OPTIMAL STRATEGIES IN PERFECT-INFORMATION STOCHASTIC GAMES WITH TAIL WINNING CONDITIONS

HUGO GIMBERT AND FLORIAN HORN

LaBRI, CNRS, Bordeaux, France *e-mail address*: hugo.gimbert@labri.fr

LIAFA, Université Paris 7, Paris, France *e-mail address*: florian.horn@liafa.jussieu.fr

ABSTRACT. We prove that optimal strategies exist in perfect-information stochastic games with finitely many states and actions and tail winning conditions.

INTRODUCTION

We prove that optimal strategies exist in perfect-information stochastic games with finitely many states and actions and tail winning conditions.

This proof is different from the algorithmic proof sketched in [Hor08].

1. Perfect-Information Stochastic Games

In this section we give formal definitions of perfect-information stochastic games, values and optimal strategies.

1.1. Games, plays and strategies. A (perfect-information stochastic) game is a tuple $(V, V_{\text{Max}}, V_{\min}, V_{\text{R}}, E, W, p)$, where (V, E) is a finite graph, $(V_{\text{Max}}, V_{\min}, V_{\text{R}})$ is a partition of $V, W \subseteq V^{\omega}$ is a measurable set called the winning condition and for every $v \in V_{\text{R}}$ and $w \in V$, $p(w|v) \ge 0$ is the transition probability from v to w, with the property $\sum_{w \in V} p(w|v) = 1$.

A play is an infinite sequence $v_0v_1 \dots \in V^{\omega}$ of vertices such that if $v_n \in (V_{\text{Max}} \cup V_{\min})$ then $(v_n, v_{n+1}) \in E$ and if $v_n \in V_{\text{R}}$ then $p(v_{n+1}|v_n) > 0$. A play is won by Max if it belongs to W otherwise the play is won by Min. A finite play is a finite prefix of a play.

A strategy for player Max is a mapping $\sigma : V^*V_{\text{Max}} \to V$ such that for each finite play $h = v_0 \dots v_n$ such that $v_n \in V_{\text{Max}}$, we have $(v_n, \sigma(h)) \in E$. A play $v_0 v_1 \dots$ is consistent with σ if for every n, if $v_n \in V_{\text{Max}}$ then v_{n+1} is $\sigma(v_0 \dots v_n)$. A strategy for player Min is defined similarly, and is generally denoted τ .

Once the initial vertex v and two strategies σ, τ for player Max and Min are fixed, we can measure the probability that a given set of plays occurs. This probability measure is denoted

²⁰⁰⁰ ACM Subject Classification: Games, Stochastic Processes.

Key words and phrases: perfect-information stochastic games, optimal strategies.

This research was partially supported by french project ANR "DOTS".

GIMBERT AND HORN

 $\mathbb{P}_{v}^{\sigma,\tau}$. For every $n \in \mathbb{N}$, we denote by V_n the random variable defined by $V_n(v_0v_1\cdots) = v_n$, the set of plays is equipped with the σ -algebra generated by random variables $(V_n)_{n\in\mathbb{N}}$. Then there exists a probability measure $\mathbb{P}_{v}^{\sigma,\tau}$ with the following properties:

Π

$$v^{\sigma,\tau}_{v}(V_0 = v) = 1$$
 (1.1)

$$\mathbb{P}_{v}^{\sigma,\tau}\left(V_{n+1} = \sigma(V_0 \cdots V_n) \mid V_n \in V_{\mathrm{Max}}\right) = 1 \quad , \tag{1.2}$$

$$\mathbb{P}_{v}^{\sigma,\tau}\left(V_{n+1} = \tau(V_{0}\cdots V_{n}) \mid V_{n} \in V_{\min}\right) = 1 \quad , \tag{1.3}$$

$$\mathbb{P}_{v}^{\sigma,\tau}\left(V_{n+1} \mid V_{n} \in V_{\mathrm{R}}\right) = p(V_{n+1}|V_{n}) \quad . \tag{1.4}$$

Expectation of a real-valued, measurable and bounded function ϕ under $\mathbb{P}_v^{\sigma,\tau}$ is denoted $\mathbb{E}_v^{\sigma,\tau}[\phi]$. For an event $W \subseteq V^{\omega}$, we denote $\mathbb{1}_W$ the indicator function of W. We will often use implicitely the following formula, which gives the expectation of ϕ once a finite prefix $h = v_0 v_1 \cdots v_n$ of the play is fixed:

$$\mathbb{E}_{v}^{\sigma,\tau}\left[\phi \mid V_{0}\cdots V_{n}=h\right] = \mathbb{E}_{v_{n}}^{\sigma[h],\tau[h]}\left[\phi[h]\right] \quad , \tag{1.5}$$

where $\sigma[h](w_0w_1w_2\cdots) = \sigma(v_0\cdots v_nw_1w_2\cdots)$ and $\tau[h]$ and $\phi[h]$ are defined similarly.

1.2. Values. The goal of player Max is to satisfy the winning condition with the highest probability possible, whereas player Min has the opposite goal. Given a starting vertex v and a strategy σ for player Max, whatever strategy τ is chosen by Min, the play will be won with probability at least:

$$\inf_{\tau} \mathbb{P}_{v}^{\sigma,\tau}\left(W\right)$$

Thus, starting from v, player Max can ensure winning the game with probability arbitrarily close to:

$$\operatorname{val}_{*}(v) = \sup \inf_{\tau} \mathbb{P}_{v}^{\sigma,\tau}(W)$$

and symmetrically, player Min can ensure the play is not won with probability much higher than:

$$\operatorname{val}^{*}(v) = \inf_{\tau} \sup_{\sigma} \mathbb{P}_{v}^{\sigma,\tau}(W) \quad .$$

Clearly $\operatorname{val}_*(v) \leq \operatorname{val}^*(v)$. According to Martin's theorem [Mar98] these values are equal, and this common value is called the value of vertex v and denoted $\operatorname{val}(v)$

1.3. Optimal and ϵ -optimal strategies. By definition of the value, for each $\epsilon > 0$ there exist ϵ -optimal strategies σ_{ϵ} for player Max and τ_{ϵ} for player Min such that for every vertex v,

$$\inf_{\tau} \mathbb{P}_{v}^{\sigma_{\epsilon},\tau}(W) \geq \operatorname{val}(v) - \epsilon \;\;,$$

and symmetrically for player 2,

$$\sup_{\sigma} \mathbb{P}_{v}^{\sigma, \tau_{\epsilon}}(W) \leq \operatorname{val}(v) + \epsilon \;\;.$$

For several classes of winning conditions, it is known that there exists *optimal strategies*, i.e. strategies that are ϵ -optimal for every ϵ .

In this paper, we prove that optimal strategies exist in games whose winning condition has the following property. **Definition 1.1.** A winning condition $W \subseteq V^{\omega}$ is a *tail winning condition* if for every finite play $p \in V^*$ and infinite play $q \in V^{\omega}$,

$$(q \in W) \iff (pq \in W)$$

Games with tail winning conditions have the following properties.

Lemma 1.2. Let G be a game with a tail winning condition W. Then for every vertex $v \in V$,

 $\begin{cases} \operatorname{val}(v) = \max_{(v,w)\in E} \operatorname{val}(w) & \text{if } v \in V_{Max} \\ \operatorname{val}(v) = \min_{(v,w)\in E} \operatorname{val}(w) & \text{if } v \in V_{min} \\ \operatorname{val}(v) = \sum_{(v,w)\in E} p(w|v) \operatorname{val}(w) & \text{if } v \in V_R \end{cases}.$

Proof. This comes from (1.5), and the fact that $\mathbb{1}_W[h] = \mathbb{1}_W$, because W is a tail winning condition.

2. Optimal strategies in games with tail winning conditions

Our main result is:

Theorem 2.1. In every perfect-information stochastic game with tail winning condition and finitely many states and actions, both players have optimal strategies.

The proof of this theorem relies on several intermediary results.

2.1. Consistent games. Next lemma states that it is enough to prove Theorem 2.1 in the case where no move of player Max can decrease the value of a vertex and no move of player Min can increase the value of a vertex.

Lemma 2.2. Let G be a game with a tail winning condition W. We say an edge (v, w) is superfluous when either $v \in V_{Max}$ and $\operatorname{val}_G(w) < \operatorname{val}_G(v)$ or $v \in V_{min}$ and $\operatorname{val}_G(w) > \operatorname{val}_G(v)$. Let G' the game obtained from G by removing all superfluous edges. If there are optimal strategies in G' then there are optimal strategies in G as well.

Proof. We prove that there exists optimal strategies in the game G' obtained by removing only one of the superfluous edges, Lemma 2.2 then results from a trivial induction.

Let (v_s, w_s) be the superfluous edge removed. Without loss of generality, suppose $v_s \in V_{\text{Max}}$, and let

$$m = \operatorname{val}_G(v_s) - \operatorname{val}_G(w_s) > 0$$
.

Suppose there exists optimal strategies σ', τ' in G'.

In game G, player Max has more freedom than in game G', and from every vertex v player Max can guarantee the probability to win to be at least $\operatorname{val}_{G'}(v)$, for that player Max can use its strategy σ' for G', which is a strategy in G as well.

We are going to show that this is the best that player Max can expect in G: we are going to build a strategy τ that prevents the probability to win to be greater than $\operatorname{val}_{G'}$. As a consequence, σ' and τ are a couple of optimal strategies in G, which proves the lemma.

The strategy τ is as follows. As long as player Max does not choose the superfluous edge (v_s, w_s) , the play is a play in G' and strategy τ consists in playing like the strategy τ' in G'. If at some moment player Max chooses the superfluous edge (v_s, w_s) then strategy τ forgets the prefix of the play and switches definitively to a $\frac{m}{2}$ -optimal strategy $\tau_{\frac{m}{2}}$ in G.

GIMBERT AND HORN

If subsequently player Max chooses the superfluous edge again, nothing special happens, τ keeps playing accordingly to τ' . Let Superf be the event defined by:

Superf = {
$$\exists n \in \mathbb{N}, (V_n, V_{n+1}) = (v_s, w_s)$$
},

then the definition of τ and m ensures that for any strategy σ and vertex v,

$$\mathbb{P}_{v}^{\sigma,\tau}\left(W \mid \text{Superf}\right) \le \operatorname{val}_{G}(w_{s}) + \frac{m}{2} = \operatorname{val}_{G}(v_{s}) - \frac{m}{2} \quad .$$

$$(2.1)$$

That way we have an upper bound on the probability to win when the plays does go through the superfluous edge. In case the play does *not* go through the superfluous edge, we prove:

$$\mathbb{P}_{v_s}^{\sigma,\tau}(W \mid \neg \operatorname{Superf}) \le \operatorname{val}_{G'}(v_s) \quad . \tag{2.2}$$

For this, we use the following transformation of σ into a strategy σ_s in G'. Strategy σ_s plays similarly to σ as long as strategy σ does not plays the superfluous edge (v_s, w_s) . If after a finite play v_0, \ldots, v_n , with $v_n = v_s$, strategy σ is about to choose the superfluous edge (v_s, w_s) , then σ_s stops playing similarly to σ . Instead, strategy σ_s forgets the past and switches definitively to the strategy σ' optimal in G', in other words for every play p, $\sigma_s(v_0 \cdots v_n p) = \sigma'(p)$. We denote Switch σ the event:

Switch_{$$\sigma$$} = { $\exists n \in \mathbb{N}, V_n = v_s$ and $\sigma(V_0, \dots, V_n) = w_s$ }

Then by definition of σ_s , for every strategy σ and vertex v,

$$\mathbb{P}_{v}^{\sigma_{s},\tau}\left(W\mid\neg\operatorname{Switch}_{\sigma}\right) = \mathbb{P}_{v}^{\sigma,\tau}\left(W\mid\neg\operatorname{Superf}\right)$$
(2.3)

$$\mathbb{P}_{v}^{\sigma_{s},\tau}\left(W \mid \text{Switch}_{\sigma}\right) \ge \operatorname{val}_{G'}(v_{s}) \quad . \tag{2.4}$$

Since σ_s is a strategy in G' then $\mathbb{P}_v^{\sigma_s,\tau}(W) = \mathbb{P}_v^{\sigma_s,\tau'}(W) \leq \operatorname{val}(G')(v_s)$ because τ' is optimal in G'. Since $\mathbb{P}_v^{\sigma_s,\tau}(W)$ is a convex combination of $\mathbb{P}_v^{\sigma_s,\tau}(W \mid \neg \operatorname{Switch}_{\sigma})$ and $\mathbb{P}_v^{\sigma_s,\tau}(W \mid \operatorname{Switch}_{\sigma})$ then according to (2.4) it implies that $\mathbb{P}_v^{\sigma_s,\tau}(W \mid \neg \operatorname{Switch}_{\sigma}) \leq \operatorname{val}_{G'}(v_s)$. Together with (2.3) it proves (2.2).

We can now prove that the value of v_s in G and G' are the same:

$$\operatorname{val}_G(v_s) = \operatorname{val}_{G'}(v_s) \quad . \tag{2.5}$$

Indeed, for every strategy σ , $\mathbb{P}_{v}^{\sigma_{s},\tau}(W)$ is a convex combination of $\mathbb{P}_{v}^{\sigma_{s},\tau}(W \mid \text{Superf})$ and $\mathbb{P}_{v}^{\sigma_{s},\tau}(W \mid \neg \text{Superf})$ hence according to (2.1) and (2.2), $\mathbb{P}_{v}^{\sigma_{s},\tau}(W) \leq \max\{\text{val}_{G} v_{s} - \frac{m}{2}, \text{val}_{G'}(v_{s})\}$. Taking the supremum over σ , since m > 0 it proves (2.5).

To conclude we prove that (2.5) holds not only for v_s but for any vertex v. Let v be a vertex, σ be a strategy and σ_s the associated switch strategy. Then, since σ and σ_s coincide when event Superf does not occur,

$$\mathbb{P}_{v}^{\sigma,\tau}(W) = \mathbb{P}_{v}^{\sigma,\tau}(W \wedge \neg \operatorname{Superf}) + \mathbb{P}_{v}^{\sigma,\tau}(W \mid \operatorname{Superf}) \cdot \mathbb{P}_{v}^{\sigma,\tau}(\operatorname{Superf}) \\
= \mathbb{P}_{v}^{\sigma_{s},\tau}(W \wedge \neg \operatorname{Superf}) + \mathbb{P}_{v}^{\sigma,\tau}(W \mid \operatorname{Superf}) \cdot \mathbb{P}_{v}^{\sigma_{s},\tau}(\operatorname{Superf}) \quad .$$
(2.6)

According to (2.1), $\mathbb{P}_{v}^{\sigma,\tau}(W \mid \text{Superf}) \leq \text{val}_{G}(v_{s}) = \text{val}_{G'}(v_{s})$ according to (2.5). By definition of τ and σ_{s} , $\mathbb{P}_{v}^{\sigma_{s},\tau}(W \mid \text{Superf}) = \text{val}_{G'}(v_{s})$ because when the event Superf occurs the play is consistent with optimal strategies σ' and τ' in G'. Finally, $\mathbb{P}_{v}^{\sigma,\tau}(W \mid \text{Superf}) \leq \mathbb{P}_{v}^{\sigma_{s},\tau}(W \mid \text{Superf})$, which together with (2.6) gives $\mathbb{P}_{v}^{\sigma,\tau}(W) \leq \mathbb{P}_{v}^{\sigma_{s},\tau}(W)$. Since σ_{s} is a strategy in G' and τ is optimal in G', $\mathbb{P}_{v}^{\sigma_{s},\tau}(W) \leq \text{val}_{G'}(v)$. Taking the supremum over σ , we get $\text{val}_{G}(v) \leq \text{val}_{G'}(v)$ which achieves the proof.

We say that a game G is consistent when for every edge (v, w), if $v \in V_{\text{Max}} \cup V_{\min}$ then $\operatorname{val}_G(v) = \operatorname{val}_G(w)$. consistent games have the following properties.

Lemma 2.3. Let G be a consistent game with a tail winning condition W. Then for every initial vertex v_0 and strategies σ, τ , and every $n \in \mathbb{N}$,

$$\mathbb{E}_{v_0}^{\sigma,\tau}\left[\operatorname{val}(V_{n+1}) \mid V_0, \dots, V_n\right] = \operatorname{val}(V_n)$$

Proof. Comes from lemma 1.2 and the fact that the game is consistent.

2.2. **Deviations.** To detect bad behaviours of a strategy, we use the notions of quality and deviations.

The quality of a strategy σ after a finite play is

$$h_{\sigma}(v_0,\ldots,v_n) = \inf \mathbb{P}_v^{\sigma,\tau} \left(W \mid V_0 = v_0,\ldots,V_n = v_n \right)$$

A *deviation* occurs when the quality of the strategy drops significantly below the value of the current vertex. Formally, let

$$m = \min_{v \in V} \{ \operatorname{val}(v), \operatorname{val}(v) > 0 \} ,$$

be the smallest strictly positive value¹ of a vertex in G, the deviation date is denoted dev_{σ} and defined by:

$$\operatorname{dev}_{\sigma} = \min\left\{n \mid h_{\sigma}(V_0, \dots, V_n) \le \operatorname{val}(V_n) - \frac{m}{2}\right\}$$

with the convention $\min \emptyset = \infty$.

Next lemma states that when player Max plays ϵ -optimally, with ϵ small enough, deviations occur with probability strictly less than 1.

Lemma 2.4. Let G be a consistent game with a tail winning condition W. Let $\epsilon > 0$ and σ be an ϵ -optimal strategy. For every vertex v and strategy τ ,

$$\mathbb{P}_{v_0}^{\sigma,\tau}\left(\operatorname{dev}_{\sigma}<\infty\right) \le \frac{1+\epsilon}{1+\frac{m}{2}} \quad . \tag{2.7}$$

Proof. We start with proving:

$$\mathbb{E}_{v_0}^{\sigma,\tau}\left[\operatorname{val}\left(V_{\operatorname{dev}_{\sigma}}\right) \cdot \mathbb{1}_{\operatorname{dev}_{\sigma} < \infty}\right] \le \operatorname{val}(v_0) \quad .$$

$$(2.8)$$

For every $n \in \mathbb{N}$ let $\operatorname{dev}_{\sigma}^{(n)} = \min\{n, \operatorname{dev}_{\sigma}\}$. According to lemma 2.3, $\mathbb{E}_{v_0}^{\sigma,\tau} \left[\operatorname{val}\left(V_{\operatorname{dev}_{\sigma}^{(n)}}\right)\right] = \operatorname{val}(v_0)$ hence $\mathbb{E}_{v_0}^{\sigma,\tau} \left[\operatorname{val}\left(V_{\operatorname{dev}_{\sigma}^{(n)}}\right) \cdot \mathbb{1}_{\operatorname{dev}_{\sigma} < n}\right] \leq \operatorname{val}(v_0)$. Taking the limit of the left hand-side of this inequality, we obtain (2.8).

Main step in the proof consists in establishing:

$$\mathbb{P}_{v_0}^{\sigma,\tau}\left(W \wedge \operatorname{dev}_{\sigma} < \infty\right) \le \operatorname{val}(v_0) - \frac{m}{2} \cdot \mathbb{P}_{v_0}^{\sigma,\tau}\left(\operatorname{dev}_{\sigma} < \infty\right) \quad . \tag{2.9}$$

For that, we introduce an auxiliary strategy τ' . Let $\epsilon' > 0$. Strategy τ' plays like strategy τ as long as there is no deviation i.e. as long as $h_{\sigma}(v_0, \ldots, v_n) > \operatorname{val}(v_n) - \frac{m}{2}$. In case a deviation occurs i.e. $h_{\sigma}(v_0, \ldots, v_n) \leq \operatorname{val}(v_n) - \frac{m}{2}$ then strategy τ' forgets the past and switches definitively to an ϵ' -optimal response to $\sigma[v_0, \ldots, v_n]$, so that

$$\mathbb{P}_{v_0}^{\sigma,\tau'}(W \mid \operatorname{dev}_{\sigma} = n \text{ and } V_0 \cdots V_n = v_0 \cdots v_n) \le \operatorname{val}(v_n) - \frac{m}{2} + \epsilon' \quad . \tag{2.10}$$

¹if $\forall v \in V$, val(v) = 0 then $m = \infty$ however this case has no interest.

Then,

$$\begin{split} \mathbb{P}_{v_0}^{\sigma,\tau} \left(W \wedge \operatorname{dev}_{\sigma} < \infty \right) &= \mathbb{E}_{v_0}^{\sigma,\tau} \left[\mathbbm{1}_W \cdot \mathbbm{1}_{\operatorname{dev}_{\sigma} < \infty} \right] \\ &= \mathbb{E}_{v_0}^{\sigma,\tau} \left[\mathbb{E}_{v_0}^{\sigma,\tau} \left[\mathbbm{1}_W \cdot \mathbbm{1}_{\operatorname{dev}_{\sigma} < \infty} \mid \operatorname{dev}_{\sigma}, V_0, \dots, V_{\operatorname{dev}_{\sigma}} \right] \right] \\ &= \mathbb{E}_{v_0}^{\sigma,\tau} \left[\mathbb{E}_{v_0}^{\sigma,\tau} \left[\mathbbm{1}_W \mid \operatorname{dev}_{\sigma}, V_0, \dots, V_{\operatorname{dev}_{\sigma}} \right] \cdot \mathbbm{1}_{\operatorname{dev}_{\sigma} < \infty} \right] \\ &\leq \mathbb{E}_{v_0}^{\sigma,\tau} \left[\left(\operatorname{val}(V_{\operatorname{dev}_{\sigma}}) - \frac{m}{2} + \epsilon' \right) \cdot \mathbbm{1}_{\operatorname{dev}_{\sigma} < \infty} \right] \\ &= \mathbb{E}_{v_0}^{\sigma,\tau} \left[\operatorname{val}(V_{\operatorname{dev}_{\sigma}}) \cdot \mathbbm{1}_{\operatorname{dev}_{\sigma} < \infty} \right] + \left(-\frac{m}{2} + \epsilon' \right) \cdot \mathbb{P}_{v_0}^{\sigma,\tau} \left(\operatorname{dev}_{\sigma} < \infty \right) \\ &\leq \operatorname{val}(v_0) + \left(-\frac{m}{2} + \epsilon' \right) \cdot \mathbb{P}_{v_0}^{\sigma,\tau} \left(\operatorname{dev}_{\sigma} < \infty \right) \quad , \end{split}$$

where the three first equalities are properties of conditional expectations, the first inequality is (2.10) and the second inequality is (2.8). Since this holds for every ϵ' , we obtain (2.9) as promised.

Now we can conclude. Since σ is ϵ -optimal,

$$\operatorname{val}(v_0) - \epsilon \leq \mathbb{P}_{v_0}^{\sigma,\tau}(W) = \mathbb{P}_{v_0}^{\sigma,\tau}(W \wedge \operatorname{dev}_{\sigma} < \infty) + \mathbb{P}_{v_0}^{\sigma,\tau}(W \wedge \operatorname{dev}_{\sigma} = \infty)$$
$$\leq \mathbb{P}_{v_0}^{\sigma,\tau}(W \wedge \operatorname{dev}_{\sigma} < \infty) + 1 - \mathbb{P}_{v_0}^{\sigma,\tau}(\operatorname{dev}_{\sigma} < \infty) \quad .$$
(2.11)

Together with (2.9) we obtain (2.7), which achieves the proof of the lemma.

2.3. Construction of an optimal strategy. We can now proceed with the second and last step in the proof of Theorem 2.1. From an ϵ -optimal strategy σ , with ϵ small enough, we construct an optimal strategy, by resetting the memory of σ at right moments. A similar construction has been used in [Cha06] for proving a zero–one law in concurrent games with tail winning conditions.

Lemma 2.5. Let G be a consistent game with a tail winning condition W. Then player Max has an optimal strategy in G.

Proof. If all vertices in G have value 0, there is nothing to prove.

Otherwise, let m be the smallest strictly positive value of a vertex and σ be an $\frac{m}{4}$ optimal strategy. Using σ , we are going to define a strategy σ' and prove that σ' is optimal
in G. For that, we define $t(v_0, \ldots, v_n)$ the date of the latest deviation before date n by $t(v_0) = 0$ and

$$t(v_0, \dots, v_n, v_{n+1}) = \begin{cases} t(v_0, \dots, v_n) & \text{if } h_{\sigma}(v_{t(v_0, \dots, v_n)}, \dots, v_{n+1}) \ge \operatorname{val}(v_{n+1}) - \frac{m}{2} \\ n+1 & \text{otherwise.} \end{cases}$$

By definition the sequence $(t(V_0, \ldots, V_n))_{n \in \mathbb{N}}$ is increasing, we denote T its limit in $\mathbb{N} \cup \{\infty\}$. Strategy σ' consists in forgetting everything before the last deviation and applying σ , i.e.

$$\sigma'(v_0,\ldots,v_n)=\sigma(v_{t(v_0,\ldots,v_n)},\ldots,v_n) \ .$$

To prove that σ' is optimal, we start with proving for every strategy τ and vertex v,

$$\mathbb{P}_v^{\sigma',\tau}\left(T<\infty\right) = 1 \quad . \tag{2.12}$$

Let $D = \min\{n \mid t(V_0, \ldots, V_n) \ge 1\}$ be the date of the first deviation, then since σ and σ' coincide until the first deviation,

$$\mathbb{P}_{v}^{\sigma',\tau}\left(D<\infty\right) = \mathbb{P}_{v}^{\sigma,\tau}\left(D<\infty\right) \quad , \tag{2.13}$$

and by definition of σ' for every $n \in \mathbb{N}$,

$$\mathbb{P}_{v}^{\sigma',\tau}(T=\infty \mid D=n, V_{0}=v_{0}, \dots, V_{n}=v_{n}) = \mathbb{P}_{v_{n}}^{\sigma',\tau[v_{0},\dots,v_{n}]}(T=\infty) \quad .$$
(2.14)

Let $\epsilon > 0$ and τ and v such that:

$$\sup_{\tau',v'} \mathbb{P}_{v'}^{\sigma',\tau'} \left(T = \infty\right) \le \mathbb{P}_{v}^{\sigma',\tau} \left(T = \infty\right) + \epsilon \quad .$$
(2.15)

According to lemma 2.4, since σ is $\frac{m}{4}$ -optimal,

$$\mathbb{P}_{v}^{\sigma,\tau}\left(D<\infty\right) \le \frac{1+\frac{m}{4}}{1+\frac{m}{2}} < 1 \quad . \tag{2.16}$$

By properties of conditional expectations,

$$\begin{split} \mathbb{P}_{v}^{\sigma',\tau}\left(T=\infty\right) &= \mathbb{E}_{v}^{\sigma',\tau}\left[\mathbb{P}_{v}^{\sigma',\tau}\left(T=\infty\mid D, V_{0},\ldots,V_{D}\right)\right] \\ &= \mathbb{E}_{v}^{\sigma',\tau}\left[\mathbbm{1}_{D<\infty}\cdot\mathbb{P}_{v}^{\sigma',\tau}\left(T=\infty\mid D, V_{0},\ldots,V_{D}\right)\right] \\ &= \mathbb{E}_{v}^{\sigma',\tau}\left[\mathbbm{1}_{D<\infty}\cdot\mathbb{P}_{V_{D}}^{\sigma',\tau\left[V_{0},\ldots,V_{D}\right]}\left(T=\infty\right)\right] \\ &\leq \mathbb{E}_{v}^{\sigma',\tau}\left[\mathbbm{1}_{D<\infty}\cdot\left(\mathbb{P}_{v}^{\sigma',\tau}\left(T=\infty\right)+\epsilon\right)\right] \\ &= \mathbb{P}_{v}^{\sigma',\tau}\left(D<\infty\right)\cdot\left(\mathbb{P}_{v}^{\sigma',\tau}\left(T=\infty\right)+\epsilon\right) \\ &= \frac{1+\frac{m}{4}}{1+\frac{m}{2}}\cdot\left(\mathbb{P}_{v}^{\sigma',\tau}\left(T=\infty\right)+\epsilon\right) \ , \end{split}$$

where the second equality is because $\mathbb{P}_{v}^{\sigma,\tau}$ $(D < \infty \mid T = \infty) = 1$, the third equality is (2.14), the inequality is (2.15), and the last equality is (2.13) and (2.16). Since this holds for any ϵ , we obtain $\mathbb{P}_{v}^{\sigma',\tau}$ $(T = \infty) = 0$ i.e. (2.12).

Second step of the proof is to establish:

$$\mathbb{P}_{v_0}^{\sigma',\tau}\left(\operatorname{val}(V_n) \underset{n \to \infty}{\longrightarrow} 0 \mid h_{\sigma'}(V_0, \dots, V_n) \underset{n \to \infty}{\longrightarrow} 0\right) = 1 \quad .$$
(2.17)

When playing with σ' , suppose $h_{\sigma'}(V_0, \ldots, V_n)$ converges to 0 then by definition of σ' , $h_{\sigma}(V_{t(V_0,\ldots,V_n)},\ldots,V_n)$ converges to 0 as well. According to (2.12), $t(V_0,\ldots,V_n)$ has limit $T < \infty$ hence $h_{\sigma}(V_T,\ldots,V_n)$ converges to 0 as well. By definition of T, for every $n \ge T$, $t(V_0,\ldots,v_n) = T$, hence $h_{\sigma}(V_T,\ldots,V_n) \ge \operatorname{val}(V_n) - \frac{m}{2}$. Since $h_{\sigma}(V_T,\ldots,V_n)$ converges to 0, $\limsup_n \operatorname{val}(V_n) \le \frac{m}{2}$. But hence $\operatorname{val}(V_n) \longrightarrow_n 0$ because by definition of m, $(\operatorname{val}(v) < m) \Longrightarrow (\operatorname{val}(v) = 0)$. This proves (2.17).

We can now achieve the proof of the optimality of σ' . Since W is a tail winning condition, Levy's law [Dur96] implies,

$$\mathbb{P}_{v_0}^{\sigma',\tau}(\neg W) = \mathbb{P}_{v_0}^{\sigma',\tau}\left(\mathbb{P}_{v_0}^{\sigma',\tau}(W \mid V_0, \dots, V_n) \xrightarrow{n} 0\right) \\
\leq \mathbb{P}_{v_0}^{\sigma',\tau}\left(h_{\sigma'}(V_0, \dots, V_n) \xrightarrow{n} 0\right) \\
\leq \mathbb{P}_{v_0}^{\sigma',\tau}\left(\operatorname{val}(V_n) \xrightarrow{n} 0\right) \\
\leq \mathbb{E}_{v_0}^{\sigma',\tau}\left[1 - \limsup_{n} \operatorname{val}(V_n)\right] \\
\leq 1 - \limsup_{n} \mathbb{E}_{v_0}^{\sigma',\tau}\left[\operatorname{val}(V_n)\right] \\
= 1 - \operatorname{val}(v_0) ,$$

where the first inequality holds by definition of $h_{\sigma'}(v_0, \ldots, v_n)$, the second is (2.17), the third and fourth are basic properties of expectation and the last equality holds according to lemma 2.3. This proves that σ' is optimal in G.

2.4. Proof of Theorem 2.1. According to lemma 2.2 we can suppose without loss of generality that G is consistent. Since both the winning condition W and its complement $V^{\omega} \setminus W$ are tail winning conditions, lemma 2.5 implies that both players have optimal strategies in G.

CONCLUSION

We have proved the existence of optimal strategies in any perfect-information game with a tail winning condition. We relied heavily on the finiteness of the game, actually the result does not hold in general for infinite arenas. Extension of this result to certain classes of games with partial information or with infinitely many vertices seems to be an interesting research direction.

References

- [Cha06] K. Chatterjee. Concurrent games with tail objectives. In CSL'06, 2006.
- [Dur96] R. Durett. Probability Theory and Examples. Duxbury Press, 1996.
- [Hor08] F. Horn. Random Games. PhD thesis, Université Denis-Diderot and RWTH Aachen, 2008.
- [Mar98] D.A. Martin. The determinacy of Blackwell games. Journal of Symbolic Logic, 63(4):1565–1581, 1998.