



HAL
open science

Constant Payoff Property in Zero-Sum Stochastic Games with a Finite Horizon

Thomas Ragel, Bruno Ziliotto

► **To cite this version:**

Thomas Ragel, Bruno Ziliotto. Constant Payoff Property in Zero-Sum Stochastic Games with a Finite Horizon. 2024. hal-04786399

HAL Id: hal-04786399

<https://hal.science/hal-04786399v1>

Preprint submitted on 15 Nov 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Constant Payoff Property in Zero-Sum Stochastic Games with a Finite Horizon

Thomas Ragel* and Bruno Ziliotto†

Abstract

This paper examines finite zero-sum stochastic games and demonstrates that when the game's duration is sufficiently long, there exists a pair of approximately optimal strategies such that the expected average payoff at any point in the game remains close to the value. This property, known as the *constant payoff property*, was previously established only for absorbing games and discounted stochastic games.

Introduction

Zero-sum stochastic games [8] model dynamic interactions between two adversarial players. At each stage $m \geq 1$, Player 1 selects an action i_m and Player 2 selects an action j_m , possibly using randomization. Player 1 receives a stage payoff $g_m := g(\omega_m, i_m, j_m)$, and Player 2 receives $-g_m$, where ω_m is a random variable called the *state*. The distribution of ω_m depends only on $(\omega_{m-1}, i_{m-1}, j_{m-1})$. Moreover, at the end of each stage m , players are informed of i_m, j_m and ω_{m+1} . In the *n-stage game*, the total payoff is the expectation of the Cesaro mean $\frac{1}{n} \sum_{m=1}^n g_m$, while in the *λ -discounted game*, the total payoff is the expectation of the Abel mean $\sum_{m \geq 1} \lambda(1 - \lambda)^{m-1} g_m$. In both games, Player 1 aims to maximize the total payoff while Player 2 seeks to minimize it. The state space and action sets are assumed to be finite. The value of the *n-stage game* is denoted by v_n , and the value of the *λ -discounted game* is denoted by v_λ . Significant effort has been dedicated to studying the asymptotic behavior of stochastic games as n tends to infinity and λ goes to 0. A seminal result by Bewley and Kohlberg [2] shows that (v_λ) and (v_n) converge to the same limit v^* , called the *limit value*. Recently, Attia and Oliu-Barton [1] provided a characterization of v^* . However, the limit value may not exist when the state is unobserved or when the state space is infinite [15]. Similarly, it may not exist if one of the action sets is infinite [14, 13, 16].

In this paper, we study the *constant payoff property*, which originated in the work [11]. That work shows that, in the context of single decision-maker problems, the uniform convergence of (v_n) implies that, for all $t \in]0, 1[$ and n sufficiently large, when players use ε -optimal strategies in the *n-stage game*, the expectation of $\frac{1}{n} \sum_{m=1}^{\lceil tn \rceil} g_m$ is approximately equal to tv^* , up to a $O(\varepsilon)$ term: the average payoff is *constant*. In particular, it can not happen that one player gets far more than the value during the first $n/2$ stages, and then far less during the last $n/2$ stages. A similar property holds for discounted evaluations: for any $M \in \mathbb{N}$ and sufficiently small λ , if players use ε -optimal strategies, the expectation of the cumulated payoff $\sum_{m=1}^M \lambda(1 - \lambda)^{m-1} g_m$ is

*CEREMADE, Paris Dauphine University, PSL Research Institute, France.

†Institut Mathématique de Toulouse, CNRS, Toulouse School of Economics, France.

close to $\left[\sum_{m=1}^M \lambda(1-\lambda)^{m-1}\right] v^*$, up to a $O(\varepsilon)$. However, [11] also provides an example showing that these results do not generalize to the two-player setting for any pair of ε -optimal strategies. Nevertheless, they propose the following conjecture: in any two-player stochastic game where (v_n) converges uniformly (possibly with infinite state space or action sets), for each $\varepsilon > 0$ and n large enough, *there exists* a pair of ε -optimal strategies that satisfies the constant payoff property.

Regarding finite stochastic games, the conjecture holds due to the existence of a uniform value [5]. However, the strategies built in [5] are notably complex. This complexity has led subsequent research to explore whether the conjecture remains valid for Markovian strategies, which depend only on the stage and current state. The property has been established for absorbing games [12, 7], and for stochastic games with discounted payoffs [3]. This paper adds to this body of work by demonstrating that the constant payoff property also holds for n -stage stochastic games. The proof uses a class of Markovian strategies, called *adapted strategies*, similar to the asymptotically optimal strategies considered in [17]. At each stage m , these strategies play optimally in a discounted game with a discount factor of $\frac{1}{n-m+1}$. This discount factor reflects the importance of stage m relative to the remaining stages of the game. We consider a variation where the discount factor is piecewise constant, updated only at regular intervals. We prove that these adapted strategies are asymptotically optimal in the n -stage game and satisfy the constant payoff property. To establish their asymptotical optimality, we employ the same operator-based approach used in [17]. That technique is also reminiscent of an argument by Neyman, that used the operator approach to prove that if (v_λ) has bounded variation, then (v_n) converges (see [10, Theorem C.8, p.177] and [6, Theorem 4, p.401]).

To prove the constant payoff property, we leverage the fact that by construction, adapted strategies behave locally as optimal strategies in the corresponding discounted game. Utilizing results from [3], we conclude that for these strategies, the payoff is constant within blocks, from which the constant payoff property follows.

The structure of the paper is as follows. In [Section 1](#), we introduce the stochastic game model and state the main results, [Theorems 1.5 and 1.6](#). [Section 2](#) is dedicated to proving [Theorem 1.5](#). In [Section 3](#), we present preliminary results necessary for the proof of [Theorem 1.6](#), which is completed in [Section 4](#). Finally, [Section 5](#) discusses possible generalizations and future research directions.

1 Model and Main Result

1.1 Stochastic Games

Given a finite set A , we denote by $\Delta(A)$ the set of probability distributions over A .

Model A zero-sum stochastic game is described by a tuple $\Gamma = (\Omega, I, J, g, q)$, where Ω is the state space, I is Player 1's action set, J is Player 2's action set, $g : \Omega \times I \times J \rightarrow \mathbb{R}$ is the payoff function, and $q : \Omega \times I \times J \rightarrow \Delta(\Omega)$ is the transition function. We assume that Ω , I and J are finite sets.

The game proceeds as follows: at each stage $m \geq 1$, simultaneously and independently, Player 1 picks $i_m \in I$ and Player 2 picks $j_m \in J$. Player 1 receives the stage payoff $g_m := g(\omega_m, i_m, j_m)$, while Player 2 receives $-g_m$. The next state ω_{m+1} is drawn from the distribution $q(\omega_m, i_m, j_m)$. Players observe (ω_{m+1}, i_m, j_m) , and the game proceeds to the next stage.

Strategies An element of $H_m := (\Omega \times I \times J)^{m-1} \times \Omega$ is called an *history before stage m* . A *strategy* for Player 1 is a collection of mappings $(\sigma_m)_{m \geq 1}$, where $\sigma_m : H_m \rightarrow \Delta(I)$, with the following interpretation: if the history before stage m is $h_m := (\omega_1, i_1, j_1, \dots, \omega_{m-1}, i_{m-1}, j_{m-1}, \omega_m)$, then Player 1 draws i_m according to the distribution $\sigma(h_m)$. Similarly, a strategy for Player 2 is a collection of mappings $(\rho_m)_{m \geq 1}$, where $\rho_m : H_m \rightarrow \Delta(J)$.

A *Markov strategy* is a strategy that plays according to the current state and the current stage only. A Markov strategy for Player 1 can be identified with a mapping $\sigma : \mathbb{N} \times \Omega \rightarrow \Delta(I)$, and a Markov strategy for Player 2 can be identified with a mapping $\rho : \mathbb{N} \times \Omega \rightarrow \Delta(J)$.

A *stationary strategy* is a strategy that plays according to the current state only. A stationary strategy for Player 1 can be identified with a mapping $x : \Omega \rightarrow \Delta(I)$, and a stationary strategy for Player 2 can be identified with a mapping $y : \Omega \rightarrow \Delta(J)$.

n -stage game and λ -discounted game An initial state ω and a pair of strategies (σ, ρ) induce a probability measure on the set of finite histories $\cup_{m \geq 1} H_m$. Thanks to the Kolmogorov extension theorem, it can be extended into a probability measure on $(\Omega \times I \times J)^{\mathbb{N}}$, that is denoted by $\mathbb{P}_{\sigma, \rho}^{\omega}$. The expectation with respect to $\mathbb{P}_{\sigma, \rho}^{\omega}$ is denoted by $\mathbb{E}_{\sigma, \rho}^{\omega}$.

The n -stage game $\Gamma_n(\omega)$ is the normal-form game $(\Sigma, T, \gamma_n^{\omega})$, where γ_n^{ω} is the payoff function defined by

$$\gamma_n^{\omega}(\sigma, \rho) = \mathbb{E}_{\sigma, \rho}^{\omega} \left(\frac{1}{n} \sum_{m=1}^n g_m \right).$$

This game has a value [8], that is denoted by $v_n(\omega)$:

$$v_n(\omega) := \max_{\sigma \in \Sigma} \min_{\rho \in T} \gamma_n^{\omega}(\sigma, \rho) = \min_{\rho \in T} \max_{\sigma \in \Sigma} \gamma_n^{\omega}(\sigma, \rho).$$

A strategy $\sigma \in \Sigma$ of Player 1 is *optimal* (resp., ε -optimal) in $\Gamma_n(\omega)$ if for all $\rho \in T$, $\gamma_n^{\omega}(\sigma, \rho) \geq v_n(\omega)$ (resp., $\gamma_n^{\omega}(\sigma, \rho) \geq v_n(\omega) - \varepsilon$). A strategy $\rho \in T$ of Player 2 is *optimal* (resp., ε -optimal) in $\Gamma_n(\omega)$ if for all $\sigma \in \Sigma$, $\gamma_n^{\omega}(\sigma, \rho) \leq v_n(\omega)$ (resp., $\gamma_n^{\omega}(\sigma, \rho) \leq v_n(\omega) + \varepsilon$). An optimal strategy in Γ_n is a strategy that is optimal in $\Gamma_n(\omega)$ for any ω .

The λ -discounted game is the normal-form game $(\Sigma, T, \gamma_{\lambda}^{\omega})$, where $\gamma_{\lambda}^{\omega}$ is the payoff function defined by

$$\gamma_{\lambda}^{\omega}(\sigma, \rho) = \mathbb{E}_{\sigma, \rho}^{\omega} \left(\frac{1}{n} \sum_{m=1}^n g_m \right).$$

This game has a value [8], that is denoted by $v_{\lambda}(\omega)$:

$$v_{\lambda}(\omega) := \max_{\sigma \in \Sigma} \min_{\rho \in T} \gamma_{\lambda}^{\omega}(\sigma, \rho) = \min_{\rho \in T} \max_{\sigma \in \Sigma} \gamma_{\lambda}^{\omega}(\sigma, \rho).$$

The notion of optimal strategy and ε -optimal strategy can be defined similarly as in the n -stage game.

Let us recall that, using the fact that Ω , I and J are finite, (v_n) and (v_{λ}) both converge (as n tends to $+\infty$ and λ tends to 0) to the same limit v^* , called *limit value* [2].

1.2 Discounted Optimal Profiles, Adaptive Profiles and Main Results

We now introduce two concepts of strategy families that play a crucial role in the paper.

Definition 1.1. A family of stationary strategy pairs $(x_\lambda, y_\lambda)_\lambda$ is a *discounted optimal profile* if for all $\lambda \in]0, 1]$, (x_λ, y_λ) is a pair of optimal strategies in Γ_λ .

Notation. Let $n \geq 1$. Given $a_n \in \mathbb{N}$, let us define,

- for all $m \in \llbracket 1; n \rrbracket$, $k(m) := \left\lfloor \frac{m-1}{a_n} \right\rfloor$,
- $p_n := \left\lfloor \frac{n}{a_n} \right\rfloor$,
- for all $k \in \llbracket 0; p_n - 1 \rrbracket$, $\lambda_k^n := 1/(n - ka_n)$.

Definition 1.2. A family of strategy pairs $(\sigma_n, \rho_n)_n$ is an *adapted profile* if there exists a discounted optimal profile $(x_\lambda, y_\lambda)_\lambda$ and a sequence of positive integers $(a_n)_{n \geq 1}$ such that a_n/n tends to 0 and such that for all $n \geq 1$, the strategy σ_n (resp., ρ_n) plays $x_{\lambda_{k(m)}^n}$ (resp., $y_{\lambda_{k(m)}^n}$) at each stage $m \in \llbracket 1; n \rrbracket$.

Consider an adapted profile $(\sigma_n, \rho_n)_n$ and $n \in \mathbb{N}$. At each stage $m \in \llbracket 1; m \rrbracket$, the strategies σ_n and ρ_n play optimally in a discounted game with discounted factor $\frac{1}{n-k(m)+1}$. The discounted factor is thus updated at the beginning at each block of size a_n , and the integer $n - k(m) + 1$ represents the number of remaining stages at the beginning of the current block.

Definition 1.3. A family of strategy pairs $(\sigma_n, \rho_n)_n$ is an *asymptotically optimal profile* if there exists $\varepsilon_n \rightarrow 0$ such that for all $n \in \mathbb{N}$, the strategies σ_n and ρ_n are ε_n -optimal in Γ_n .

Definition 1.4. A family of strategy pairs $(\sigma_n, \rho_n)_n$ satisfies the *constant payoff property* if for all $t \in]0, 1[$ and $\omega \in \Omega$,

$$\lim_{n \rightarrow +\infty} \mathbb{E}_{\sigma_n, \rho_n}^\omega \left[\frac{1}{n} \sum_{m=1}^{\lceil tn \rceil} g_m \right] = tv^*(\omega).$$

Our main results are the following:

Theorem 1.5. *Any adapted profile is asymptotically optimal.*

Theorem 1.6. *There exists an adapted profile that satisfies the constant payoff property.*

Consequently, there exists an asymptotically optimal profile that satisfies the constant payoff property.

2 Proof of Theorem 1.5

Proof. Let $(\sigma_n, \rho_n)_n$ be an adapted profile, let $(x_\lambda, y_\lambda)_\lambda$ be the corresponding discounted optimal profile and $(a_n)_n$ be the corresponding sequence. Let us recall the following definitions:

$$k(m) := \left\lfloor \frac{m-1}{a_n} \right\rfloor, \quad p_n := \lfloor n/a_n \rfloor \quad \text{and} \quad \lambda_k^n := 1/(n - ka_n).$$

Let us also define $\mu_m^n := \lambda_{k(m)}^n$. Hence, at any stage $m \in \llbracket 1; n \rrbracket$, the strategies σ_n and ρ_n play respectively $x_{\mu_m^n}$ and $y_{\mu_m^n}$.

Define w_m^n as the payoff guaranteed by σ_n in the n -stage game starting from stage m . Let $\omega \in \Omega$. Using Shapley's equations [8], we have

$$w_m^n(\omega) = \min_{j \in J} \left\{ \frac{1}{n-m+1} \sum_{i \in I} x_{\mu_m^n}(i|\omega) g(\omega, i, j) + \left(1 - \frac{1}{n-m+1}\right) \sum_{(i, \omega') \in I \times \Omega} x_{\mu_m^n}(i|\omega) q(\omega'|\omega, i, j) w_{m+1}^n(\omega') \right\},$$

where $x_{\mu_m^n}(i|\omega)$ designates the probability that action i is chosen at state ω , under the stationary strategy $x_{\mu_m^n}$. Similarly, we have

$$v_{\mu_m^n}(\omega) = \min_{j \in J} \left\{ \mu_m^n \sum_{i \in I} x_{\mu_m^n}(i|\omega) g(\omega, i, j) + (1 - \mu_m^n) \sum_{(i, \omega') \in I \times \Omega} x_{\mu_m^n}(i|\omega) q(\omega'|\omega, i, j) v_{\mu_m^n}(\omega') \right\}.$$

Consider two sets of real numbers $A = \{a_j\}_{j \in J}$ and $B = \{b_j\}_{j \in J}$. The following inequality holds:

$$|\min A - \min B| \leq \max_{j \in J} |a_j - b_j|.$$

We deduce that

$$\begin{aligned} \|w_m^n - v_{\mu_m^n}\|_\infty &\leq 2 \|g\|_\infty \left| \frac{1}{n-m+1} - \mu_m^n \right| + \left(1 - \frac{1}{n-m+1}\right) \|w_{m+1}^n - v_{\mu_m^n}\|_\infty \\ &\leq 2 \|g\|_\infty \left| \frac{1}{n-m+1} - \mu_m^n \right| + \left(1 - \frac{1}{n-m+1}\right) \|w_{m+1}^n - v_{\mu_{m+1}^n}\|_\infty + \|v_{\mu_{m+1}^n} - v_{\mu_m^n}\|_\infty. \end{aligned}$$

Multiplying both sides by $\frac{n-m+1}{n}$ yields

$$\begin{aligned} \frac{n-m+1}{n} \|w_m^n - v_{\mu_m^n}\|_\infty - \frac{n-m}{n} \|w_{m+1}^n - v_{\mu_{m+1}^n}\|_\infty \\ \leq \frac{n-m+1}{n} \left(2 \|g\|_\infty \left| \frac{1}{n-m+1} - \mu_m^n \right| + \|v_{\mu_{m+1}^n} - v_{\mu_m^n}\|_\infty \right). \end{aligned}$$

Hence,

$$\frac{n-m+1}{n} \|w_m^n - v_{\mu_m^n}\|_\infty - \frac{n-m}{n} \|w_{m+1}^n - v_{\mu_{m+1}^n}\|_\infty \leq 2 \|g\|_\infty \left| \frac{1}{n-m+1} - \mu_m^n \right| + \|v_{\mu_{m+1}^n} - v_{\mu_m^n}\|_\infty.$$

Fix $\varepsilon > 0$, and define $k_0 := \lfloor (1 - \varepsilon)p_n \rfloor$. Summing the above inequality from $m = 1$ to $k_0 a_n$, we obtain

$$\|w_1^n - v_{\mu_1^n}\|_\infty - \frac{n - k_0 a_n}{n} \|w_{k_0 a_n + 1}^n - v_{\lambda_n}\|_\infty \leq 2 \|g\|_\infty \sum_{m=1}^{k_0 a_n} \left| \frac{1}{n-m+1} - \mu_m^n \right| + \sum_{m=1}^{k_0 a_n} \|v_{\mu_{m+1}^n} - v_{\mu_m^n}\|_\infty.$$

We deduce that

$$\|w_1^n - v_{\lambda_1^n}\|_\infty \leq \left(\varepsilon + \frac{2a_n}{n} \right) \|g\|_\infty + 2 \|g\|_\infty \sum_{m=1}^{k_0 a_n} \left| \frac{1}{n-m+1} - \mu_m^n \right| + \sum_{m=1}^{k_0 a_n} \|v_{\mu_{m+1}^n} - v_{\mu_m^n}\|_\infty. \quad (2.1)$$

Let us bound the two sums, starting with the first one. For $k \in \llbracket 0; k_0 - 1 \rrbracket$ and $m \in \llbracket ka_n + 1; (k + 1)a_n \rrbracket$, we have

$$\begin{aligned} \left| \frac{1}{n - m + 1} - \mu_m^n \right| &= \left| \frac{1}{n - m + 1} - \lambda_k^n \right| \\ &\leq \frac{1}{n - (k + 1)a_n} - \frac{1}{n - ka_n}. \end{aligned}$$

We deduce that

$$\begin{aligned} \sum_{m=1}^{k_0 a_n} \left| \frac{1}{n - m + 1} - \mu_m^n \right| &\leq a_n \sum_{k=0}^{k_0 - 1} \left(\frac{1}{n - (k + 1)a_n} - \frac{1}{n - ka_n} \right) \\ &\leq a_n \frac{1}{n - k_0 a_n} \\ &\leq \frac{a_n}{\varepsilon n}. \end{aligned}$$

Because $a_n/n \rightarrow 0$, we deduce that the second sum vanishes as n tends to infinity.

Let us now bound the second sum in [Section 2](#). The family $(v_\lambda)_\lambda$ has bounded variation, the sequence $(\mu_m^n)_{m \geq 1}$ is increasing, and $\mu_{k_0 a_n + 1}^n = \frac{1}{n - k_0 a_n} \leq \frac{1}{\varepsilon n}$. We deduce that the second sum vanishes as n tends to infinity.

Using [Section 2](#), we deduce that

$$\limsup_{n \rightarrow +\infty} \|w_1^n - v_{\lambda_1^n}\|_\infty \leq \varepsilon \|g\|_\infty.$$

Because $v_{\lambda_1^n} = v_{1/n}$ goes to v^* , and ε is arbitrary, we deduce that $\lim_{n \rightarrow +\infty} w_1^n = v^*$, hence the family $(\sigma_n)_n$ is asymptotically optimal. Exchanging the roles of the players, we obtain the proof of [Theorem 1.5](#). \square

3 Preliminary Results for the Proof of [Theorem 1.6](#)

This section presents two technical results, primarily adapted from the paper [\[3\]](#), which will be useful for the proof of [Theorem 1.6](#). First, we recall the results from [\[3\]](#) that we will exploit.

3.1 Results from the Paper [\[3\]](#)

For all $\lambda \in]0, 1[$ and $t \in [0, 1[$, define

$$\varphi(\lambda, t) := \inf \left\{ M \geq 1, \sum_{m=1}^M \lambda(1 - \lambda)^{m-1} \geq t \right\} = \left\lceil \frac{\ln(1 - t)}{\ln(1 - \lambda)} \right\rceil \in \mathbb{N} \cup \{+\infty\}.$$

Let us state Proposition 4.4 from [\[3\]](#):

Proposition 3.1. *Let $(x_\lambda, y_\lambda)_\lambda$ be a discounted optimal profile. The following properties are equivalent:*

1. The family $(x_\lambda, y_\lambda)_\lambda$ satisfies the discounted constant payoff property : for all $t \in]0, 1[$ and $\omega \in \Omega$,

$$\lim_{\lambda \rightarrow 0} \mathbb{E}_{x_\lambda, y_\lambda}^\omega \left[\sum_{m=1}^{\varphi(\lambda, t)} \lambda (1 - \lambda)^{m-1} g_m \right] = tv^*(\omega).$$

2. For all $\omega \in \Omega$ and $t \in]0, 1[$, $\lim_{\lambda \rightarrow 0} (\mathbb{E}_{\sigma_\lambda, \rho_\lambda}^\omega [v_\lambda(\omega_{\varphi(\lambda, t)})] - v_\lambda(\omega)) = 0$.

Moreover, in [3], it has been shown that any discounted optimal profile verifies the discounted constant payoff property:

Theorem 3.2. Any discounted optimal profile $(x_\lambda, y_\lambda)_\lambda$ verifies the discounted constant property. In particular, for all $\omega \in \Omega$ and $t \in]0, 1[$,

$$\lim_{\lambda \rightarrow 0} (\mathbb{E}_{\sigma_\lambda, \rho_\lambda}^\omega [v_\lambda(\omega_{\varphi(\lambda, t)})] - v_\lambda(\omega)) = 0.$$

3.2 Uniform Convergence in Proposition 3.1

Proposition 3.3. For all $\omega \in \Omega$ and $T \in]0, 1[$, we have

$$\sup_{t \in]0, T]} \mathbb{E}_{x_\lambda, y_\lambda}^\omega [v^*(\omega_{\varphi(\lambda, t)}) - v^*(\omega)] \xrightarrow{\lambda \rightarrow 0} 0.$$

We first need the following lemma:

Lemma 3.4. Let $(f_\lambda)_{\lambda \in]0, 1[}$ be a family of functions such that:

- For all $\lambda \in]0, 1[$, f_λ is a function $f_\lambda :]0, 1[\rightarrow \mathbb{R}$.
- $f_\lambda \xrightarrow{\lambda \rightarrow 0} 0$ pointwise.
- There exists $K > 0$ such that, for all $\lambda > 0$ small enough, and all $t, t' \in]0, 1[$ with $t < t'$:

$$|f_\lambda(t') - f_\lambda(t)| \leq K((t' - t) + \lambda). \quad (3.1)$$

Then $f_\lambda \xrightarrow{\lambda \rightarrow 0} 0$ uniformly.

Proof. Let $\varepsilon > 0$ and consider the finite set $S = \{k\varepsilon \mid k \in \llbracket 1; \lfloor \frac{1}{\varepsilon} \rfloor \rrbracket\}$.

The limit $f_\lambda \xrightarrow{\lambda \rightarrow 0} 0$ is uniform on the finite set S . Let us show that it is uniform on $]0, 1[$.

Let $\Lambda < \varepsilon$ small enough such that Eq. (3.1) is verified and:

$$\forall \lambda \in]0, \Lambda[, \forall s \in S, |f_\lambda(s)| \leq \varepsilon. \quad (3.2)$$

Let $t \in]0, 1[$ and an integer k such that $|t - k\varepsilon| \leq \varepsilon$. Combining Eq. (3.1) and Eq. (3.2), one has, for all $\lambda < \Lambda$:

$$\begin{aligned} |f_\lambda(t)| &\leq |f_\lambda(t) - f_\lambda(k\varepsilon)| + |f_\lambda(k\varepsilon)| \\ &\leq ((t - k\varepsilon) + \lambda) K + \varepsilon \\ &\leq (2K + 1)\varepsilon, \end{aligned}$$

which proves that $\|f_\lambda\|_\infty \xrightarrow{\lambda \rightarrow 0} 0$. □

Proof of Proposition 3.3. Using [Theorem 3.2](#), for all $\omega \in \Omega$ and $t \in]0, 1[$:

$$\lim_{\lambda \rightarrow 0} \mathbb{E}_{x_\lambda, y_\lambda}^\omega [v_\lambda(\omega_{\varphi(\lambda, t)}) - v_\lambda(\omega)] = 0. \quad (3.3)$$

Let us use [Lemma 3.4](#) to show that this limit is uniform on $]0, T[$ for all $T \in]0, 1[$. Fix $\omega \in \Omega$ and let us define, for all $\lambda \in]0, 1[$, the function

$$f_\lambda : t \mapsto (1 - \lambda)^{\varphi(\lambda, t) - 1} \mathbb{E}_{x_\lambda, y_\lambda}^\omega [v_\lambda(\omega_{\varphi(\lambda, t)}) - v_\lambda(\omega)],$$

defined on $]0, 1[$. [Eq. \(3.3\)](#) implies that $f_\lambda(t) \xrightarrow{\lambda \rightarrow 0} 0$ for all $t \in]0, 1[$.

Shapley's equations [\[8\]](#) yield, for any $p \in \mathbb{N}^*$:

$$v_\lambda(\omega) = \mathbb{E}_{\sigma_\lambda, \rho_\lambda}^\omega \left[\sum_{m=1}^{p-1} \lambda(1 - \lambda)^{m-1} g_m \right] + (1 - \lambda)^{p-1} \mathbb{E}_{\sigma_\lambda, \rho_\lambda}^\omega [v_\lambda(\omega_p)] \quad (\text{E}_p)$$

Making the difference of [Eq. \(E_p\)](#) with $p = b$ and [Eq. \(E_p\)](#) with $p = a$ yields:

$$\begin{aligned} & \left| (1 - \lambda)^{b-1} \mathbb{E}_{\sigma_\lambda, \rho_\lambda}^\omega [v_\lambda(\omega_b) - v_\lambda(\omega)] - (1 - \lambda)^{a-1} \mathbb{E}_{\sigma_\lambda, \rho_\lambda}^\omega [v_\lambda(\omega_a) - v_\lambda(\omega)] \right| \\ &= \left| \mathbb{E}_{\sigma_\lambda, \rho_\lambda}^\omega \left[\sum_{m=a}^{b-1} \lambda(1 - \lambda)^{m-1} g_m \right] + \left((1 - \lambda)^{a-1} - (1 - \lambda)^{b-1} \right) v_\lambda(\omega) \right| \\ &\leq 2 \|g\|_\infty \left((1 - \lambda)^{a-1} - (1 - \lambda)^{b-1} \right). \end{aligned} \quad (3.4)$$

Let us consider $t, t' \in]0, 1[$ with $t < t'$. We recall that $\varphi(\lambda, t) = \left\lceil \frac{\ln(1-t)}{\ln(1-\lambda)} \right\rceil$. Using [Eq. \(3.4\)](#):

$$\begin{aligned} |f_\lambda(t') - f_\lambda(t)| &\leq \frac{2}{1 - \lambda} \|g\|_\infty \left((1 - \lambda)^{\varphi(\lambda, t)} - (1 - \lambda)^{\varphi(\lambda, t')} \right) \\ &\leq \left((1 - \lambda)^{\frac{\ln(1-t)}{\ln(1-\lambda)}} - (1 - \lambda)^{\frac{\ln(1-t')}{\ln(1-\lambda)} + 1} \right) \\ &= \frac{2}{1 - \lambda} \|g\|_\infty [(1 - t) - (1 - t')(1 - \lambda)] \\ &= \frac{2}{1 - \lambda} \|g\|_\infty [(t' - t) + \lambda(1 - t')] \\ &\leq \frac{2}{1 - \lambda} \|g\|_\infty ((t' - t) + \lambda). \end{aligned}$$

For $\lambda \leq \frac{1}{2}$, we hence have:

$$|f_\lambda(t') - f_\lambda(t)| \leq 4 \|g\|_\infty ((t' - t) + \lambda).$$

Using [Lemma 3.4](#), for all $\omega \in \Omega$, f_λ converges uniformly on $]0, 1[$. It follows that for all $T \in]0, 1[$,

$$\sup_{t \in]0, T[} \mathbb{E}_{x_\lambda, y_\lambda}^\omega [v_\lambda(\omega_{\varphi(\lambda, t)}) - v_\lambda(\omega)] \xrightarrow{\lambda \rightarrow 0} 0.$$

As $v_\lambda \xrightarrow{\lambda \rightarrow 0} v^*$ uniformly, [Proposition 3.3](#) is proved. \square

3.3 A Sufficient Condition for the Constant Payoff Property

We show an analogous result to the implication (2. \implies 1.) in [Proposition 3.1](#) in the finite-horizon framework, as stated in the following proposition:

Proposition 3.5. *Let $(\sigma_n, \rho_n)_n$ be a family of Markov strategies. If for all $\omega \in \Omega$ and $t \in]0, 1[$,*

$$\mathbb{E}_{\sigma_n, \rho_n}^\omega (v^*(\omega_{\lceil tn \rceil + 1}) - v^*(\omega)) \xrightarrow{n \rightarrow +\infty} 0,$$

then $(\sigma_n, \rho_n)_n$ satisfies the constant payoff property.

Proof. Let $\omega \in \Omega$ and $n \geq 1$. Define $\varepsilon_n = \mathbb{E}_{\sigma_n, \rho_n}^\omega \left[\frac{1}{n} \sum_{m=1}^n g_m \right] - v_n(\omega)$. First, by [Theorem 1.5](#), the strategies σ_n and ρ_n are asymptotically optimal, hence $\varepsilon_n \xrightarrow{n \rightarrow +\infty} 0$.

Now, take $t \in]0, 1[$. One has:

$$\begin{aligned} \mathbb{E}_{\sigma_n, \rho_n}^\omega \left[\frac{1}{n} \sum_{m=1}^n g_m \right] &= \mathbb{E}_{\sigma_n, \rho_n}^\omega \left[\frac{1}{n} \sum_{m=1}^{\lceil tn \rceil} g_m \right] + \mathbb{E}_{\sigma_n, \rho_n}^\omega \left[\frac{1}{n} \sum_{m=\lceil tn \rceil + 1}^n g_m \right] \\ &= \mathbb{E}_{\sigma_n, \rho_n}^\omega \left[\frac{1}{n} \sum_{m=1}^{\lceil tn \rceil} g_m \right] + \frac{n - \lceil tn \rceil}{n} \mathbb{E}_{\sigma_n, \rho_n}^\omega \left[\frac{1}{n - \lceil tn \rceil} \sum_{m=1}^{n - \lceil tn \rceil} g_{m + \lceil tn \rceil} \right]. \end{aligned} \quad (3.5)$$

As a result,

$$\begin{aligned} \varepsilon_n &= \left(\mathbb{E}_{\sigma_n, \rho_n}^\omega \left[\frac{1}{n} \sum_{m=1}^{\lceil tn \rceil} g_m \right] - t v_n(\omega) \right) \\ &\quad + \frac{n - \lceil tn \rceil}{n} \left(\mathbb{E}_{\sigma_n, \rho_n}^\omega \left[\frac{1}{n - \lceil tn \rceil} \sum_{m=1}^{n - \lceil tn \rceil} g_{m + \lceil tn \rceil} \right] - \mathbb{E}_{\sigma_n, \rho_n}^\omega [v_{n - \lceil tn \rceil}(\omega_{\lceil tn \rceil + 1})] \right) \\ &\quad + \frac{n - \lceil tn \rceil}{n} (\mathbb{E}_{\sigma_n, \rho_n}^\omega [v_{n - \lceil tn \rceil}(\omega_{\lceil tn \rceil + 1})] - v_n(\omega)) \\ &\quad + \left(\frac{n - \lceil tn \rceil}{n} - (1 - t) \right) v_n(\omega). \end{aligned}$$

Clearly,

$$\left(\frac{n - \lceil tn \rceil}{n} - (1 - t) \right) v_n(\omega) \xrightarrow{n \rightarrow +\infty} 0,$$

and, using our hypothesis,

$$\mathbb{E}_{\sigma_n, \rho_n}^\omega (v_{n - \lceil tn \rceil}(\omega_{\lceil tn \rceil + 1}) - v_n(\omega)) \xrightarrow{n \rightarrow +\infty} 0.$$

To prove the proposition, it is thus enough to show that

$$\mathbb{E}_{\sigma_n, \rho_n}^\omega \left[\frac{1}{n - \lceil tn \rceil} \sum_{m=1}^{n - \lceil tn \rceil} g_{m + \lceil tn \rceil} \right] - \mathbb{E}_{\sigma_n, \rho_n}^\omega [v_{n - \lceil tn \rceil}(\omega_{\lceil tn \rceil + 1})] \xrightarrow{n \rightarrow +\infty} 0. \quad (3.6)$$

If this was not the case, there should be some n such that either

$$\frac{n - \lceil tn \rceil}{n} \left[\mathbb{E}_{\sigma_n, \rho_n}^\omega \left[\frac{1}{n - \lceil tn \rceil} \sum_{m=1}^{n - \lceil tn \rceil} g_{m + \lceil tn \rceil} \right] - \mathbb{E}_{\sigma_n, \rho_n}^\omega [v_{n - \lceil tn \rceil}(\omega_{\lceil tn \rceil + 1})] \right] < -2\varepsilon_n,$$

or the reverse inequality with ε_n instead of $-\varepsilon_n$. Assume the first inequality holds. Combining with Eq. (3.5), we get

$$\gamma_n^\omega(\sigma_n, \rho_n) < \mathbb{E}_{\sigma_n, \rho_n}^\omega \left[\frac{1}{n} \sum_{m=1}^{\lceil tn \rceil} g_m \right] + \frac{n - \lceil tn \rceil}{n} \mathbb{E}_{\sigma_n, \rho_n}^\omega [v_{n - \lceil tn \rceil}(\omega_{\lceil tn \rceil + 1})] + 2\varepsilon_n.$$

By playing σ_n until stage $\lceil tn \rceil$, and then an optimal strategy in $\Gamma_{n - \lceil tn \rceil}$, Player 1 would thus get strictly more than $\gamma_n^\omega(\sigma_n, \rho_n) + 2\varepsilon_n$. This contradicts the fact that σ_n and ρ_n are ε_n -optimal, which concludes the proof of the proposition. \square

4 Proof of Theorem 1.6

Let us recall Theorem 1.6.

Theorem. *There exists an adapted profile that satisfies the constant payoff property.*

Proof.

Step 1 (Definition of the strategy). Let $(x_\lambda, y_\lambda)_\lambda$ be a discounted optimal profile.

Using Proposition 3.3, for all $p \geq 1$, there exists $\mu \in]0, 1[$ such that:

$$\forall \lambda \in]0, \mu], \forall t \in]0, 7/8], \forall \omega \in \Omega, \quad \left| \mathbb{E}_{x_\lambda, y_\lambda}^\omega (v^*(\omega_{\varphi(\lambda, t)}) - v^*(\omega)) \right| \leq p^{-2}. \quad (4.1)$$

One can define a sequence $(\mu_p)_p$ in $]0, 1/2]$ such that $\mu_p \xrightarrow{p \rightarrow +\infty} 0$, and, for all $p \geq 1$, μ_p verifies Eq. (4.1). There exists $n_0 \in \mathbb{N}$ such that, for all $n \geq n_0$, the set $\{a \in \llbracket 2; n \rrbracket \mid 1/a \leq \mu_{\lfloor n/a \rfloor}\}$ is non-empty and thus has a minimum a_n . By definition, for all $n \geq n_0$,

$$2 \leq a_n \leq n \quad \text{and} \quad 1/a_n \leq \mu_{\lfloor n/a_n \rfloor}. \quad (4.2)$$

To respect the definition of an adapted profile, one can complete the sequence by defining $a_1, \dots, a_{n_0-1} \in \mathbb{N}$. Given that the property we are interested in is asymptotic, the exact values of $a_1, \dots, a_{n_0-1} \in \mathbb{N}$ are not relevant.

For each $\varepsilon > 0$ and n large enough, we have $1/(\varepsilon n) \leq \mu_{\lfloor n/(\varepsilon n) \rfloor}$, hence $a_n \leq \varepsilon n$. This shows that a_n/n tends to 0.

Let us reuse previously defined notations:

Notation. Let $n \geq 1$. Given $a_n \in \mathbb{N}$, let us define,

- for all $m \in \llbracket 1; n \rrbracket$, $k(m) := \left\lfloor \frac{m-1}{a_n} \right\rfloor$,
- $p_n := \left\lfloor \frac{n}{a_n} \right\rfloor$,
- for all $k \in \llbracket 0; p_n - 1 \rrbracket$, $\lambda_k^n := 1/(n - ka_n)$.

Let us define $(\sigma_n, \rho_n)_n$ as the adapted profile corresponding to the discounted optimal profile $(x_\lambda, y_\lambda)_\lambda$ and the sequence of positive integers $(a_n)_n$. More precisely, the strategy σ_n (resp., ρ_n) plays $x_{\lambda_{k(m)}^n}$ (resp., $y_{\lambda_{k(m)}^n}$) at each stage $m \in \llbracket 1; n \rrbracket$.

Let us show that it satisfies the constant payoff property.

Step 2 (A bound on the variation of the expected value within a block).

Notation. We define the following notations:

$$\forall \omega \in \Omega, \quad \mathbb{E}_n^\omega := \mathbb{E}_{\sigma_n, \rho_n}^\omega \quad \text{and} \quad \mathbb{E}_\lambda^\omega := \mathbb{E}_{x_\lambda, y_\lambda}^\omega.$$

Let $n \geq n_0$, $k \in \llbracket 0; p_n - 1 \rrbracket$, $j \in \llbracket 1; a_n + 1 \rrbracket$, and $\omega \in \Omega$. Let us show that

$$|\mathbb{E}_n^\omega [v^*(\omega_{ka_n+j}) - v^*(\omega_{ka_n+1})]| \leq p_n^{-2}. \quad (4.3)$$

Since

$$\mathbb{E}_n^\omega [v^*(\omega_{ka_n+j}) - v^*(\omega_{ka_n+1})] = \mathbb{E}_n^\omega \left[\mathbb{E}_{\lambda_k^n}^{\omega_{ka_n+1}} (v^*(\omega_j) - v^*(\omega_1)) \right],$$

it is sufficient to show that for all $\omega' \in \Omega$, all $j' \in \llbracket 1; a_n + 1 \rrbracket$,

$$\left| \mathbb{E}_{\lambda_k^n}^{\omega'} [v^*(\omega_{j'})] - v^*(\omega') \right| \leq p_n^{-2}. \quad (4.4)$$

Let $j' \in \llbracket 1; a_n + 1 \rrbracket$ and $\omega' \in \Omega$. The function $\begin{cases}]0, 1[& \rightarrow \mathbb{N} \cup \{0\} \\ t & \mapsto \varphi(\lambda_k^n, t) \end{cases}$ is surjective. Moreover, by definition of λ_k^n , we have

$$\lambda_k^n \leq \frac{1}{a_n}. \quad (4.5)$$

Using Eq. (4.5), one has:

$$\begin{aligned} \sum_{m=1}^{j'} \lambda_k^n (1 - \lambda_k^n)^{m-1} &= 1 - (1 - \lambda_k^n)^{j'} \\ &\leq 1 - \left(1 - \frac{1}{a_n}\right)^{a_n+1}. \end{aligned}$$

As the function $\lambda \mapsto 1 - (1 - \lambda)^{\frac{1}{\lambda}+1}$ is increasing on $]0, 1/2]$ and $a_n \geq 2$, one has:

$$\sum_{m=1}^{j'} \lambda_k^n (1 - \lambda_k^n)^{m-1} \leq \frac{7}{8}.$$

As a result, there exists $t \in]0, 7/8]$ such that $\varphi(\lambda_k^n, t) = j'$. Eq. (4.2) and Eq. (4.5) yield $\lambda_k^n \leq \mu_{p_n}$, and, combined with Eq. (4.1), we obtain

$$\begin{aligned} \left| \mathbb{E}_{\lambda_k^n}^{\omega'} [v^*(\omega_{j'})] - v^*(\omega') \right| &= \left| \mathbb{E}_{\lambda_k^n}^{\omega'} [v^*(\omega_{\varphi(\lambda_k^n, t)})] - v^*(\omega') \right| \\ &\leq p_n^{-2}. \end{aligned}$$

We deduce that [Eq. \(4.4\)](#) holds, hence [Eq. \(4.3\)](#) holds.

Step 3 (A general bound on the variation of the expected value).

Let $n \geq n_0$, $m \in \llbracket 1; p_n \cdot a_n \rrbracket$ and $\omega \in \Omega$. Let us show that

$$|\mathbb{E}_n^\omega [v^*(\omega_m)] - v^*(\omega)| \leq p_n^{-1}. \quad (4.6)$$

Indeed, let $k' \in \llbracket 0; p_n - 1 \rrbracket$ such that $k'a_n + 1 \leq m \leq (k' + 1)a_n$. We have

$$|\mathbb{E}_n^\omega [v^*(\omega_m)] - v^*(\omega)| \leq \sum_{\ell=0}^{k'-1} |\mathbb{E}_n^\omega [v^*(\omega_{(\ell+1)a_n+1})] - v^*(\omega_{\ell a_n+1})| + |\mathbb{E}_n^\omega [v^*(\omega_m)] - v^*(\omega_{k'a_n+1})|$$

Applying [Eq. \(4.3\)](#) to $k = \ell$ and $j = a_n + 1$, we obtain

$$|\mathbb{E}_n^\omega [v^*(\omega_{(\ell+1)a_n+1})] - v^*(\omega_{\ell a_n+1})| \leq p_n^{-2}.$$

Let $j = m - (k'a_n + 1)$. Note that $j \leq a_n$, $k' \leq p_n - 1$ and $m = k'a_n + (1 + j)$, hence, applying [Eq. \(4.3\)](#) yields:

$$|\mathbb{E}_n^\omega [v^*(\omega_m)] - v^*(\omega_{k'a_n+1})| \leq p_n^{-2}.$$

As a result,

$$|\mathbb{E}_n^\omega [v^*(\omega_m)] - v^*(\omega)| \leq 2p_n^{-1}.$$

Step 4 (Proof of the constant payoff property).

Let $t \in]0, 1[$. Because a_n/n tends to 0, for $n \geq n_0$ large enough, we have $a_n \leq n - \lceil tn \rceil$, hence $p_n a_n \geq n - a_n \geq \lceil tn \rceil$. Using [Eq. \(4.6\)](#), we get

$$|\mathbb{E}_{\sigma_n, \rho_n}^\omega (v^*(\omega_{\lceil tn \rceil+1}) - v^*(\omega))| \leq 2p_n^{-1}.$$

This implies that $|\mathbb{E}_{\sigma_n, \rho_n}^\omega (v^*(\omega_{\lceil tn \rceil+1}) - v^*(\omega))| \xrightarrow{n \rightarrow +\infty} 0$. Finally, [Proposition 3.5](#) proves the constant payoff property. \square

5 Perspectives

In the single-player case, the constant payoff property remains valid even when the state space and action sets are infinite [\[11\]](#), provided that (v_n) converges uniformly. A natural extension of our result would be to address two-player stochastic games with infinite action sets, infinite state space and/or imperfect observation of the state, for which (v_n) converges uniformly. Examples of such classes can be found in [\[9, 18, 4\]](#). To date, the only positive result in this direction pertains to discounted absorbing games with compact action sets [\[12\]](#).

Another direction is to investigate whether the constant payoff property holds for *any* asymptotically optimal profile.

Acknowledgments

This work was supported by the French Agence Nationale de la Recherche (ANR) under reference ANR-21-CE40-0020 (CONVERGENCE project). Part of this work was done during a 1-year visit of Bruno Ziliotto to the Center for Mathematical Modeling (CMM) at University of Chile in 2023, under the IRL program of CNRS. The authors are grateful to Rida Laraki and Guillaume Vigeral for fruitful discussions.

References

- [1] L. Attia and M. Oliu-Barton. A formula for the value of a stochastic game. *Proceedings of the National Academy of Sciences*, 116(52):26435–26443, 2019.
- [2] T. Bewley and E. Kohlberg. The asymptotic theory of stochastic games. *Mathematics of Operations Research*, 1(3):197–208, 1976.
- [3] O. Catoni, M. Oliu-Barton, and B. Ziliotto. Constant payoff in zero-sum stochastic games. *Annales de l’Institut Henri Poincaré (B) Probabilités et statistiques*, 57(4):1888–1900, 2021.
- [4] R. Laraki and S. Sorin. Advances in zero-sum dynamic games. In *Handbook of game theory with economic applications*, volume 4, pages 27–93. Elsevier, 2015.
- [5] J. Mertens and A. Neyman. Stochastic games have a value. *The Proceedings of the National Academy of Sciences*, 79:2145–2146, 1981.
- [6] A. Neyman and S. Sorin. *Stochastic games and applications*. Kluwer Academic Publishers, 2003.
- [7] M. Oliu-Barton. Weighted-average stochastic games with constant payoff. *Operational Research*, 22(3):1675–1696, 2022.
- [8] L.S. Shapley. Stochastic games. *Proceedings of the National Academy of Sciences of the United States of America*, 39(10):1095–1100, 1953.
- [9] E. Solan and B. Ziliotto. Stochastic games with signals. In *Advances in Dynamic and Evolutionary Games*, pages 77–94. Springer, 2016.
- [10] S. Sorin. *A first course on zero-sum repeated games*, volume 37. Mathématiques et Applications, Springer, 2002.
- [11] S. Sorin, X. Venel, and G. Vigeral. Asymptotic properties of optimal trajectories in dynamic programming. *Sankhya A*, 72(1):237–245, 2010.
- [12] S. Sorin and G. Vigeral. Limit optimal trajectories in zero-sum stochastic games. *Dynamic Games and Applications*, 10(2):555–572, 2020.
- [13] Sylvain Sorin and Guillaume Vigeral. Reversibility and oscillations in zero-sum discounted stochastic games. *Journal of Dynamics and Games (JDG)*, 2(1):103–115, 2015.
- [14] G. Vigeral. A zero-sum stochastic game with compact action sets and no asymptotic value. *Dynamic Games and Applications*, 3(2):172–186, 2013.
- [15] B. Ziliotto. A tauberian theorem for nonexpansive operators and applications to zero-sum stochastic games. *Mathematics of Operations Research*, 41(4):1522–1534, 2016.
- [16] B. Ziliotto. Zero-sum repeated games: Counterexamples to the existence of the asymptotic value and the conjecture $\max\min = \lim v_n$. *The Annals of Probability*, 44(2):1107–1133, 2016.
- [17] B. Ziliotto. A tauberian theorem for general iterations of nonexpansive operators. *Games and Economic Behavior*, 108:486–503, 2018.

- [18] B. Ziliotto. Mertens conjectures in absorbing games with incomplete information. *The Annals of Applied Probability*, 34(2):1948–1986, 2024.