad (8)$$

and the *limit part*

$$u_\lambda^{\text{lim}} = \left(\prod_{i=0}^{\infty} \boldsymbol{\lambda}_i \right) \cdot r(\liminf_n \varphi_n). \quad (9)$$

Some remarks concerning this definition are in order. Let

$$T = \inf\{i \mid \boldsymbol{\lambda}_j = 1 \text{ for all } j \geq i\}. \quad (10)$$

Since, by convention, the infimum of the empty set is ∞ , $\{T = \infty\}$ consists of all infinite histories $h \in \mathcal{H}^\omega$ for which $\boldsymbol{\lambda}_i < 1$ for infinitely many i . Thus we can rewrite u_λ as:

$$u_\lambda = \sum_{i < T} \boldsymbol{\lambda}_0 \cdots \boldsymbol{\lambda}_{i-1} (1 - \boldsymbol{\lambda}_i) \mathbf{r}_i + \left(\prod_{i < T} \boldsymbol{\lambda}_i \right) \cdot r(\liminf_n \varphi_n). \quad (11)$$

Moreover, if $T = \infty$ then the product $\prod_{i < T} \boldsymbol{\lambda}_i$, containing infinitely many factors smaller than 1, is equal to 0 and for such infinite histories the limit part u_λ^{lim}

disappears while the discount part is (a sum of) an infinite series. The other extreme case is $T = 0$, i.e. when the discount factor is 1 for all visited states. Then it is the the discount part that disappears from 11 and the payoff is just $r(\liminf_n \varphi_n)$, the same as for priority games of Section 3.

Let (\mathcal{A}, u_λ) be a discounted priority game on \mathcal{A} . As explained in the preceding section, a stopping priority game on \mathcal{A} with stopping probabilities given by means of λ can be identified with the priority game $(\mathcal{A}_\lambda^\sharp, u)$ on the transformed arena $\mathcal{A}_\lambda^\sharp$. As noted also in the preceding section, there is a natural correspondence allowing to identify strategies in both arenas. We shall note by $\mathbb{P}_{\sigma, \tau}^{\sharp s}$ the probability generated by strategies σ and τ on $\mathcal{A}_\lambda^\sharp$ and $\mathbb{P}_{\sigma, \tau}^s$ the similar probability generated by the same strategies on \mathcal{A} . The corresponding expectations are denoted $\mathbb{E}_{\sigma, \tau}^{\sharp s}$ and $\mathbb{E}_{\sigma, \tau}^s$. Having all this facts in mind, the following proposition shows that stopping priority games and discounted priority games are equivalent in the sense that the same strategies yield the same payoffs in both games:

Proposition 2. *For all strategies σ, τ of players 1, 2 and all states $s \in S$, $\mathbb{E}_{\sigma, \tau}^{\sharp s}[u] = \mathbb{E}_{\sigma, \tau}^s[u_\lambda]$.*

Proof. (sketch) Let $T = \inf\{i \mid \mathbf{S}_i \in S^\sharp\}$ be the first moment in the game $(\mathcal{A}_\lambda^\sharp, u)$ when we enter a stopping state. Direct calculations show that $\mathbb{P}_{\sigma, \tau}^{\sharp s}(\mathbf{S}_{i+1} = s_{i+1} \mid \mathbf{S}_0 = s_0, \dots, \mathbf{S}_i = s_i) = \lambda(\varphi(s_i))\mathbb{P}_{\sigma, \tau}^s(\mathbf{S}_{i+1} = s_{i+1} \mid \mathbf{S}_0 = s_0, \dots, \mathbf{S}_i = s_i)$ if all states s_0, \dots, s_i, s_{i+1} are not stopping. This can be used to show that³ $\mathbb{E}_{\sigma, \tau}^{\sharp s}[u; T = \infty] = \mathbb{E}_{\sigma, \tau}^s[u_\lambda^{\text{lim}}]$. On the other hand, $\mathbb{E}_{\sigma, \tau}^{\sharp s}[u; T = m] = \mathbb{E}_{\sigma, \tau}^s[\lambda_0 \cdots \lambda_{m-1}(1 - \lambda_m)r_m]$, implying $\mathbb{E}_{\sigma, \tau}^{\sharp s}[u; T < \infty] = \mathbb{E}_{\sigma, \tau}^s[u_\lambda^{\text{disc}}]$. \square

We can note that in the special case when all discount factors are strictly smaller than 1 (i.e. all stopping probabilities are greater than 0) Proposition 2 reduces to a well-known folklore fact: stopping (Shapley) games[1] and discounted games are equivalent.

6 Limits of Priority Discounted Games

The main aim of this section is to prove that discounted priority games (\mathcal{A}, u_λ) admit pure stationary optimal strategies for both players. Of course, due to Shapley's theorem, we already know that this is true for discounted mappings λ such that $\lambda(i) < 1$ for all priorities i . Our proof will use in an essential way the concept of uniformly optimal strategies, which is of independent interest.

Let $\lambda_1, \dots, \lambda_m, 1 \leq m \leq k$, be a sequence of constants, all belonging to the right-open interval $[0, 1)$. Let λ be the following discount mapping:

$$\text{for all } i \in \{1, \dots, k\}, \quad \lambda(i) = \begin{cases} \lambda_i & \text{if } i \leq m, \\ 1 & \text{if } i > m. \end{cases} \quad (12)$$

³ By $\mathbb{E}_{\sigma, \tau}^{\sharp s}[u; A]$ we denote the integral of u over the set A .

In the sequel we shall write $u_{\lambda_1, \dots, \lambda_m}^{(k)}$ to denote the discounted priority payoff mapping u_λ , where λ is given by (12). (Note, however, that one should not confuse $\lambda_1, \lambda_2, \dots$ which are used to denote real numbers from $[0, 1)$ with bold $\lambda_1, \lambda_2, \dots$ that are used to denote a stochastic process (6)).

Defining $u_{\lambda_1, \dots, \lambda_m}^{(k)}$ we have assumed that $m \geq 1$, however it is convenient to include the case $m = 0$ in this notation: if $m = 0$ then $\lambda(i) = 1$ for all priorities $i \in \{1, \dots, k\}$ and thus $u^{(k)}$ is just the priority payoff mapping u of Section 3. In particular if $m = 1$ then $u_{\lambda_1, \dots, \lambda_{m-1}}^{(k)}$ denotes just $u^{(k)} = u$.

Which strategies are optimal in the game $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_m}^{(k)})$ usually depends heavily on the discount factors $\lambda_1, \dots, \lambda_m$. But, in an important paper [13] Blackwell observed that in discounted Markov decision processes optimal strategies are independent of the discount factor if this factor is close to 1. This leads to the concept of uniformly optimal strategies:

Definition 3. *Let \mathcal{A} be a finite arena, $m \in \{1, \dots, k\}$. Let us fix values of the first $m - 1$ discount factors $\lambda_1, \dots, \lambda_{m-1} \in [0, 1)$. Strategies σ, τ for players 1, 2 are said to be uniformly optimal for $\lambda_1, \dots, \lambda_{m-1}$ if there exists an $\epsilon > 0$ (that can depend on $\lambda_1, \dots, \lambda_{m-1}$) such that σ, τ are optimal for all games $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)})$ with $1 - \epsilon < \lambda_m < 1$.*

Now we are prepared to announce the main result of the paper:

Theorem 4. *For each $m \in \{1, \dots, k\}$ the games $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)})$ admit pure stationary uniformly optimal strategies for both players. Moreover, if $(\sigma^\#, \tau^\#)$ is a pair of such strategies then $\sigma^\#, \tau^\#$ are also optimal in the game $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_{m-1}}^{(k)})$.*

Proposition 5 below establishes the following chain of implications:

if $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_m}^{(k)})$ admits pure stationary optimal strategies then $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_m}^{(k)})$ admits pure stationary *uniformly* optimal strategies which in turn implies that $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_{m-1}}^{(k)})$ admits pure stationary optimal strategies. Since, by Shapley's theorem, $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_k}^{(k)})$ has pure stationary optimal strategies, trivial backward induction on m will yield immediately Theorem 4.

Proposition 5. *Let \mathcal{A} be a finite arena with states labelled by priorities from $\{1, \dots, k\}$. Let $m \in \{1, \dots, k\}$ and $\lambda_1, \dots, \lambda_{m-1}$ be a sequence of discount factors for priorities $1, \dots, m$, all belonging to the interval $[0, 1)$. Suppose that the game $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_m}^{(k)})$ has pure stationary strategies for both players. Then the following conditions hold:*

- (i) *for both players there exist pure stationary uniformly optimal strategies in the game $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)})$,*
- (ii) *there exists an $\epsilon > 0$ such that, for each pair of pure stationary strategies (σ, τ) for players 1 and 2, whenever σ and τ are optimal in the game $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)})$ for some $1 - \epsilon < \lambda_m < 1$ then σ and τ optimal for all games $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)})$ with $1 - \epsilon < \lambda_m < 1$, in particular σ and τ are uniformly optimal,*

- (iii) if σ, τ are pure stationary uniformly optimal strategies in the game $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_m}^{(k)})$ then they are optimal in the game $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_{m-1}}^{(k)})$,
- (iv) $\lim_{\lambda_m \uparrow 1} \text{val}_s(\mathcal{A}, u_{\lambda_1, \dots, \lambda_m}^{(k)}) = \text{val}_s(\mathcal{A}, u_{\lambda_1, \dots, \lambda_{m-1}}^{(k)})$, where $\text{val}_s(\mathcal{A}, u_{\lambda_1, \dots, \lambda_m}^{(k)})$ is the value of the game $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_m}^{(k)})$ for an initial state s .

We precede the proof of Proposition 5 by auxiliary lemmas.

Lemma 6. *Suppose that $\lambda_1, \dots, \lambda_k$, the discount factors for all priorities, are strictly smaller than 1. Let σ, τ be pure stationary strategies for players 1 and 2 in the game $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_k}^{(k)})$. Then the expectation $\mathbb{E}_{\sigma, \tau}^s[u_{\lambda_1, \dots, \lambda_k}^{(k)}]$ is a rational function of $\lambda_1, \dots, \lambda_n$ bounded on $[0, 1]^k$.*

In fact, if we fix pure stationary strategies then we get a finite Markov chain with discounted evaluation. In this context Lemma 6 is standard, at least for one discount factor, see for example [14], and the extension to several discount factors is trivial.

The lack of space compels us to skip the the proof of the following (intuitively obvious) lemma:

Lemma 7. *Let $f(x_1, \dots, x_k)$ be a rational function well-defined and bounded on $[0, 1]^k$. Then, for each $0 \leq m < k$, the iterated limit $\lim_{x_{m+1} \uparrow 1} \dots \lim_{x_k \uparrow 1} f(x_1, \dots, x_k)$ exists and is finite. Moreover, for every fixed $(x_1, \dots, x_{m-1}) \in [0, 1]^{m-1}$ there exists $\epsilon > 0$ such that the one-variable mapping*

$$x_m \mapsto \lim_{x_{m+1} \uparrow 1} \dots \lim_{x_k \uparrow 1} f(x_1, \dots, x_{m-1}, x_m, x_{m+1}, \dots, x_k)$$

is rational on the interval $[1 - \epsilon, 1)$.

For any infinite history $h \in \mathcal{H}^\omega$ the value $u_{\lambda_1, \dots, \lambda_m}^{(k)}(h)$ can be seen as a function of discount factors $\lambda_1, \dots, \lambda_m$. It turns out that

Lemma 8. *For each $m \in \{1, \dots, k\}$ and for each $h \in \mathcal{H}^\omega$,*

$$\lim_{\lambda_m \uparrow 1} u_{\lambda_1, \dots, \lambda_m}^{(k)}(h) = u_{\lambda_1, \dots, \lambda_{m-1}}^{(k)}(h) . \quad (13)$$

Proof. (Sketch) Let $u_{\lambda_1, \dots, \lambda_m}^{(k)} = u_{\lambda_1, \dots, \lambda_m}^{\text{disc}} + u_{\lambda_1, \dots, \lambda_m}^{\text{lim}}$ be the decomposition of $u_{\lambda_1, \dots, \lambda_m}^{(k)}$ onto the discount and limit parts. Let λ and λ^* be discount factor mappings from $\{1, \dots, k\}$ into $[0, 1]$ such that for $i \in \{1, \dots, k\}$, $\lambda(i) = \lambda_i$ for $i \leq m$ and $\lambda(i) = 1$ for $i > m$, while $\lambda^*(i) = \lambda(i)$ for $i \neq m$ and $\lambda^*(i) = 1$ for $i = m$. As usually, $\lambda_i = \lambda \circ \varphi_i$ and $\lambda_i^* = \lambda^* \circ \varphi_i$ are the corresponding stochastic processes.

We examine three cases:

Case 1: $m < \liminf_i \varphi_i(h)$.

In this case, all priorities appearing infinitely often in the sequence $\varphi_i(h)$, $i =$

$0, 1, \dots$ have the corresponding discount factors equal to 1. Thus $T(h) = \min\{j \mid \lambda_l(h) = 1 \text{ for all } l \geq j\}$ is finite. Then $\lim_{\lambda_m \uparrow 1} u_{\lambda_1, \dots, \lambda_m}^{\text{disc}}(h) = u_{\lambda_1, \dots, \lambda_{m-1}}^{\text{disc}}(h)$ since $u_{\lambda_1, \dots, \lambda_m}^{\text{disc}}(h)$ is just a polynomial of variables $\lambda_1, \dots, \lambda_m$. Similarly, $\lim_{\lambda_m \uparrow 1} u_{\lambda_1, \dots, \lambda_m}^{\text{lim}}(h) = u_{\lambda_1, \dots, \lambda_{m-1}}^{\text{lim}}(h)$, and we get (13).

Case 2: $m = \liminf_i \varphi_i(h)$.

Since for infinitely many i , $\lambda_i(h) = \lambda_m < 1$, we have $\prod_{i=0}^{\infty} \lambda_i(h) = 0$, and then $u_{\lambda_1, \dots, \lambda_m}^{\text{lim}}(h) = 0$.

Let $T_0(h) := \max_j \{\varphi_j(h) < m\}$ be the last moment when a priority strictly smaller than m appears in the sequence $\varphi_i(h), i \in \mathbb{N}$, of visited priorities. Notice that $T_0(h) < \infty$ and $\sum_{0 \leq l \leq T_0(h)} \lambda_0(h) \cdots \lambda_{l-1}(h) (1 - \lambda_l(h)) r_l(h) \xrightarrow{\lambda_m \uparrow 1} u_{\lambda_1, \dots, \lambda_{m-1}}^{\text{disc}}(h)$. On the other hand, $\sum_{l=T_0(h)+1}^{\infty} \lambda_0(h) \cdots \lambda_{l-1}(h) (1 - \lambda_l(h)) r_l(h) = (\prod_{j=0}^{T_0(h)} \lambda_j(h)) \cdot \sum_{l=0}^{\infty} (\lambda_m)^l (1 - \lambda_m) r(m) = (\prod_{j=0}^{T_0(h)} \lambda_j(h)) \cdot r(m) \xrightarrow{\lambda_m \uparrow 1} (\prod_{j=0}^{T_0(h)} \lambda_j^*(h)) r(m) = (\prod_{j=0}^{\infty} \lambda_j^*(h)) r(\liminf_i \varphi_i(h)) = u_{\lambda_1, \dots, \lambda_{m-1}}^{\text{lim}}(h)$. Thus we have shown that $\lim_{\lambda_m \uparrow 1} u_{\lambda_1, \dots, \lambda_m}^{\text{disc}}(h) = u_{\lambda_1, \dots, \lambda_{m-1}}^{\text{disc}}(h) + u_{\lambda_1, \dots, \lambda_{m-1}}^{\text{lim}}(h)$.

Case 3: $m > \liminf_i \varphi_i(h)$.

Since $m > m - 1 \geq \liminf_i \varphi_i(h)$ both $u_{\lambda_1, \dots, \lambda_m}^{\text{lim}}(h)$ and $u_{\lambda_1, \dots, \lambda_{m-1}}^{\text{lim}}(h)$ are equal to 0. Thus it suffices to show that

$$\lim_{\lambda_m \uparrow 1} u_{\lambda_1, \dots, \lambda_m}^{\text{disc}}(h) = u_{\lambda_1, \dots, \lambda_{m-1}}^{\text{disc}}(h) . \quad (14)$$

For a subset Z of \mathbb{N} set $f_Z(\lambda_1, \dots, \lambda_m) = \sum_{i \in Z} (1 - \lambda_i(h)) \lambda_0(h) \cdots \lambda_{i-1}(h) r_i(h)$ and consider $f_X(\lambda_1, \dots, \lambda_m)$ and $f_Y(\lambda_1, \dots, \lambda_m)$, where $X = \{i \mid \varphi_i(h) = m\}$ and $Y = \mathbb{N} \setminus X$. We show that

$$\lim_{\lambda_m \uparrow 1} f_X(\lambda_1, \dots, \lambda_m) = 0 . \quad (15)$$

This is obvious if X is finite, thus assume X infinite. Define a process T_i : $T_0(h) = -1$, $T_{i+1}(h) = \min\{j \mid j > T_i(h) \text{ and } \varphi_j(h) = m\}$. Thus $T_i(h), i = 1, 2, \dots$, gives the time of the i -th visit to a state with priority m . Set $p(h) = \liminf_i \varphi_i(h)$ and define another process⁴: $W_i(h) = \sum_{j=0}^{T_i(h)-1} \mathbf{1}_{\{\varphi_j(h)=p(h)\}}$. Thus $W_i(h)$ gives the number states with priority $p(h)$ that were visited prior to the moment $T_i(h)$. Notice that, for all $i \geq 1$, $\lambda_0(h) \cdots \lambda_{T_i(h)-1}$ contains $i - 1$ factors λ_m and $W_i(h)$ factors $\lambda_{p(h)}$ (and possibly other discount factors) whence $\lambda_0(h) \cdots \lambda_{T_i(h)-1} \leq (\lambda_m)^{i-1} (\lambda_{p(h)})^{W_i(h)}$ implying $f_X(\lambda_1, \dots, \lambda_m) = (1 - \lambda_m) r(m) \sum_{i=0}^{\infty} \lambda_0(h) \cdots \lambda_{T_i(h)-1}(h) \leq (1 - \lambda_m) r(m) \sum_{i=0}^{\infty} (\lambda_{p(h)})^{W_{i+1}(h)} (\lambda_m)^{i-1}$. Now notice that $\lim_{i \rightarrow \infty} W_i(h) = \infty$ since $p(h)$ is visited infinitely often in h . Since $p(h) < m$, we have $\lambda_{p(h)} < 1$ and $\lim_{i \rightarrow \infty} (\lambda_{p(h)})^{W_{i+1}(h)} = 0$. This implies easily (15) (applying the well-know and easy fact that summable series are Abel summable, [15]).

Now let us examine $f_Y(\lambda_1, \dots, \lambda_m)$.

Note that $f_Y(\lambda_1, \dots, \lambda_{m-1}, 1) = u_{\lambda_1, \dots, \lambda_{m-1}}^{\text{disc}}(h)$. Then $\lim_{\lambda_m \uparrow 1} f_Y(\lambda_1, \dots, \lambda_m) =$

⁴ We use the usual notation, $\mathbf{1}_A$ is the indicator function of an event A , $\mathbf{1}_A(h) = 1$ if $\mathcal{H}^\omega \ni h \in A$ and $\mathbf{1}_A(h) = 0$ otherwise.

$f_Y(\lambda_1, \dots, \lambda_{m-1}, 1)$ follows directly from the well-know Abel's lemma for power series, see [15]. This and (15) yield (14). \square

Proof of Proposition 5. Since the payoff mappings $u_{\lambda_1, \dots, \lambda_{i+1}}^{(k)}$ are bounded and Borel-measurable, Lebesgue's dominated convergence theorem and Lemma 8 imply that for all strategies σ and τ for players 1 and 2, $\lim_{\lambda_{i+1} \uparrow 1} \mathbb{E}_{\sigma, \tau}^s(u_{\lambda_1, \dots, \lambda_{i+1}}^{(k)}) = \mathbb{E}_{\sigma, \tau}^s(\lim_{\lambda_{i+1} \uparrow 1} u_{\lambda_1, \dots, \lambda_{i+1}}^{(k)}) = \mathbb{E}_{\sigma, \tau}^s(u_{\lambda_1, \dots, \lambda_i}^{(k)})$. Iterating, we get

$$\lim_{\lambda_{m+1} \uparrow 1} \dots \lim_{\lambda_k \uparrow 1} \mathbb{E}_{\sigma, \tau}^s(u_{\lambda_1, \dots, \lambda_k}^{(k)}) = \mathbb{E}_{\sigma, \tau}^s(u_{\lambda_1, \dots, \lambda_m}^{(k)}) . \quad (16)$$

Suppose that strategies σ and τ are pure stationary. Then, by Lemma 6, the mapping $[0, 1]^k \ni (\lambda_1, \dots, \lambda_k) \mapsto \mathbb{E}_{\sigma, \tau}^s(u_{\lambda_1, \dots, \lambda_k}^{(k)})$ is rational and bounded. Lemma 7 applied to the left hand side of (16) allows us to deduce that, for fixed $\lambda_1, \dots, \lambda_{m-1}$, the mapping

$$(0, 1) \ni \lambda_m \mapsto \mathbb{E}_{\sigma, \tau}^s(u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)}) \quad (17)$$

is a rational mapping (of λ_m) for λ_m sufficiently close to 1.

For pure stationary strategies σ and $\sigma^\#$ for player 1 and τ , $\tau^\#$ for player 2 and fixed discount factors $\lambda_1, \dots, \lambda_{m-1}$ we consider the mapping

$$[0, 1) \ni \lambda_m \mapsto \Phi_{\sigma^\#, \tau^\#, \sigma, \tau}(\lambda_m) := \mathbb{E}_{\sigma^\#, \tau^\#}^s(u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)}) - \mathbb{E}_{\sigma, \tau}^s(u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)}) .$$

As a difference of rational mappings, all mappings $\Phi_{\sigma^\#, \tau^\#, \sigma, \tau}$ are rational for λ_m sufficiently close to 1. Since rational mappings are continuous and have finitely many zeros, for each $\Phi_{\sigma^\#, \tau^\#, \sigma, \tau}$ we can find $\epsilon > 0$ such that $\Phi_{\sigma^\#, \tau^\#, \sigma, \tau}$ does not change the sign for $1 - \epsilon < \lambda_m < 1$, i.e. $\forall \lambda_m \in (1 - \epsilon, 1)$,

$$\Phi_{\sigma^\#, \tau^\#, \sigma, \tau}(\lambda_m) \geq 0, \quad \text{or} \quad \Phi_{\sigma^\#, \tau^\#, \sigma, \tau}(\lambda_m) = 0, \quad \text{or} \quad \Phi_{\sigma^\#, \tau^\#, \sigma, \tau}(\lambda_m) \leq 0 . \quad (18)$$

Moreover, since there is only a finite number of pure stationary strategies, we can choose in (18) the same ϵ for all mappings $\Phi_{\sigma^\#, \tau^\#, \sigma, \tau}$, where $\sigma, \sigma^\#$ range over pure stationary strategies of player 1 while $\tau, \tau^\#$ range over pure stationary strategies of player 2.

Suppose that $\sigma^\#, \tau^\#$ are optimal pure stationary strategies in the game $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)})$ for some $\lambda_m \in (1 - \epsilon, 1)$. This means that for all strategies σ, τ for both players

$$\mathbb{E}_{\sigma, \tau^\#}^s(u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)}) \leq \mathbb{E}_{\sigma^\#, \tau^\#}^s(u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)}) \leq \mathbb{E}_{\sigma^\#, \tau}^s(u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)}) . \quad (19)$$

For pure stationary strategies σ, τ , Eq. (19) is equivalent with $\Phi_{\sigma^\#, \tau^\#, \sigma, \tau^\#}(\lambda_m) \geq 0$ and $\Phi_{\sigma^\#, \tau^\#, \sigma^\#, \tau}(\lambda_m) \leq 0$. However, if these two inequalities are satisfied for some λ_m in $(1 - \epsilon, 1)$ then they are satisfied for all such λ_m , i.e. (19) holds for all λ_m in $(1 - \epsilon, 1)$ for all all pure stationary strategies σ, τ . (Thus we have proved that $\sigma^\#$ and $\tau^\#$ are optimal for all $\lambda_m \in (1 - \epsilon, 1)$ but only if we restrict ourselves to the class of pure stationary strategies.) But we have assumed that for each

λ_m the game $(\mathcal{A}, u_{\lambda_1, \dots, \lambda_{m-1}, \lambda_m}^{(k)})$ has optimal pure stationary strategies (now we take into account all strategies), and under this assumption it is straightforward to prove that if (19) holds for all pure stationary strategies σ, τ then it holds for all strategies σ, τ , i.e. $\sigma^\#$ and $\tau^\#$ are optimal in the class of all strategies and for all $\lambda_m \in (1 - \epsilon, 1)$. In this way we have proved conditions (i) and (ii) of Proposition 5.

Applying the limit $\lambda_m \uparrow 1$ to (19) and taking into account (16) we get

$$\mathbb{E}_{\sigma, \tau^\#}^s(u_{\lambda_1, \dots, \lambda_{m-1-1}, \lambda_{m-1}}^{(k)}) \leq \mathbb{E}_{\sigma^\#, \tau^\#}^s(u_{\lambda_1, \dots, \lambda_{m-1-1}, \lambda_{m-1}}^{(k)}) \leq \mathbb{E}_{\sigma^\#, \tau}^s(u_{\lambda_1, \dots, \lambda_{m-1\epsilon-1}, \lambda_{m-1\epsilon}}^{(k)}) ,$$

which proves (iii). It is obvious that this implies also (iv). \square

References

1. Shapley, L.S.: Stochastic games. *Proceedings Nat. Acad. of Science USA* **39** (1953) 1095–1100
2. Emerson, E., Jutla, C.: Tree automata, μ -calculus and determinacy. In: *FOCS'91*, IEEE Computer Society Press (1991) 368–377
3. de Alfaro, L., Henzinger, T.A., Majumdar, R.: Discounting the future in systems theory. In: *ICALP 2003*. Volume 2719 of LNCS., Springer (2003) 1022–1037
4. de Alfaro, L., Majumdar, R.: Quantitative solution to omega-regular games. *Journal of Computer and System Sciences* **68** (2004) 374–397
5. Gimbert, H., Zielonka, W.: Deterministic priority mean-payoff games as limits of discounted games. In: *ICALP 2006*. Volume 4052, part II of LNCS., Springer (2006) 312–323
6. Chatterjee, K., Jurdziński, M., Henzinger, T.: Quantitative stochastic parity games. In: *Proceedings of the 15th Annual Symposium on Discrete Algorithms SODA*. (2004) 114–123
7. McIver, A., Morgan, C.: Games, probability and the quantitative μ -calculus qmu. In: *Proc. LPAR*. Volume 2514 of LNAI., Springer (2002) 292–310 full version arxiv.org/abs/cs.LG/0309024.
8. de Alfaro, L.: Formal Verification of Probabilistic Systems. PhD thesis, Stanford University (december 1997)
9. Zielonka, W.: Perfect-information stochastic parity games. In: *FOSSACS 2004*. Volume 2987 of LNCS., Springer (2004) 499–513
10. Osborne, M.J., Rubinstein, A.: *A Course in Game Theory*. The MIT Press (2002)
11. Filar, J., Vrieze, K.: *Competitive Markov Decision Processes*. Springer (1997)
12. McIver, A., Morgan, C.: A novel stochastic game via the quantitative mu-calculus. In Cerone, A., Wiklicky, H., eds.: *Proc. of the Third Workshop on Quantitative Aspects of Programming Languages (QAPL 2005)*. Volume 153, Issue 2 of ENTCS., Elsevier (2005) 195–212
13. Blackwell, D.: Discrete dynamic programming. *Annals of Mathematical Statistics* **33** (1962) 719–726
14. Hordijk, A., Yushkevich, A.: Blackwell optimality. In Feinberg, E., Schwartz, A., eds.: *Handbook of Markov Decision Processes*. Kluwer (2002)
15. Titchmarsh, E.: *The Theory of Functions*. Oxford University Press second edition.