

Undiscounted Equilibrium in Positive Recursive Absorbing Games with Non-Rectangular Absorption Structure*

Eilon Solan[†] and Nicolas Vieille[‡]

December 5, 2025

Abstract

An *absorbing game* is a stochastic game with a single nonabsorbing state. Such a game is called *recursive* if all players receive a payoff of 0 in the nonabsorbing state, and *positive* if all payoffs in absorbing states are positive. An action profile is *nonabsorbing* if, when it is played, the game remains in the nonabsorbing state with probability 1. The set of nonabsorbing action profiles can be partitioned into the connected components of an undirected graph, whose vertices are these profiles, with two vertices joined by an edge whenever the corresponding profiles differ in the action of a single player. A connected component is said to be *rectangular* if it is the Cartesian product of subsets of the players' action sets.

We prove that every positive recursive absorbing game whose non-absorbing components are all non-rectangular admits an undiscounted equilibrium payoff.

Keywords: Stochastic games, absorbing games, positive recursive games, undiscounted equilibrium, uniform equilibrium, absorption structure.

MSC2020 Classification: 91A06, 91A10, 91A15.

1 Introduction

One of the main open problems in game theory to date is whether every multiplayer stochastic game admits an undiscounted equilibrium payoff. This open problem was answered affirmatively for two-player zero-sum games (Mertens and Neyman (1981)), for two-player non-zero-sum games (Vieille (2000a,b)),

*The authors thank János Flesch for commenting on an earlier version of the paper. Solan acknowledges the support of the Israel Science Foundation grant #211/22.

[†]School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: eilons@tauex.tau.ac.il.

[‡]Department of Economics and Decision Sciences, HEC Paris, 1, rue de la Libération, 78 351 Jouy-en-Josas, France. E-mail: vieille@hec.fr.

and for various classes of stochastic games with more than two players (see, e.g., Solan (1999); Solan and Vieille (2001), Simon (2007, 2012), Flesch et al. (2007, 2008, 2009), and Solan et al. (2020).

Absorbing games are stochastic games with a single nonabsorbing state. This class of games was introduced by Kohlberg (1974), who proved the existence of the limit value in two-player zero-sum absorbing games, and has since been instrumental in the study of stochastic games. Indeed, the result of Kohlberg (1974) was extended to all stochastic games by Mertens and Neyman (1981). Similarly, the existence of an undiscounted equilibrium payoff in two-player non-zero-sum absorbing games, proved by Vrieze and Thuijsman (1989), was extended to all two-player stochastic games by Vieille (2000a,b). And the characterization of the value in zero-sum absorbing games (Laraki (2010)) was extended to all zero-sum stochastic games by Attia and Oliu-Barton (2019).

Some results that have been proven for absorbing games have so far resisted extension to stochastic games, e.g., the existence of an undiscounted equilibrium in three-player absorbing games (Solan (1999)) and in absorbing team games (Solan (2000)), the existence of an undiscounted normal-form correlated equilibrium in multiplayer absorbing games (Solan and Vohra (2002)), and weak approachability in absorbing games (Ragel (2024)). Some subclasses of absorbing games, such as quitting games (Solan and Vieille (2001)) and positive recursive absorbing games (Solan and Solan (2021)) were also studied and provided tools and useful insights for the study of stochastic games.

Main result and main idea behind the proof. In this paper, we focus on multiplayer absorbing games for which the stage payoff of all players in the nonabsorbing state is 0, and all absorbing payoffs are positive. We prove that an undiscounted equilibrium payoff exists in these games, under an additional restriction on the transition structure.

Our proof technique is an extension of the approach of Solan and Vieille (2001). Assuming the absorbing payoffs are bounded from above by 1, and denoting by I the set of players, we define a function f from a subset of $[0, 1]^I$ to itself, with the following property. Given w and $\varepsilon > 0$, there is an integer $T(w)$ and a strategy profile $\sigma(w)$ such that (i) $\sigma(w)$ is an ε -equilibrium in the $T(w)$ -stage version of the game that ends with a terminal payoff of w in case the game did not absorb before, (ii) the probability given $\sigma(w)$ that the game absorbs is positive and (iii) the payoff induced by $\sigma(w)$ is ε -close to $f(w)$.

We use a finite orbit of f to construct a sequence $(w^{(k)})_{k=0}^{K_0}$ such that $w^{(k)} = f(w^{(k+1)})$, with $w^{(K_0)} = (1, \dots, 1)$. Finally, we define a strategy profile σ^* by letting the players follow $\sigma(w^{(0)})$ for $T(w^{(0)})$ stages, then $\sigma(w^{(1)})$ for $N(w^{(1)})$ stages, and so on.

Significance of the paper. Beyond advancing the long-standing open problem of whether every multiplayer stochastic game admits an undiscounted equilibrium payoff, the paper makes several conceptual contributions.

First, we extend the dynamical-system approach to stochastic games, initiated by Solan and Vieille (2001) for a restricted class of quitting games, to a broad class of absorbing games.

Second, together with Solan et al. (2020), our result delineates where the main difficulty in establishing equilibrium in recursive absorbing games lies. The non-rectangularity condition identifies a large, easily checkable class of games for which the equilibrium problem is now resolved.

Third, the ideas developed here can be used to study more general stochastic games and other solution concepts. In a companion paper Solan and Vieille (2025), building on and refining the techniques introduced here, we show that in the presence of a public correlation device, every positive recursive stochastic game (not necessarily absorbing) admits an undiscounted equilibrium payoff. Thus, the present paper provides the core building block for a more general theory.

Fourth, our arguments can be used to substantially simplify existing proofs in the literature, such as the proof of the main result of Solan and Vieille (2001), thereby clarifying the underlying structure of those results.

Structure of the paper. The model and the main result are presented in Section 2. Section 3 presents some new and old concepts that will be used in the proof. The proof appears in Section 4, and open problems are discussed in Section 5.

2 Model and Main Result

2.1 The Model

Definition 2.1 (Recursive absorbing game) A recursive absorbing game is a tuple $\Gamma = (I, (A_i)_{i \in I}, (r_i)_{i \in I}, p)$, where

- I is a finite set of players.
- A_i is a finite set of actions, for each player $i \in I$. Denote the set of action profiles by $A := \prod_{i \in I} A_i$.
- $r_i : A \rightarrow \mathbf{R}$ is player i 's absorbing payoff function, for each $i \in I$.
- $p : A \rightarrow [0, 1]$ is the absorption probability function.

The game is played as follows. At every stage $n \in \mathbf{N}$, each player i selects an action $a_i^n \in A_i$ if the game was not absorbed before. With probability $p(a^n)$, where $a^n := (a_i^n)_{i \in I}$, the game is absorbed, and each player $i \in I$ receives the terminal payoff $r_i(a^n)$. With probability $1 - p(a^n)$ the game is not absorbed.

Definition 2.2 (Positive recursive absorbing game) A recursive absorbing game $\Gamma = (I, (A_i)_{i \in I}, (r_i)_{i \in I}, p)$ is positive if $r_i(a) > 0$, for every $i \in I$ and every $a \in A$.

For convenience, when the game is positive and recursive, we will assume that the payoff function r is bounded between 0 and 1.

We denote by $\theta \in \mathbf{N} \cup \{+\infty\}$ the stage at which absorption occurs. The game terminates in stage θ , yet it will be convenient to assume that players choose actions in all stages $n \in \mathbf{N}$, even after θ .

The set of *histories* is the set $H := \bigcup_{n=0}^{\infty} A^n$. A history $(a^1, a^2, \dots, a^n)$ is interpreted as the sequence of actions in the first n stages. A (behavior) *strategy* of player i is a function $\sigma_i : H \rightarrow \Delta(A_i)$. The set of all strategies of player i is denoted Σ_i , and the set of all strategy profiles of all players except i is denoted $\Sigma_{-i} = \prod_{j \neq i} \Sigma_j$. A *play* is an infinite sequence of action profiles $(a^1, a^2, \dots)$.

Every strategy profile $\sigma = (\sigma_i)_{i \in I}$ induces a probability distribution $\mathbf{P}_\sigma$ on the set of plays, endowed with the cylinder σ -algebra. We denote by $\mathbf{E}_\sigma$ the corresponding expectation operator.

The (expected) *undiscounted payoff* under strategy profile σ is

$$\gamma(\sigma) := \mathbf{E}_\sigma[r(a^\theta)1_{\theta < +\infty}] \in [0, 1]^I.$$

Note that for action profiles $a \in A$ with $p(a) = 0$, the value $r(a)$ does not affect the undiscounted payoff.

For every vector $b = (b_j)_{j \in I}$ we denote $b_{-i} = (b_j)_{j \neq i}$, and for every subset $J \subseteq I$ we denote $b_J = (b_j)_{j \in J}$ and $b_{-J} = (b_j)_{j \notin J}$.

Definition 2.3 (Undiscounted equilibrium) *Let $\varepsilon > 0$. The strategy profile σ^* is an undiscounted ε -equilibrium if for every $i \in I$ and every $\sigma_i \in \Sigma_i$,*

$$\gamma_i(\sigma^*) \geq \gamma_i(\sigma_i, \sigma_{-i}^*) - \varepsilon.$$

A vector $z \in \mathbf{R}^I$ is an undiscounted equilibrium payoff if $z = \lim_{k \rightarrow \infty} \gamma(\sigma^k)$, for some sequence $(\sigma^k)_{k \in \mathbf{N}}$ such that σ^k is a $\frac{1}{k}$ -equilibrium, for each $k \in \mathbf{N}$.

2.2 The Absorption Structure

Definition 2.4 (Set of nonabsorbing action profiles B) *The set of nonabsorbing action profiles is*

$$B := \{a \in A : p(a) = 0\}.$$

Definition 2.5 (Connected components of nonabsorbing action profiles)

Consider the graph whose set of vertices is the set of nonabsorbing action profiles B , and there is an edge between a and a' if these two action profiles differ in the action of at most one player; that is, there is $i \in I$ such that $a_{-i} = a'_{-i}$. Any connected component of this graph is called a connected component of nonabsorbing action profiles.

For simplicity, we will shorten the term “connected component of nonabsorbing action profiles” to “connected component”. Denote by $(B^\ell)_{\ell=1}^L$ the collection of connected components, so that $B = \bigcup_{\ell=1}^L B^\ell$.

Definition 2.6 (rectangular connected component) *The connected component B^ℓ is rectangular if it is a product set $B^\ell = \prod_{i \in I} B_i^\ell$.*

Example 2.7 *In Figure 1 we depict the absorption structure of a three-player absorbing game, where Player 1 selects a row, Player 2 selects a column, Player 3 selects a matrix, and absorbing entries are marked with an asterisk. In this game there are three connected components of nonabsorbing action profiles:*

- *One contains a single nonabsorbing action profile, and in particular is rectangular: $\{(a'_1, a'_2, a'_3)\}$.*
- *One is rectangular and contains six nonabsorbing action profiles: $\{a_1'''\} \times \{a_2'', a_2'''\} \times \{a_3, a'_3, a_3''\}$.*
- *One contains six nonabsorbing action profiles and is not rectangular: $\{(a_1, a_2), (a_1, a'_2), (a'_1, a_2)\} \times \{a_3, a'_3\}$.*

a_3					a_3'					a_3''				
	a_2	a_2'	a_2''	a_2'''		a_2	a_2'	a_2''	a_2'''		a_2	a_2'	a_2''	a_2'''
a_1			*	*	a_1			*	*	a_1	*	*	*	*
a_1'		*	*	*	a_1'		*	*	*	a_1'	*		*	*
a_1''	*	*	*	*	a_1''	*	*	*	*	a_1''	*	*	*	*
a_1'''	*	*			a_1'''	*	*			a_1'''	*	*		

Figure 1: The absorption structure of the three-player recursive absorbing game in Example 2.7.

2.3 Main Result

It follows from Solan and Vieille (2001) that all recursive absorbing games that satisfy the following restrictive conditions admit an undiscounted equilibrium payoff:

- Exactly one action profile $\hat{a}$ is nonabsorbing.
- For every player $i \in I$ and every action profile $b_{-i} \in A_{-i}$, one has $r_i(a_i, \hat{a}_{-i}) \geq r_i(a_i, b_{-i})$, for each a_i in the set

$$\hat{A}_i := \arg \max_{b_i \in A_i : p(b_i, \hat{a}_{-i}) > 0} r_i(b_i, \hat{a}_{-i})$$

of best absorbing actions against $\hat{a}$.

Solan et al. (2020) proved that all positive recursive absorbing games that satisfy the following conditions admit an undiscounted equilibrium payoff:

- There is exactly one connected component which is rectangular, so that the set of nonabsorbing entries is a Cartesian product $\prod_{i \in I} B_i$. Moreover, $|B_i| \geq 2$ for at least two players $i \in I$.
- Each player has exactly one action that is not part of any nonabsorbing entry: $|A_i \setminus B_i| = 1$, for every $i \in I$.
- The probability of absorption in all entries is either 0 or 1, that is, $p(a) \in \{0, 1\}$ for every $a \in A$.

Our main result shows that an undiscounted equilibrium payoff exists when *all* connected components are not rectangular.

Theorem 2.8 (Existence of undiscounted equilibrium) *Every positive recursive absorbing game that has no rectangular connected component admits an undiscounted equilibrium payoff.*

Remark 2.9 (Uniform equilibrium) *Theorem 2.8 also applies to the stronger concept of uniform equilibrium payoff, see Mertens et al. (2015), Section VII.4. A vector $z \in \mathbf{R}^I$ is a uniform equilibrium payoff if for every $\varepsilon > 0$ there is a strategy profile σ which is an ε -equilibrium in all games with finite horizon, and induces a payoff vector within ε of z in all such games, provided the horizon is sufficiently long.*

In positive recursive absorbing games, the expected average payoff $\gamma_n(\sigma)$ over the first n stages is nondecreasing in n , and converges to the undiscounted payoff $\gamma(\sigma)$. This implies that the two concepts of equilibrium payoffs coincide for such games.

Thus, Theorem 2.8 implies that every positive recursive absorbing game that has no rectangular connected component admits a uniform equilibrium payoff as well.

3 Basic Concepts

3.1 Mixed-Action Profiles and Connected Components

For every mixed action $x_i \in \Xi_i := \Delta(A_i)$, denote

$$\text{supp}(x_i) := \{a_i \in A_i : x_i(a_i) > 0\}.$$

Set $\Xi := \prod_{i \in I} \Xi_i$ and $\Xi_{-i} := \prod_{j \neq i} \Xi_j$. For every mixed action profile $x \in \Xi$ set $\text{supp}(x) := \prod_{i \in I} \text{supp}(x_i)$.

The set of nonabsorbing mixed action profiles whose support is contained in the connected component B^ℓ is

$$X^\ell := \{x \in \Xi : \text{supp}(x) \subseteq B^\ell\}, \quad \forall \ell \in [L] := \{1, 2, \dots, L\},$$

and the set of all nonabsorbing mixed action profiles is

$$X := \bigcup_{l=1}^L X^\ell.$$

We note that each X^ℓ is a connected subset of Ξ : if $x, x' \in X^\ell$, then there is a continuous path $(x_t)_{t \in [0,1]}$ in X^ℓ that starts at $x_0 = x$ and ends at $x_1 = x'$.

3.2 The minmax value

For every $i \in I$, denote by $v_i := \inf_{\sigma_{-i} \in \Sigma_{-i}} \sup_{\sigma_i \in \Sigma_i} \gamma_i(\sigma)$ player i 's minmax value. Since payoffs are non-negative and bounded by 1, we have $0 \leq v_i \leq 1$ for each $i \in I$.

A *minmaxing profile against i* is a profile (σ_i, σ_{-i}) such that $\gamma_i(\tilde{\sigma}_i, \sigma_{-i}) \leq v_i$ for each $\tilde{\sigma}_i \in \Sigma_i$.

In the proof, it is convenient to use stationary minmaxing profiles. The following result is a generalization of Flesch et al. (1996) to multiplayer positive recursive absorbing games.

Lemma 3.1 *For each $i \in I$ there exists a stationary minmaxing profile.*

Proof. For every $\lambda \in (0, 1]$ denote by $v_{i,\lambda} := \min_{\sigma_{-i} \in \Sigma_{-i}} \max_{\sigma_i \in \Sigma_i} \gamma_{i,\lambda}(\sigma)$ player i 's minmax value in the λ -discounted game, where $\gamma_{i,\lambda}(\sigma) := \mathbf{E}_\sigma [(1 - \lambda)^\theta r_i(a^\theta)]$. It is well known that there is a stationary profile $x_{-i,\lambda} \in \Xi_{-i}$ that achieves the minimum, and that $v_i = \lim_{\lambda \rightarrow 0} v_{i,\lambda} \geq 0$, see Neyman (2003). Any limit point $x_{-i,0}$ of the family $(x_{-i,\lambda})_{\lambda > 0}$ as $\lambda \rightarrow 0$ is a stationary minmaxing profile against i . Indeed, let $x_i \in \Xi_i$ be arbitrary. If the probability of absorption under $(x_i, x_{-i,0})$ is positive, then

$$\gamma_i(x_i, x_{-i,0}) = \lim_{\lambda \rightarrow 0} \gamma_{i,\lambda}(x_i, x_{-i,\lambda}) \leq \lim_{\lambda \rightarrow 0} v_{i,\lambda} = v_i.$$

If the profile $(x_i, x_{-i,0})$ is nonabsorbing, then $\gamma_i(x_i, x_{-i,0}) = 0 \leq v_i$. ■

As in much of the literature on dynamic games, our equilibrium constructions will be supported by the threat of switching to a minmaxing profile in case a player does not abide to a prespecified behavior.

3.3 The Maximal Absorbing Response

For every mixed action profile $x \in \Xi$, the expected absorption probability under x is

$$p(x) := \sum_{a \in A} \left(p(a) \prod_{j \in I} x_j(a_j) \right),$$

and player i 's expected absorption payoff under x given absorption is

$$r_i(x) := \sum_{a \in A} \left(r_i(a) p(a) \prod_{j \in I} x_j(a_j) \right) / p(x), \quad \forall i \in I.$$

The expected absorption payoff is defined only for absorbing mixed action profiles x .

Player i 's *payoff under a best absorbing response* to x_{-i} is

$$\rho_i(x) := \max\{r_i(a_i, x_{-i}) : a_i \in A_i, p(a_i, x_{-i}) > 0\}. \quad (1)$$

By convention, the maximum on an empty set is $-\infty$. Since the absorbing payoffs are positive, $\rho_i(x) > 0$ as soon as $\rho_i(x) \neq -\infty$. Note that $\rho_i(x)$ is independent of x_i . Note also that the functions $x \mapsto \rho_i(x)$ are lower semicontinuous: $\liminf_{k \rightarrow \infty} \rho_i(x^{(k)}) \geq \rho_i(x)$ for every sequence $(x^{(k)})_{k \in \mathbf{N}}$ that converges to x . This is so because when $(x^{(k)})_{k \in \mathbf{N}}$ is a sequence in Ξ that converges to x , if $p(a_i, x_{-i}) > 0$, then $p(a_i, x_{-i}^{(k)}) > 0$ for every $k \in \mathbf{N}$ sufficiently large, and hence the maximum in Eq. (1) contains fewer actions for x than for $x^{(k)}$, for every large k .

When $\rho_i(x) > -\infty$, let $a_i(x) \in A_i$ be a *best absorbing response* to x_{-i} , namely, an action that attains the maximum in Eq. (1).

We note that if there is a nonabsorbing mixed action profile such that $\rho_i(x) = -\infty$, then $v_i = 0$. Indeed, since the game is positive and recursive, $v_i \geq 0$, and by playing the mixed action profile x_{-i} , the other players lower the best payoff player i can obtain to 0.

If there is $x \in X$ such that $\rho_i(x) = -\infty$ for every $i \in I$, then the stationary strategy x is an undiscounted 0-equilibrium: it yields each player the payoff 0, and all unilateral deviations are nonabsorbing and hence not profitable. Since our goal is to prove the existence of an undiscounted equilibrium payoff, we can make the following assumption.

Assumption 3.2 (No nonabsorbing equilibrium) *For every $x \in X$ there is $i \in I$ such that $\rho_i(x) > 0$.*

We note¹ that $\rho_i(x) \geq v_i$ whenever $\rho_i(x) > 0$. It will be convenient to denote

$$\rho(x) := (\rho_i(x))_{i \in I}, \quad \forall x \in X.$$

3.4 Partitions of the Space of Payoff Vectors

For every nonabsorbing mixed action profile $x \in X$ we partition $\mathbf{R}^I$ into three subsets, $W(x)$, $W_L(x)$, and $W_H(x)$. Define (see Figure 2 for an illustration with two players):

$$\begin{aligned} W_H(x) &:= \{w \in \mathbf{R}^I : w_i > \rho_i(x), \quad \forall i \in I\}, \\ W(x) &:= \left\{ w \in \mathbf{R}^I : \begin{array}{ll} w_i \geq \rho_i(x) & \forall i \in I, \\ w_i = \rho_i(x), & \text{for some } i \in I \end{array} \right\}, \\ W_L(x) &:= \{w \in \mathbf{R}^I : w_i < \rho_i(x), \quad \text{for some } i \in I\}. \end{aligned}$$

¹Assumption 3.2 and the inequality $\rho_i(x) \geq v_i$ whenever the left-hand side is finite are two properties that are used in the proof and require the assumption that the game is recursive and positive.




Figure 2: The sets $W(x)$, $W_L(x)$, and $W_H(x)$.

Set

$$W_H := \bigcup_{x \in X} W_H(x), \quad W := \bigcup_{x \in X} W(x), \quad W_L := \bigcup_{x \in X} W_L(x).$$

We note that the three sets W , $W_H \setminus W$, and $\bigcap_{x \in X} W_L(x)$ form a partition of $\mathbf{R}^I$. Indeed, by De Morgan's laws,

$$\begin{aligned} \bigcap_{x \in X} W_L(x) &= \bigcap_{x \in X} (\mathbf{R}^I \setminus (W_H(x) \cup W(x))) \\ &= \mathbf{R}^I \setminus \left(\bigcup_{x \in X} (W_H(x) \cup W(x)) \right) \\ &= \mathbf{R}^I \setminus (W_H \cup W). \end{aligned}$$

3.5 Exits

We here present and study the notions of exits and joint exits, which were introduced by Solan (1999).

Definition 3.3 (Exit, joint exit) *Let $x \in X$ be a nonabsorbing mixed action profile. A pair (J, a_J) where $\emptyset \neq J \subseteq I$ and $a_J \in \prod_{i \in J} A_i$ is an exit at x if*

- $p(a_J, x_{-J}) > 0$.
- $p(a_{J'}, x_{-J'}) = 0$ for every $J' \subset J$.

An exit (J, a_J) is joint if $|J| \geq 2$. The set of joint exits from x is denoted $\mathcal{E}_J(x)$.

For future reference we note that if a sequence of nonabsorbing mixed action profiles that have joint exits converges to a limit, then the limit also has a joint exit.

Lemma 3.4 *If $(x^{(k)})_{k \in \mathbf{N}}$ is a sequence of nonabsorbing mixed actions that converges to x , and if $\mathcal{E}_J(x^{(k)}) \neq \emptyset$, then $\mathcal{E}_J(x) \neq \emptyset$.*

The following observation relates joint exits to rectangular connected components. We omit the proof as well.

Lemma 3.5 *The connected component B^ℓ is rectangular if and only if there is no joint exit at any mixed action profile in X^ℓ .*

Example 3.6 *Consider the positive recursive absorbing game in Figure 3, where Player 1 has 3 actions, Player 2 has 4 actions, empty entries are nonabsorbing, and each nonempty entry a is absorbing with positive probability and contains the vector $r(a)$. The four pairs $(\{1, 2\}, (a'_1, a'_2))$, $(\{1, 2\}, (a'_1, a''_2))$, $(\{1, 2\}, (a''_1, a'_2))$, $(\{1, 2\}, (a''_1, a''_2))$ are joint exits from (a_1, a_2) . The pair $(\{2\}, (a_1, a''_2))$ is an exit of (a_1, a_2) but not a joint exit. There is no joint exit at (a'_1, a_2) .*

	a_2	a'_2	a''_2	a'''_2
a_1				$(\frac{1}{2}, \frac{1}{4})$
a'_1		$(\frac{1}{4}, \frac{1}{3})$	$(\frac{1}{2}, \frac{1}{2})$	$(\frac{1}{3}, \frac{3}{5})$
a''_1		$(\frac{1}{5}, \frac{3}{4})$	$(\frac{2}{5}, \frac{3}{5})$	$(\frac{2}{3}, \frac{2}{3})$

Figure 3: The positive recursive absorbing game in Example 3.6.

The following result states that at all discontinuity points of $x \mapsto \rho(x)$ there is a joint exit.

Lemma 3.7 *If $\rho(\cdot)$ is not continuous at $x \in X$, then $\mathcal{E}_J(x) \neq \emptyset$.*

Proof. Suppose that there is no joint exit at x . We will show that $\rho_i(x) = \lim_{k \rightarrow \infty} \rho_i(x^{(k)})$ for every $i \in I$.

For every $y \in X$ denote the set of all absorbing responses of player i to y_{-i} by

$$\hat{A}_i(y) := \{a_i \in A_i : p(a_i, y_{-i}) > 0\}.$$

Assume w.l.o.g. that $\text{supp}(x^{(k)})$ is independent of k , and therefore $\hat{A}_i(x^{(k)})$ is independent of k . In this proof we will denote this set simply by $\hat{A}_i(x^{(k)})$.

Fix $i \in I$. To prove that $\rho_i(x) = \lim_{k \rightarrow \infty} \rho_i(x^{(k)})$ it is sufficient to show that $\hat{A}_i(x) = \hat{A}_i(x^{(k)})$.

We note that the inequality $\hat{A}_i(x) \subseteq \hat{A}_i(x^{(k)})$ holds because $\text{supp}(x^{(k)}) \supseteq \text{supp}(x)$ for all $k \in \mathbf{N}$. We will next prove the reverse inclusion.

We argue by contradiction and assume that there is $a_i \in \widehat{A}_i(x^{(k)}) \setminus \widehat{A}_i(x)$. Since $p(a_i, x_{-i}^{(k)}) > 0$ for all $k \in \mathbf{N}$, while $p(a_i, x_{-i}) = 0$, there is an action profile $\widehat{a}_{-i} \in \text{supp}(x_{-i}^{(k)})$ for all $k \in \mathbf{N}$, such that $p(a_i, \widehat{a}_{-i}) > 0$. Denoting by $J \subseteq I \setminus \{i\}$ the set of all players j such that $\widehat{a}_j \notin \text{supp}(x_j)$, we obtain that $|J| \geq 1$ and $(J \cup \{i\}, (a_i, \widehat{a}_J))$ is a joint exit at x . ■

The next result relates joint exits to the set $W_H(x)$: it states that if a payoff vector w lies in $W_H \setminus W$, then there is $x \in X$ such that not only does w lie in $W_H(x)$, but also x has a joint exit.

Lemma 3.8 *If $w \in W_H \setminus W$, then there are $\ell \in [L]$ and $x \in X^\ell$ such that $w \in W_H(x)$ and $\mathcal{E}_J(x) \neq \emptyset$.*

Proof. Since $w \in W_H \setminus W$, there are $\ell \in [L]$ and $x' \in X^\ell$ such that $w \in W_H(x') \setminus W$. If $\mathcal{E}_J(x') \neq \emptyset$, we are done. Assume then that $\mathcal{E}_J(x') = \emptyset$.

Since B^ℓ is not rectangular, there is a mixed-action profile in X^ℓ at which there is a joint exit. Since X^ℓ is connected, there is a continuous path $(x_t)_{t \in [0,1]}$ in X^ℓ that starts at x' and ends at some $x'' \in X^\ell$ such that $\mathcal{E}_J(x'') \neq \emptyset$. By Lemma 3.4 we can furthermore assume that x'' is the *only* point along the path where $\mathcal{E}_J(x_t)$ is nonempty.

By Lemma 3.7, the function ρ is continuous along the path, except, possibly at its end. Since $w \in W_H(x')$, it follows that $w_i > \rho_i(x_0)$ for all $i \in I$. Since $w \notin W(x_t)$ for all $t \in [0, 1)$ and since $t \mapsto \rho(x_t)$ is continuous on $[0, 1)$, it follows that $w_i > \rho_i(x_t)$ for all $t \in [0, 1)$. Since ρ_i is lower semicontinuous, this inequality holds also at $t = 1$. This implies that $w \in W_H(x'')$, which is what we want to show. ■

The following result, which is a special case of Lemma 5.3 in Solan (1999), asserts that if there exists a joint exit that yields all players a high payoff, then an undiscounted equilibrium payoff exists.

Lemma 3.9 (Solan (1999), Lemma 5.3) *Suppose there exists a nonabsorbing mixed action profile $x \in X$ and a joint exit $(J, a_J) \in \mathcal{E}_J(x)$ such that $r_i(a_J, x_{-J}) \geq \rho_i(x)$ for each $i \in I$. Then $r(a_J, x_{-J})$ is an undiscounted equilibrium payoff.*

For a given $\varepsilon > 0$, one undiscounted ε -equilibrium under the conditions of Lemma 3.9 is the following. Until some player is declared the deviator, the players select the mixed action profile ξ defined by

$$\xi_i := \begin{cases} x_i, & i \notin J, \\ (1 - \beta)x_i + \beta a_i, & i \in J, \end{cases} \quad (2)$$

where $\beta > 0$ is sufficiently small. Thus, if no player is declared the deviator, the play is eventually absorbed, and the absorbing payoff is $r(a_J, x_{-J})$.

Since $r_i(a_J, x_{-J}) \geq \rho_i(x)$ for each $i \in I$, no player i can profit more than β by deviating in an absorbing way; namely, by selecting an action $a_i \in A_i$ such

that $a_i \notin \text{supp}(x_i)$ and $p(a_i, x_{-i}) > 0$. Since the game is positive recursive, no player in J can profit by never selecting a_i .

Players not in J can still profit by changing the frequency in which they select actions in $\text{supp}(x_i)$. Such a deviation can be identified by a standard statistical test. For more details, see the proof of Lemma 5.3 in Solan (1999).

4 Proof of Theorem 2.8

If the condition of Lemma 3.9 holds, then the game admits an undiscounted equilibrium payoff. Assume then that this condition does not hold.

4.1 The Auxiliary One-Shot Game $G(w)$

For every vector $w \in \mathbf{R}^I$ denote by $G(w)$ the one-shot game derived from Γ by assuming that if the game does not absorb in the first stage, then the players get a terminal payoff w . Formally, $G(w)$ is the following one-shot game:

- The set of players is I , as in Γ .
- The set of actions of each player $i \in I$ is A_i , as in Γ .
- The payoff function of each player $i \in I$ is

$$u_i^w(a) := p(a)r_i(a) + (1 - p(a))w_i, \quad \forall a \in A.$$

4.2 Definition of Two Functions f and μ

Fix $\varepsilon \in (0, 1)$, and set $Y := \prod_{i \in I} [v_i - \varepsilon, 1]$. The set Y contains all payoff profiles that are approximately individually rational for the players. In this section we define two functions, $f : Y \rightarrow Y$ and $\mu : Y \rightarrow (0, 1]$.

Recall that the three sets W , $W_H \setminus W$, and $\bigcap_{x \in X} W_L(x)$ form a partition of $\mathbf{R}^I$. For $w \in Y$, the definition of $f(w)$ and $\mu(w)$ depends on which of these three sets contains w .

In each of the following three sections we define f on one of these sets. For every nonabsorbing mixed action profile $x \in X$ and every player $i \in I$, denote by $a_i(x)$ a best absorbing response of player i at x .

Fix $w \in Y$.

4.2.1 Definition of f and μ on $W \cap Y$

Assume $w \in W \cap Y$. Then there is $x \in X$ such that $w \in W(x)$. It follows that $w_i \geq \rho_i(x)$ for each $i \in I$, with equality for at least one player, denoted i_0 . Set

$$f(w) := \varepsilon p(a_{i_0}(x), x_{-i_0}) r(a_{i_0}(x), x_{-i_0}) + (1 - \varepsilon p(a_{i_0}(x), x_{-i_0})) w, \quad (3)$$

and

$$\mu(w) := \varepsilon p(a_{i_0}(x), x_{-i_0}).$$

The quantity $f(w)$ is the payoff in $G(w)$ under the strategy profile where player i_0 selects the mixed action $(1 - \varepsilon)x_{i_0} + \varepsilon a_{i_0}(x)$, and the other players select the mixed action profile x_{-i_0} .

We argue that $f(w) \in Y$. Fix for a moment $i \in I$. If $\rho_i(x) > -\infty$, then $\rho_i(x) \geq v_i > 0$. Since $w_i \geq \rho_i(x) \geq v_i$ and since $\|f(w) - w\|_\infty \leq \varepsilon$, it follows that $f_i(w) \geq v_i - \varepsilon$. Suppose next that $\rho_i(x) = -\infty$, so that $v_i = 0$. Since the game is positive, $f_i(w) \geq (1 - \varepsilon)w_i \geq -(1 - \varepsilon)\varepsilon > -\varepsilon = v_i - \varepsilon$. Since i is arbitrary, $f(w) \in Y$.

4.2.2 Definition of f and μ on $(W_H \setminus W) \cap Y$

Assume $w \in (W_H \setminus W) \cap Y$. By Lemma 3.8, there is $x \in X$ such that $w \in W_H(x)$ and $\mathcal{E}_J(x) \neq \emptyset$. Let $(J, a_J) \in \mathcal{E}_J(x)$. Since the conditions of Lemma 3.9 do not hold, there is $i_0 \in I$ such that $r_{i_0}(a_J, x_{-J}) < \rho_{i_0}(x) - \varepsilon$.

Since $w \in W_H(x)$, we have $w_i > \rho_i(x)$ for every $i \in I$. Let $\alpha \in (0, 1)$ be the unique real number such that $(1 - \alpha)w_i + \alpha r_i(a_J, x_{-J}) \geq \rho_i(x)$ for every $i \in I$, with an equality for at least one i .

Define

$$f(w) := \alpha r(a_J, x_{-J}) + (1 - \alpha)w, \quad (4)$$

and

$$\mu(w) = \alpha.$$

Note that $f(w) \geq \rho(x) \geq v$ (component-wise), and hence $f(w) \in Y$.

The quantity $f(w)$ is the payoff in $G(w)$ when with probability α the play is absorbed through (a_J, x_{-J}) , and with probability $1 - \alpha$ the game is not absorbed.

4.2.3 Definition of f and μ on $(\bigcap_{x \in X} W_L(x)) \cap Y$

Assume $w \in (\bigcap_{x \in X} W_L(x)) \cap Y$. Let $\hat{x}$ be a Nash equilibrium of the auxiliary one-shot game $G(w)$. We argue that $p(\hat{x}) > 0$. Indeed, suppose by contradiction that $p(\hat{x}) = 0$, so that $\hat{x} \in X$. Since $w \in \bigcap_{x \in X} W_L(x) \subseteq W_L(\hat{x})$, there is $i \in I$ such that $w_i < \rho_i(\hat{x})$. But then player i can profit by deviating at $\hat{x}$ to $a_i(\hat{x})$, since in $G(w)$ her payoff under $\hat{x}$ is w_i , while her payoff under $(a_i(\hat{x}), \hat{x}_{-i})$ is a convex combination of w_i and $\rho_i(\hat{x})$. This contradicts the fact that $\hat{x}$ is a Nash equilibrium of $G(w)$.

Set

$$f(w) := u^w(\hat{x}),$$

and

$$\mu(w) := p(\hat{x}).$$

We argue that $f_i(w) \geq v_i - \varepsilon$ for each player $i \in I$, and hence $f(w) \in Y$. To this end, we will show that player i 's minmax value in $G(w)$ is at least $v_i - \varepsilon$. It is well known that for every $\lambda \in (0, 1]$, player i 's minmax value in $G((1 - \lambda)v_\lambda)$ is $v_{i,\lambda}$. By continuity of the minmax value operator, player i 's minmax value in $G(v)$ is v_i . Since the minmax value operator is nondecreasing and nonexpansive, and since $w_i \geq v_i - \varepsilon$, player i 's minmax value in $G(w)$ is at least $v_i - \varepsilon$.

4.3 Constructing Equilibria in Finite Stage Games

In this section, we explain how to implement $f(w)$ as an approximate equilibrium payoff in a finite-stage version of Γ , in which the terminal payoff is $w \in Y$ in case the game does not absorb. In Section 4.4 we explain how to concatenate this construction over several blocks.

Given $T \in \mathbf{N}$ and $w \in Y$, we denote by $G_T(w)$ the T -stage game that is derived from Γ by assuming that the players get a terminal payoff w if the game does not absorb in the first T stages. The game $G(w)$ introduced in Section 4.1 thus coincides with $G_1(w)$.

The following result states that for every $w \in Y$ there is $T \in \mathbf{N}$ sufficiently large, such that in $G_T(w)$ there is an approximate equilibrium with payoff close to $f(w)$. The statement of this lemma involves a stopping time τ_i , which is interpreted as the first stage in which player i is identified as a deviator.

Lemma 4.1 *For each $w \in Y$ and $\eta > 0$, there is $T = T(w, \eta) \in \mathbf{N}$ such that the game $G_T(w)$ has the following property: There is a stationary strategy profile $\sigma = \sigma(w, \eta)$ and a family $(\tau_i)_{i \in I}$ of stopping times such that:*

- $\mathbf{P}_\sigma(\min_{i \in I} \tau_i \leq T) \leq \eta$.
- $f(w) = \mathbf{E}_\sigma[r(a^\theta)1_{\theta \leq T} + w1_{\theta > T}]$ and $\mu(w) = \mathbf{P}_\sigma(\theta \leq T)$.
- For each player i and each alternative strategy $\tilde{\sigma}_i \in \Sigma_i$,

$$\begin{aligned} \mathbf{E}_{\tilde{\sigma}_i, \sigma_{-i}}[r_i(a^\theta)1_{\theta \leq \min\{T, \tau_i\}} + w_i1_{T < \min\{\theta, \tau_i\}} + v_i1_{\tau_i \leq \min\{T, \theta - 1\}}] \\ \leq f(w) + \eta + \varepsilon \mathbf{P}_{\tilde{\sigma}_i, \sigma_{-i}}(\tau_i \leq \min\{T, \theta - 1\}). \end{aligned}$$

The first property states that under σ , it is highly unlikely that some player will be incorrectly identified as a deviator. The second property states that, ignoring $(\tau_i)_{i \in I}$, the payoff and probability of absorption induced by σ coincide with $f(w)$ and $\mu(w)$. The last property bounds the expected gain obtained by deviating, assuming that a deviator is punished at her minmax value if it is detected. The formulation of the last property is important when concatenating several such finite-stage games, to avoid small gains from accumulating into a larger one.

Proof. Fix $w \in W$ and $\eta > 0$. The definition of σ and T will reflect the construction of f .

Case 1: $w \in \bigcap_{x \in X} W_L(x)$. In that case, there is $\hat{x} \in X$ such that $p(\hat{x}) > 0$ and $\hat{x}$ is a Nash equilibrium of $G(w)$ with payoff $f(w)$ (see Section 4.2.3). We set $T = 1$, $\sigma = \hat{x}$, and no statistical test is employed: $\tau_i = +\infty$ for every $i \in I$.

Case 2: $w \in W$. In that case, there is a nonabsorbing mixed action profile $x \in X$ and a player $i_0 \in I$ such that $w \in W(x)$ and

$$f(w) = (1 - \varepsilon p(a_{i_0}(x), x_{-i}))w + \varepsilon p(a_{i_0}(x), x_{-i})r(a_{i_0}(x), x_{-i}),$$

(see Section 4.2.1).

The profile σ plays repeatedly the mixed action profile y defined by

$$y_i := \begin{cases} x_i, & i \neq i_0, \\ (1 - \beta)x_i + \beta a_{i_0}(x), & i = i_0, \end{cases}$$

where β is specified below.

Given β , the duration $T \in \mathbf{N}$ of the game is such that the total probability that the game is absorbed is ε :

$$(1 - \beta p(a_{i_0}(x), x_{-i}))^T = 1 - \varepsilon.$$

Against σ_{-i} , a player i may potentially benefit by

- increasing or lowering the probability of ever playing a_{i_0} , if $i = i_0$;
- modifying the relative frequencies of actions in the support of x_i , for $i \neq i_0$;
- choosing an action not in the support of a_i , for any $i \in I$.

The first type of deviation is not profitable, given the properties of a_{i_0} ; the second type of deviation can be monitored by checking that the empirical distribution of each player i 's choices remains close to x_i ; the last type of deviation is immediately detected.

The stopping time τ_i deals with the latter two types of deviation and is defined to be the first stage in which player i chooses an action which is not in the support of y_i , or in which the empirical distribution of player i 's actions is not within some appropriate confidence set around x_i , that depends on η . It is set to $\tau_i = +\infty$ otherwise, if both tests are passed.

The precise definition of $(\tau_i)_{i \in I}$ uses standard statistical tests which appear in, e.g., the proof of Lemma 5.3 in Solan (1999) or in the proof of Proposition 7 in Vieille (2000b). For that reason, we omit details.

If β is small enough, and for an appropriate confidence set, σ and T satisfy the conclusions of Lemma 4.1: (i) undetected deviations modify the distribution of moves by at most η and therefore do not affect payoffs by more than η ; (ii) since $f_i(w) \geq \rho_i(w) - \varepsilon \geq v_i - \varepsilon$ for each i , detected deviations do not improve payoffs by more than ε .

Case 3: $w \in W_H \setminus W$. In that case, there is $x \in X$, a joint exit $(J, a_J) \in \mathcal{E}_J(x)$, and $\alpha \in (0, 1)$, such that $w \in W_H(x)$ and

$$f(w) = \alpha r(a_J, x_{-J}) + (1 - \alpha)w,$$

(see Section 4.2.2).

The profile σ plays repeatedly the mixed action profile y defined by

$$y_i := \begin{cases} x_i, & i \notin J, \\ (1 - \beta)x_i + \beta a_i, & i \in J, \end{cases}$$

where β will be specified below. Given β , the duration T of the block is set to be such that the total probability that the game is absorbed during the block is α :

$$(1 - \beta^{|J|} p(a_J, x_{-J}))^T = 1 - \alpha.$$

Against σ_{-i} , a player i may potentially benefit by

- modifying the frequency of a_i , if $i \in J$;
- modifying the relative frequencies of actions in the support of x_i , for $i \notin J$;
- choosing an action not in the support of a_i , for any $i \in I$.

The latter two types of deviation are prevented as in **Case 2**, but the former one is different. For β small, the definition of T ensures that $T\beta$ is of the order of $\frac{1}{\beta^{|J|-1}}$ up to constants, so that the expected number of stages in which player $i \in J$ chooses a_i increases to $+\infty$ as $\beta \rightarrow 0$. This fact allows players $-i$ to monitor the empirical frequency with which a_i is chosen. We refer to Vieille (2000b), Section 8.2 for details. ■

4.4 Constructing a Strategy profile σ^*

We use the function f to construct an approximate equilibrium σ^* . Our construction is inspired by Solan and Vieille (2001), see also Section 12.3 in Solan (2022). The play will be divided into blocks. On each block, the strategy profile σ^* is obtained by applying Lemma 4.1 to some continuation payoff $w^{(k)}$, as long as the game has not absorbed and no deviator has been identified.

The construction proceeds as follows. In Step 1 we use an approximate finite orbit of f , to construct a finite sequence $(w^{(k)})_{k=1}^{K_0}$ such that $f(w^{(k+1)})$ is close to $w^{(k)}$. In Step 2 we will use this sequence to construct a strategy profile in blocks, such that for each $k = 0, \dots, K_0$, the expected payoff from the beginning of block k will be close to $w^{(k)}$. Within block k , the strategy profile σ^* will be derived from $f(w^{(k+1)})$ using Lemma 4.1

In Section 4.5 we will prove that σ^* is an undiscounted 7ε -equilibrium.

Step 1: Definition of a sequence of (approximate) continuation payoffs.

Let $\delta \in (0, -\frac{1}{\ln(\varepsilon)})$, so that $\exp(-1/\delta) < \varepsilon$, and hence in particular $\delta < \varepsilon$. By Theorem 4 in Solan and Solan (2020), see also Theorem 12.6 in Solan (2022), f has an approximate orbit that starts at the vector of maximal payoffs.²

²Both Solan and Solan (2020) and Solan (2022) prove this result when f has no fixed point. However, the result also holds when f has fixed points. To see this, note that if $w^{(k)}$ is not δ -close to any fixed point of f , for all k , then the proof in Solan and Solan (2020) and Solan (2022) is valid. If $w^{(k)}$ is δ -close to some fixed point $\tilde{w}$ of f , set $w^{(l)} = \tilde{w}$ for every $l < k$. We then obtain (5)–(7) with 2δ instead of δ .

Formally, this theorem implies the existence of a finite sequence $(w^{(k)})_{k=0}^{K_0}$ that satisfies the following properties:

$$w^{(K_0)} := (1, 1, \dots, 1), \quad (5)$$

$$\sum_{k=0}^{K_0-1} \mu(w^{(k)}) \geq \frac{1}{\delta}, \quad (6)$$

$$\sum_{k=0}^{K_0-1} \|w^{(k)} - f(w^{(k+1)})\|_{\infty} \leq \delta. \quad (7)$$

Step 2: Definition of σ^* .

Fix $\eta^{(0)}, \dots, \eta^{(K_0-1)}$ such that $\sum_{k=0}^{K_0-1} \eta^{(k)} < \varepsilon$. For each k , $0 \leq k \leq K_0-1$, we apply Lemma 4.1 with $w = w^{(k+1)}$ and $\eta = \eta^{(k)}$. We denote by $T^{(k)}$, $\sigma^{(k)}$ and $(\tau_i^{(k)})_{i \in I}$ the horizon, the stationary profile, and the stopping times for which the conclusion of Lemma 4.1 holds.

The strategy profile σ^* is defined in K_0 blocks. The length of block k is $T^{(k)}$. The cumulative duration of the first k blocks is $N^{(k)} := \sum_{m=0}^{k-1} T^{(m)}$.

The players follow the sequence $(\sigma^{(k)})_{k=0}^{K_0-1}$ of stationary profiles in the successive blocks as long as no deviator has been identified; that is, as long as the stopping times $(\tau_i^{(k)})_{i \in I}$ are not reached. The behavior of the players on each block is monitored by means of the stopping times $(\tau_i^{(k)})$. We denote by τ_i the first stage $n \leq N^{(K_0-1)}$ in which player i is identified as a deviator, and by $\tau := \min_{i \in I} \tau_i$ the first stage in which some player is identified as the deviator. Whenever $\tau < \infty$ we let i_* be the minimal index (according to some fixed order of the players) such that $\tau_{i_*} = \tau$. At stage τ , players $-i_*$ switch to a stationary minmaxing profile against i_* . The definition of σ^* beyond block K_0 , in the event where no deviation has been detected, is irrelevant.

Eq. (6) will ensure that under σ^* the play is absorbed with probability close to 1. Eq. (7) will ensure that the fact that $w^{(k)}$ is not necessarily equal to $f(w^{(k+1)})$ does not distort the players' payoffs too much.

4.5 σ^* is an undiscounted 7ε equilibrium

For $k = 0, \dots, K_0$, we define

$$Z^{(k)} := \begin{cases} r(a^\theta), & \text{if } \theta \leq \min\{N^{(k-1)}, \tau\}, \\ w^{(k)}, & \text{if } N^{(k-1)} < \min\{\theta, \tau\}, \\ v, & \text{if } \tau \leq \min\{N^{(k-1)}, \theta - 1\}, \end{cases}$$

(with $N^{(-1)} = 0$). The random variable $Z^{(k)}$ can be thought of as the expected undiscounted payoff, conditional on the history up to block k . If the game has

absorbed, the undiscounted payoff is $r(a^\theta)$; the undiscounted payoff is otherwise $w^{(k)}$ if no deviator has been identified.

Denote by $\mathcal{H}^{(k)}$ the σ -algebra induced by histories up to the beginning of block k . Since the difference between $f(w^{(k+1)})$ and $\mathbf{E}_{\sigma^{(k)}}[w^{(k+1)}]$ stems from the stopping times $(\tau_i^{(k)})$, the following statement is a rewriting of the properties of $\sigma^{(k)}$ and $\tau^{(k)}$.

Lemma 4.2 *Under σ^* , one has with probability 1,*

$$\left\| \mathbf{E}_{\sigma^*} \left[Z^{(k+1)} \mid \mathcal{H}^{(k)} \right] - Z^{(k)} \right\|_\infty \leq \eta^{(k)} + \|w^{(k)} - f(w^{(k+1)})\|_\infty, \quad \forall k = 0, \dots, K_0 - 1. \quad (8)$$

Moreover, for each player $i \in I$ and each alternative strategy $\tilde{\sigma}_i$,

$$\begin{aligned} \mathbf{E}_{\tilde{\sigma}_i, \sigma_{-i}^*} \left[Z_i^{(k+1)} \mid \mathcal{H}^{(k)} \right] &\leq Z_i^{(k)} + \varepsilon \mathbf{P}_{\tilde{\sigma}_i, \sigma_{-i}^*} \left(N^{(k-1)} \leq \tau_i < N^{(k)} \mid \mathcal{H}^{(k)} \right) \\ &\quad + \eta^{(k)} + \left| w_i^{(k)} - f_i(w^{(k+1)}) \right|. \end{aligned} \quad (9)$$

The bound (8) implies that $Z^{(0)}$ is approximately equal to $\gamma(\sigma^*)$. Indeed, taking first expectations and then using the triangular inequality, one obtains

$$\begin{aligned} \left\| Z^{(0)} - \mathbf{E}_{\sigma^*} \left[Z^{(K_0)} \right] \right\|_\infty &\leq \sum_{k=0}^{K_0-1} \eta^{(k)} + \sum_{k=0}^{K_0-1} \|w^{(k)} - f(w^{(k+1)})\|_\infty \\ &\leq \varepsilon + \delta \leq 2\varepsilon. \end{aligned} \quad (10)$$

On the other hand, since the function $\ln(1-x)$ is concave on $[0, 1)$, Eq. (6) implies that the probability that the game terminates under σ^* in the first K_0 blocks is at least $1 - \exp(-1/\delta) > 1 - \varepsilon$. Since the probability that some player is incorrectly labeled as a deviator is at most ε , this implies that

$$\left\| \mathbf{E}_{\sigma^*} \left[Z^{(K_0)} \right] - \gamma(\sigma^*) \right\|_\infty \leq \varepsilon + \varepsilon = 2\varepsilon. \quad (11)$$

We next show, using Eq. (9), that no player can improve significantly upon σ^* . Let $i \in I$ and $\sigma_i \in \Sigma_i$ be arbitrary. Taking expectations and summing over k in Eq. (9), one has

$$\mathbf{E}_{\sigma_i, \sigma_{-i}^*} \left[Z_i^{(K_0)} \right] \leq Z_i^{(0)} + \varepsilon + \sum_{k=0}^{K_0-1} \eta^{(k)} + \sum_{k=0}^{K_0-1} \left| w_i^{(k)} - f_i(w^{(k+1)}) \right| \leq Z_i^{(0)} + 3\varepsilon. \quad (12)$$

Observe next that the expected payoff induced by $(\sigma_i, \sigma_{-i}^*)$ after the K_0 -th block is at most 1 in the event where $N^{(K_0)} < \min\{\theta, \tau\}$, which implies that $\gamma_i(\sigma_i, \sigma_{-i}^*) \leq \mathbf{E}_{\sigma_i, \sigma_{-i}^*} \left[Z_i^{(K_0)} \right]$. By (12), (10), and (11),

$$\gamma_i(\sigma_i, \sigma_{-i}^*) \leq Z_i^{(0)} + 3\varepsilon \leq \mathbf{E}_{\sigma^*} \left[Z^{(K_0)} \right] + 5\varepsilon \leq \gamma_i(\sigma^*) + 7\varepsilon.$$

Thus, σ^* is an undiscounted 7ε -equilibrium.

5 Discussion and Open Problems

As mentioned before, Solan and Solan (2021) proved that an undiscounted equilibrium payoff exists when (i) there is exactly one connected component which is rectangular, (ii) each player has exactly one action that is not a part of any nonabsorbing entry, and (iii) the probability of absorption in all entries is either 0 or 1. A natural question is whether this result can be combined with ours, to prove that every positive recursive absorbing game that does not have a rectangular connected component B^ℓ in which $|B_i^\ell| \geq 2$ for at most one player admits an undiscounted equilibrium payoff. At present, we do not know how to prove this result. The reason is that the proof approach of Solan and Solan (2021) and our approach are different, and it is not clear how to unify these approaches.

A second open problem concerns the extension of Theorem 2.8 to recursive absorbing games that are not positive and to absorbing games that are not recursive. These extensions pose difficulties because, when the game is not positive recursive, $\rho_i(x)$ may be finite but not higher than player i 's minmax value v_i . In that case, the cutoff for defining the sets $W(x)$, $W_H(x)$, and $W_L(x)$ should be $\max\{\rho_i(x), v_i\}$, and then it is not clear how to define f when $w \in W(x)$ and the only player i for whom $w_i = \max\{\rho_i(x), v_i\}$ satisfies $w_i = v_i > \rho_i(x)$.

It would also be interesting to extend Theorem 2.8 to positive recursive games, which may include more than a single nonabsorbing state.

References

- Attia, L. and M. Oliu-Barton (2019). A formula for the value of a stochastic game. *Proceedings of the National Academy of Sciences* 116(52), 26435–26443.
- Flesch, J., G. Schoenmakers, and K. Vrieze (2008). Stochastic games on a product state space. *Mathematics of Operations Research* 33, 403–420.
- Flesch, J., G. Schoenmakers, and K. Vrieze (2009). Stochastic games on a product state space: the periodic case. *International Journal of Game Theory* 38(2), 263–289.
- Flesch, J., F. Thuijsman, and O. Vrieze (1996). Recursive repeated games with absorbing states. *Mathematics of Operations Research* 21(4), 1016–1022.
- Flesch, J., F. Thuijsman, and O. J. Vrieze (2007). Stochastic games with additive transitions. *European Journal of Operational Research* 179(2), 483–497.
- Kohlberg, E. (1974). Repeated games with absorbing states. *The Annals of Statistics*, 724–738.
- Laraki, R. (2010). Explicit formulas for repeated games with absorbing states. *International Journal of Game Theory* 39(1-2), 53–69.

- Mertens, J.-F. and A. Neyman (1981). Stochastic games. *International Journal of Game Theory* 10(2), 53–66.
- Mertens, J.-F., S. Sorin, and S. Zamir (2015). *Repeated games*. Cambridge University Press.
- Neyman, A. (2003). Stochastic games: Existence of the minmax. In *Stochastic games and applications*, pp. 173–193. Springer.
- Ragel, T. (2024). Weak approachability of convex sets in absorbing games. *Mathematics of Operations Research* 49(3), 1372–1402.
- Simon, R. S. (2007). The structure of non-zero-sum stochastic games. *Advances in Applied Mathematics* 38(1), 1–26.
- Simon, R. S. (2012). A topological approach to quitting games. *Mathematics of Operations Research* 37(1), 180–195.
- Solan, E. (1999). Three-player absorbing games. *Mathematics of Operations Research* 24, 669–698.
- Solan, E. (2000). Absorbing team games. *Games and Economic Behavior* 31, 245–261.
- Solan, E. (2022). *A course in stochastic game theory*. Cambridge University Press.
- Solan, E. and O. N. Solan (2020). Quitting games and linear complementarity problems. *Mathematics of Operations Research* 45(2), 434–454.
- Solan, E. and O. N. Solan (2021). Sunspot equilibrium in positive recursive general quitting games. *International Journal of Game Theory* 50, 891–909.
- Solan, E., O. N. Solan, and R. Solan (2020). Jointly controlled lotteries with biased coins. *Games and Economic Behavior* 119, 383–391.
- Solan, E. and N. Vieille (2001). Quitting games. *Mathematics of Operations Research* 26, 265–285.
- Solan, E. and N. Vieille (2025). Public correlated equilibrium in positive recursive games. *mimeo*.
- Solan, E. and R. Vohra (2002). Correlated equilibrium payoffs and public signalling in absorbing games. *International Journal of Game Theory* 31, 91–121.
- Vieille, N. (2000a). Two-player stochastic games i: A reduction. *Israel Journal of Mathematics* 119(1), 55–91.
- Vieille, N. (2000b). Two-player stochastic games ii: The case of recursive games. *Israel Journal of Mathematics* 119(1), 93–126.
- Vrieze, O. J. and F. Thuijsman (1989). On equilibria in repeated games with absorbing states. *International Journal of Game Theory* 18(3), 293–310.