

Regularity of the minmax value and equilibria in multiplayer Blackwell games*

Galit Ashkenazi-Golan^a, János Flesch^b, Arkadi Predtetchinski^c, and Eilon Solan^d

^aLondon School of Economics and Political Science, Houghton Street London WC2A 2AE, UK

^bDepartment of Quantitative Economics, Maastricht University, P.O.Box 616, 6200 MD, The Netherlands

^cDepartment of Economics, Maastricht University, P.O.Box 616, 6200 MD, The Netherlands

^dSchool of Mathematical Sciences, Tel-Aviv University, Tel-Aviv, Israel, 6997800

October 17, 2022

Abstract

A real-valued function φ that is defined over all Borel sets of a topological space is *regular* if for every Borel set W , $\varphi(W)$ is the supremum of $\varphi(C)$, over all closed sets C that are contained in W , and the infimum of $\varphi(O)$, over all open sets O that contain W .

We study Blackwell games with finitely many players. We show that when each player has a countable set of actions and the objective of a certain player is represented by a Borel winning set, that player's minmax value is regular.

We then use the regularity of the minmax value to establish the existence of ε -equilibria in two distinct classes of Blackwell games. One is the class of n -player Blackwell games where each player has a finite action space and an analytic winning set, and the sum of the minmax values over the players exceeds $n - 1$. The other class is that of Blackwell games with bounded upper semi-analytic payoff functions, history-independent finite action spaces, and history-independent minmax values. For the latter class, we obtain a characterization of the set of equilibrium payoffs.

Keywords: Blackwell games, determinacy, value, equilibrium, regularity.

AMS classification code: Primary: 91A44 (Games involving topology, set theory, or logic). Secondary: 91A20 (Multistage and repeated games).

* Ashkenazi-Golan acknowledges the support of the Israel Science Foundation, grants #217/17 and #722/18, and the NSFC-ISF Grant #2510/17. Solan acknowledges the support of the Israel Science Foundation, grant #217/17. This work has been partly supported by COST Action CA16228 European Network for Game Theory.

1 Introduction

Blackwell games (Blackwell [7]) are dynamic multiplayer simultaneous-move games where the action sets of the players may be history dependent, and the payoff function is an arbitrary Borel-measurable function of the play. When the payoff function of a player is given by the characteristic function of a given set W , we say that W is the *winning set* of the player. These games subsume several familiar classes of dynamic games: repeated games with the discounted payoff or the limiting average payoff (e.g., Sorin [50], Mailath and Samuelson [31]), games with perfect information (e.g., Gale and Stewart [20]), and graph games arising in the computer science applications (e.g., Apt and Grädel [1], Bruyère [9, 10], Chatterjee and Henzinger [14]).

While two-player zero-sum Blackwell games and Blackwell games with perfect information are quite well understood (see, e.g., Martin [35, 36], Mertens [39], Kuipers, Flesch, Schoenmakers, and Vrieze [26]), general multiplayer nonzero-sum Blackwell games have so far received relatively little attention.

The goal of this paper is to introduce a new technique to the study of multiplayer Blackwell games: regularity of the minmax value, along with a number of related approximation results. In a nutshell, the technique amounts to the approximation of the minmax value of a winning Borel set using a closed subset. This approach allows us to establish existence of ε -equilibria in two distinct classes of Blackwell games.

REGULARITY AND APPROXIMATION RESULTS: A real-valued function φ that is defined over all Borel sets of a certain space is *inner regular* if for every Borel set W , $\varphi(W)$ is the supremum of $\varphi(C)$, over all closed sets C that are contained in W . The function φ is *outer regular* if for every Borel set W it is the infimum of $\varphi(O)$, over all open sets O that contain W . The function φ is *regular* if it is both inner regular and outer regular. Borel probability measures on metric spaces are one example of a regular function (see, e.g., Kechris [25, Theorems 17.10 and 17.11]).

When restricted to two-player zero-sum Blackwell games with finite action sets and Borel-measurable winning set for Player 1, the value function is known to be regular (Martin [36]). This result was extended to two-player zero-sum stochastic games by Maitra, Purves, and Sudderth [32].

We show that in multiplayer Blackwell games with countable action sets and Borel winning sets, the minmax value of all players is regular. We thus extend the regularity result of Martin [36] in terms of both the number of actions (countable versus finite) and the number of players (finite versus two).

A related approximation result concerns the case when a player's objective is represented by a bounded Borel-measurable payoff function. Denote by $v_i(f)$ player i 's minmax value when her payoff function is f . We show that $v_i(f)$ is the supremum of $v_i(g)$ over all bounded limsup functions $g \leq f$, and the infimum of $v_i(g)$ over all bounded limsup function $g \geq f$. A *limsup function* is a function that can be written as the limit superior of a sequence of rewards assigned to the nodes of the game tree. This too, is an extension of results by Maitra, Purves, and Sudderth [32] and Martin [36] for two-player games to multiplayer games. If, moreover, the player's minmax value is the same in every subgame, one obtains an approximation from below by an upper semi-continuous function, and an approximation from above by a lower semi-continuous function.

EXISTENCE OF ε -EQUILIBRIA: The main contribution of the paper is the application of the regularity of the minmax value to the problem of existence of an ε -equilibrium in multiplayer Blackwell games. We establish the existence in two distinct classes of Blackwell games.

One is the class of n -player Blackwell games with bounded upper semi-analytic payoff functions, history-independent finite action spaces, and history-independent minmax values. The latter assumption means that every player's minmax value is the same in each subgame. Under these assumptions, for each $\varepsilon > 0$, there is an ε -equilibrium with a pure path of play.

A prominent sufficient condition for the minmax value to be history-independent is that the payoff be tail-measurable. Roughly speaking, tail-measurability amounts to the requirement that the payoff is unaffected by a change of the action profile in any finite number of stages. We thus obtain the existence of ε -equilibria in Blackwell games with history-independent finite action spaces and bounded, upper semi-analytic, and tail-measurable payoff functions.

The second class of games for which we derive an existence result is n -player Blackwell games where each player has a finite action space at each history, her objective is represented by an analytic winning set, and the sum of the minmax values over the players exceeds $n - 1$. Under these conditions we show that there exists a play that belongs to each player's winning set; any such play induces a 0-equilibrium. At the heart of the proof is an approximation of each player's minmax value by her minmax value in an auxiliary game, where the winning set is a closed subset of the player's winning set in the original game.

The key idea of the proof of the first result is to consider an auxiliary Blackwell game with winning sets, where the winning set of player i is the set of player i 's ε -individually rational plays: the plays that yield player i a payoff no smaller than her minmax value minus ε . We show that, in the thus-defined auxiliary Blackwell game, each player's minmax value equals 1, and apply the second result.

The question whether ε -equilibria exist in multiplayer Blackwell games is a largely uncharted territory. An important benchmark is the result of Mertens and Neyman (see Mertens [39]): all games of perfect information with bounded Borel-measurable payoff functions admit an ε -equilibrium for every $\varepsilon > 0$. Zero-sum Blackwell games (where at least one of the two players has a finite set of actions) are known to be determined since the seminal work of Martin [36]. Shmaya [45] extends the latter result by showing the determinacy of zero-sum games with eventual perfect monitoring, and Arieli and Levy [2] extend Shmaya's result to stochastic signals.

Only some special classes of multiplayer dynamic games have been shown to have an ε -equilibrium. These include stochastic games with discounted payoffs (see, e.g., the survey by Jaśkiewicz and Nowak [27]), two-player stochastic games with the limiting average payoff (Vieille [52, 53]), and graph games with classical computer science objectives (e.g., Secchi and Sudderth [44], Chatterjee [15, 12], Bruyère [10], Ummels, Markey, Brenguier, and Bouyer [51]).

Both Shmaya's ([45]) and Arieli and Levy's ([2]) arguments rely on the regularity of the value in zero-sum games as established by Maitra, Purves, and Sudderth [32] and Martin [36]. A different but related argument invoking regularity of a given measure appears in various other papers; here we refer to Simon [46] and to Simon and Tomkowicz [47], who use regularity of a measure induced by a given strategy profile to derive existence of equilibria in a class of quitting games, and to construct a Bayesian game having no equilibria, respectively.

A companion paper ([3]) establishes the existence of ε -equilibria in Blackwell games with countably many players, finite action sets, and bounded, Borel-measurable, and tail-measurable

payoff functions. The present paper departs from [3] in two dimensions. Firstly, it invokes a new proof technique, the regularity of the minmax value. Secondly, it makes different assumptions on the primitives. The second of our two existence results (Theorem 4.1) has, in fact, no analogue in [3]. The first (Theorem 4.3) applies to a larger class of payoff functions than does the main result in [3]: it only requires players' minmax values to be history-independent. While tail-measurability of the payoff functions is a sufficient condition for history-independence of the minmax values, it is by no means a necessary condition. Furthermore, Borel-measurability imposed in [3] is relaxed here to upper semi-analyticity. On the other hand, [3] has a countable rather than a finite set of players, something that the methods developed here do not allow for.

CHARACTERISATION OF EQUILIBRIUM PAYOFFS: An equilibrium payoff is an accumulation point of the expected payoff vectors of ε -equilibria, as ε tends to 0. We establish a characterisation of equilibrium payoffs in games with bounded upper semi-analytic payoff functions, history-independent finite action spaces, and history-independent minmax values.

In repeated games with patient players, the folk theorem asserts that under proper conditions, the set of limiting average equilibrium payoffs (or the limit set of equilibrium payoffs, as the discount factor goes to 0 or the horizon increases to infinity) is the set of all vectors that are individually rational and lie in the convex hull of the range of the stage payoff function (see, e.g., Aumann and Shapley [5], Sorin [50], or Mailath and Samuelson [31]). Our result identifies the set of equilibrium payoffs of a Blackwell game as the set of all vectors that lie in the convex hull of the set of feasible and individually rational payoffs. The intuition for this discrepancy is that in standard repeated games, a low payoff in one stage can be compensated by a high payoff in another stage, therefore payoff vectors that are convex combinations of the stage payoff function can be equilibrium payoffs as long as this convex combination of payoffs is individually rational. In particular, these combinations can place some positive weight on payoff vectors that are not individually rational. In Blackwell games, however, the payoff is obtained only at the end of the game, hence only plays that generate individually rational payoffs can be taken into account when constructing equilibria.

Our characterization of the set of equilibrium payoffs is related to the rich literature on the folk theorem, and the study of the minmax value is instrumental to this characterizaion (see, e.g., the folk theorems in Fudenberg and Maskin [19], Mailath and Samuelson [31], or Hörner, Sugaya, Takahashi, and Vieille [24]). The minmax value of a player would often be used in the proofs of equilibrium existence to construct suitable punishments for a deviation from the supposed equilibrium play (as is done, for instance, in Aumann and Shapley [5], Rubinstein [43], Fudenberg and Maskin [19], and Solan [48]).

The paper is structured as follows. Section 2 describes the class of Blackwell games. Section 3 is devoted to the regularity of the minmax value and related approximation theorems. Section 4 applies these tools to the problem of existence of equilibrium. Section 5 is devoted to the characterisation of equilibrium payoffs. Section 6 discusses the implications of the results for games with tail-measurable payoffs. Section 7 contains a discussion, concluding remarks, and open questions.

2 Blackwell games

Blackwell games: An n -player Blackwell game is a tuple $\Gamma = (I, A, H, (f_i)_{i \in I})$. The elements of Γ are as follows.

The set of players is I , a finite set of cardinality n . For a player $i \in I$ we write $-i$ to denote the set of i 's opponents, $I \setminus \{i\}$.

The set A is a countable set and $H \subseteq \cup_{t \in \mathbb{N}} A^t$ is the game tree (throughout the paper $\mathbb{N} = \{0, 1, \dots\}$). Elements of H are called histories. The set H is assumed to have the following properties: (a) H contains the empty sequence, denoted \emptyset ; (b) a prefix of an element of H is an element of H ; that is, if for some $h \in \cup_{t \in \mathbb{N}} A^t$ and $a \in A$ the sequence (h, a) is an element of H , so is h ; (c) for each $h \in H$ there is an element $a \in A$ such that $(h, a) \in H$; we define $A(h) := \{a \in A : (h, a) \in H\}$; and (d) for each $h \in H$ and each $i \in I$ there exists a set $A_i(h)$ such that $A(h) = \prod_{i \in I} A_i(h)$. The set $A_i(h)$ is called player i 's set of actions at history h , and $A(h)$ the set of action profiles at h .

Conditions (a), (b), and (c) above say that H is a pruned tree on A . Condition (c) implies that the game has infinite horizon. Let $H_t := H \cap A^t$ denote the set of histories in stage t .

An infinite sequence $(a_0, a_1, \dots) \in A^{\mathbb{N}}$ such that $(a_0, \dots, a_t) \in H$ for each $t \in \mathbb{N}$ will be called a *play*. The set of plays is denoted by $[H]$. This is the set of infinite branches of H . For $h \in H$ let $O(h)$ denote the set of all plays of Γ having h as a prefix. We endow $[H]$ with the topology generated by the basis consisting of the sets $\{O(h) : h \in H\}$. The space $[H]$ is Polish. For $t \in \mathbb{N}$ let \mathcal{F}_t be the sigma-algebra on $[H]$ generated by the sets $\{O(h) : h \in H_t\}$. The Borel sigma-algebra of $[H]$ is denoted by \mathcal{B} . It is the minimal sigma-algebra containing the topology. A subset S of $[H]$ is *analytic* if it is the image of a continuous function from the Baire space $\mathbb{N}^{\mathbb{N}}$ to $[H]$. Each Borel set is analytic.

Each analytic set is universally measurable. Recall that a set $S \subseteq [H]$ is said to be *universally measurable* if (Kechris [25, Section 17.A]), for every Borel probability measure \mathbb{P} on $[H]$, there exist Borel sets $B, Z \in \mathcal{B}$ such that $S \Delta B \subseteq Z$ and $\mathbb{P}(Z) = 0$; here $S \Delta B = (S \setminus B) \cup (B \setminus S)$ is the symmetric difference of the sets S and B .

The last element of the game is a vector $(f_i)_{i \in I}$, where $f_i : [H] \rightarrow \mathbb{R}$ is player i 's *payoff function*. The most general class of payoff functions we allow for are bounded upper semi-analytic functions. A function $f_i : [H] \rightarrow \mathbb{R}$ is said to be *upper semi-analytic* if, for each $r \in \mathbb{R}$, the set $\{p \in [H] : r \leq f_i(p)\}$ is analytic. In particular, the indicator function 1_S of a subset $S \subseteq [H]$ of plays is upper semi-analytic if and only if S is an analytic set. Each Borel-measurable function is upper semi-analytic. Note that a bounded upper semi-analytic function is universally measurable, i.e., for each open set $U \subseteq \mathbb{R}$, the set $f_i^{-1}(U) \subseteq [H]$ is universally measurable (see, e.g., Chapter 7 in Bertsekas and Shreve [6]).

The play of the game starts at the empty history $h_0 = \emptyset$. Suppose that by a certain stage $t \in \mathbb{N}$ a history $h_t \in H_t$ has been reached. Then in stage t , the players simultaneously choose their respective actions; thus player $i \in I$ chooses an action $a_{i,t} \in A_i(h_t)$. This results in the stage t action profile $a_t = (a_{i,t})_{i \in I} \in A(h_t)$. Once chosen, the actions are revealed to all players, and the history $h_{t+1} = (h_t, a_t)$ is reached. The result of the infinite sequence of choices is the play $p = (a_0, a_1, \dots)$, an element of $[H]$. Each player $i \in I$ receives the corresponding payoff $f_i(p)$.

Given a Blackwell game Γ and a history $h \in H$, the *subgame* of Γ starting at h is the Blackwell game $\Gamma_h = (I, A, H_h, (f_{i,h})_{i \in I})$. The set H_h of histories of Γ_h consists of finite sequences $g \in$

$\bigcup_{t \in \mathbb{N}} A^t$ such that $hg \in H$, where hg is the concatenation of h and g . The payoff function $f_{i,h} : [H_h] \rightarrow \mathbb{R}$ is the composition $f_i \circ s_h$, with $s_h : [H_h] \rightarrow [H]$ given by $p \mapsto hp$, where hp is the concatenation of h and p . Note that Γ_\emptyset is just the game Γ itself.

The Blackwell game Γ is said to have *history-independent action sets* if $A_i(h) = A_i(\emptyset)$ for each history $h \in H$ and each player $i \in I$; the common action set is simply denoted by A_i . If Γ has history-independent action sets, then the set of its histories is $H = \bigcup_{t \in \mathbb{N}} A^t$, and the set of plays in Γ is $[H] = A^{\mathbb{N}}$. A Blackwell game with history-independent action sets can be described as a tuple $(I, (A_i, f_i)_{i \in I})$.

Strategies and expected payoffs: A strategy for player $i \in I$ is a function σ_i assigning to each history $h \in H$ a probability distribution $\sigma_i(h)$ on the set $A_i(h)$. The set of player i 's strategies is denoted by Σ_i . We also let $\Sigma_{-i} := \prod_{j \in -i} \Sigma_j$ and $\Sigma := \prod_{i \in I} \Sigma_i$. Each strategy profile $\sigma = (\sigma_i)_{i \in I}$ induces a unique probability measure on the Borel sets of $[H]$, denoted \mathbb{P}_σ . The corresponding expectation operator is denoted \mathbb{E}_σ . In particular, $\mathbb{E}_\sigma(f_i)$ denotes an expected payoff to player i in the Blackwell game under the strategy profile σ . It is well defined under the maintained assumptions, namely boundedness and upper semi-analyticity of f_i .

Take a history $h \in H_t$ in stage t . A strategy profile $\sigma \in \Sigma$ in Γ induces the strategy profile σ_h in Γ_h defined as $\sigma_h(g) = \sigma(hg)$ for each history $g \in H_h$. Let us define $\mathbb{E}_\sigma(f_i | h)$ as the expected payoff to player i in the Blackwell game Γ_h under the strategy profile σ_h : that is, $\mathbb{E}_\sigma(f_i | h) := \mathbb{E}_{\sigma_h}(f_{i,h})$. Note that $\mathbb{E}_\sigma(f_i | h)$, when viewed as an \mathcal{F}_t -measurable function on $[H]$, is a conditional expectation of f_i with respect to the measure \mathbb{P}_σ and the sigma-algebra \mathcal{F}_t ; whence our choice of notation.

Minmax value: Consider a Blackwell game Γ , and suppose that player i 's payoff function f_i is bounded and upper semi-analytic. Player i 's *minmax value* is defined as

$$v_i(f_i) := \inf_{\sigma_{-i} \in \Sigma_{-i}} \sup_{\sigma_i \in \Sigma_i} \mathbb{E}_{\sigma_{-i}, \sigma_i}(f_i).$$

Whenever $f_i = 1_{W_i}$ is an indicator of an analytic set $W_i \subseteq [H]$ we write $v_i(W_i)$ for $v_i(1_{W_i})$.

Player i 's minmax value is said to be *history-independent* if her minmax value in the subgame Γ_h equals that in the game Γ , for each history $h \in H$.

3 Regularity and approximation theorems

In this section we state the regularity property of the minmax: the minmax value of a Borel winning set can be approximated from below by the minmax value of closed subset and from above by the minmax value of an open superset. We also describe two related approximation results: the minmax value of a bounded Borel-measurable payoff function can be approximated from below and from above by limsup functions. If, in addition, the minmax values are history-independent, then one can choose the approximation from below to be upper semicontinuous, and the approximation from above to be lower semicontinuous. The proofs of all results are detailed in the appendix.

Theorem 3.1. (Regularity of the minmax value) *Consider a Blackwell game. Suppose that player i 's objective is given by a winning set $W_i \subseteq [H]$. Suppose that W_i is Borel. Then*

$$v_i(W_i) = \sup\{v_i(C) : C \subseteq W_i, C \text{ is closed}\} = \inf\{v_i(O) : O \supseteq W_i, O \text{ is open}\}.$$

One implication of Theorem 3.1 concerns the complexity of strategies of player i that ensures that her probability of winning is close to her minmax value. Suppose, for example, that $v_i(W_i) = \frac{1}{2}$. Then for every strategy profile σ_{-i} of the opponents of player i and every $\varepsilon > 0$, she has a response σ_i such that $\mathbb{P}_{\sigma_{-i}, \sigma_i}(W_i) \geq \frac{1}{2} - \varepsilon$. The strategy profile σ_{-i} and the winning set W_i may be complex, and accordingly the good response σ_i may be complex as well. However, take now a closed subset $C \subseteq W_i$ such that $v_i(C) > v_i(W_i) - \varepsilon = \frac{1}{2} - \varepsilon$. The complement of C , denoted C^c , is open, hence it is the union of basic open sets; that is, it can be presented as a union $C^c = \bigcup_{h \in H'} O(h)$, for some subset $H' \subseteq H$ of histories. A strategy σ'_i that satisfies $\mathbb{P}_{\sigma_{-i}, \sigma'_i}(C_i) \geq \frac{1}{2} - \varepsilon$ must aim at avoiding C^c , that is, at avoiding histories in H' . In that sense, σ' may have a simple structure.

Example 3.2. Here we consider a Blackwell game where the same stage game is being played at every stage. The stage game specifies a stage winning set for each player. A player's objective in the Blackwell game is to win the stage game infinitely often.

Thus let $\Gamma = (I, (A_i, 1_{W_i})_{i \in I})$ be a Blackwell game with history-independent countable action sets, where player i 's winning set is

$$W_i = \{(a_0, a_1, \dots) \in A^{\mathbb{N}} : a_t \in U_i \text{ for infinitely many } t \in \mathbb{N}\};$$

here U_i , called player i 's *stage winning set*, is a given subset of $\prod_{i \in I} A_i$. If $a_t \in U_i$, we say that player i *wins stage* t . The set W_i is a G_δ -set, i.e., an intersection of countably many open subsets of $A^{\mathbb{N}}$.

Fix a player $i \in I$. Let

$$d_i := \inf_{x_{-i} \in X_{-i}} \sup_{x_i \in X_i} \mathbb{P}_{x_{-i}, x_i}(U_i) \quad (1)$$

be player i 's minmax value in the stage game. As follows from the arguments below, $v_i(W_i)$ is either 0 or 1, and it is 1 exactly when $d_i > 0$. In either case, there are intuitive approximations of player i 's winning sets by a closed set from below and an open set from above.

First assume that $d_i > 0$. Take an $\varepsilon > 0$. Let us imagine that player i 's objective is not merely to win infinitely many stages in the course of the Blackwell game, but to make sure that she wins at least once in every block of stages $t_n, \dots, t_{n+1} - 1$, where the sequence of stages $t_0 < t_1 < \dots$ is chosen to satisfy

$$(1 - \frac{1}{2}d_i)^{t_{n+1} - t_n} < 2^{-n-1} \cdot \varepsilon$$

for each $n \in \mathbb{N}$. This, more demanding condition, defines an approximating set. Formally, define

$$C_i := \bigcap_{n \in \mathbb{N}} \bigcup_{t_n \leq k < t_{n+1}} \{(a_0, a_1, \dots) \in A^{\mathbb{N}} : a_k \in U_i\}.$$

As the intersection of closed sets, C_i is a closed subset of W_i . Moreover, $1 - \varepsilon \leq v_i(C_i)$. To see this, fix any strategy σ_{-i} for i 's opponents. At any history h , player i has a mixed action $\sigma_i(h)$ that, when played against $\sigma_{-i}(h)$, guarantees a win at history h with probability of at least $\frac{1}{2}d_i$. Thus, under the measure $\mathbb{P}_{\sigma_{-i}, \sigma_i}$ the probability for player i not to win at least once in a block of stages $t_n, \dots, t_{n+1} - 1$ is at most $2^{-n-1} \cdot \varepsilon$, for any history of play up to stage t_n . And hence the probability that there is a block within which player i does not win once is at most ε .

Suppose that $d_i = 0$. Let us imagine that player i 's objective is merely to win the stage game at least once. This modest objective defines the approximating set:

$$O_i = \bigcup_{t \in \mathbb{N}} \{(a_0, a_1, \dots) \in A^{\mathbb{N}} : a_t \in U_i\}.$$

As the union of open sets, O_i is an open set containing W_i . Moreover, $v_i(O_i) \leq \varepsilon$. To see this, let σ_{-i} be the strategy for i 's opponents such that, at any stage $t \in \mathbb{N}$ and any history $h \in A^t$ of stage t , the probability that the action profile a_t is an element of U_i is not greater than $2^{-t-1} \cdot \varepsilon$ regardless of the action of player i . Then, for any player i 's strategy σ_i , the probability that i wins at least once is not greater than ε . \square

We turn to two related approximation results for Blackwell games with Borel payoff functions. A function $f : [H] \rightarrow \mathbb{R}$ is said to be a *limsup function* if there exists a function $u : H \rightarrow \mathbb{R}$ such that for each play $(a_0, a_1, \dots) \in [H]$,

$$f(a_0, a_1, \dots) = \limsup_{t \rightarrow \infty} u(a_0, \dots, a_t).$$

The function $f : [H] \rightarrow \mathbb{R}$ is a *liminf function* if $-f$ is a limsup function.

Limsup and liminf payoff functions are ubiquitous in the literature on infinite dynamic games. At least since the work of Gillette [21], the so-called limiting average payoff (that is, the limit superior or the limit inferior of the average of the stage payoffs) is a standard specification of the payoffs in a stochastic game (see for example Mertens and Neyman [40], or Levy and Solan [30]). Stochastic games with limsup payoff functions have been studied in Maitra and Sudderth [33].

Limsup functions have relatively “low” set-theoretic complexity. Various characterizations of the limsup functions can be found in Hausdorff [23]. In particular, f is a limsup function if and only if, for each $r \in \mathbb{R}$, the set $\{p \in [H] : r \leq f(p)\}$ is a G_δ -set.

We now state a result on the approximation of the minmax value for Blackwell games where a player's objective is represented by a bounded Borel-measurable payoff function.

Theorem 3.3. *Consider a Blackwell game. Suppose that player i 's payoff function $f_i : [H] \rightarrow \mathbb{R}$ is bounded and Borel-measurable. Then:*

$$\begin{aligned} v_i(f_i) &= \sup\{v_i(g) : g \text{ is a bounded limsup function and } g \leq f_i\} \\ &= \inf\{v_i(g) : g \text{ is a bounded limsup function and } f_i \leq g\}. \end{aligned}$$

Theorems 3.1 and 3.3 have been proven by Martin [36] for the case $n = 2$, see [36, Theorem 5, and Remark (b)] and [36, Remark (c)]. They have been extended to two-player stochastic games by Maitra and Sudderth [34]. Theorems 3.1 and 3.3 extend the known results in two respects. First, they allow for more than two players, and second, they allow for countably many actions.

The proof of Theorem 3.1 combines and fine-tunes the arguments in Martin [36] and Maitra and Sudderth [34]. The key element of the proof is a zero-sum perfect information game, denoted $G_i(f_i, c)$, where the aim of Player I is to “prove” that the minmax value of f_i is at least c . Roughly speaking, the game proceeds as follows. Player I commences the game by proposing a fictitious continuation payoff, which one could think of as a payoff player i hopes to attain,

contingent on each possible stage 0 action profile. The number c serves as the initial threshold: player i 's minmax value of the proposed continuation payoffs is required to be at least c . Player II then chooses a stage 0 action profile, and the corresponding continuation payoff serves as the new threshold. Player I then proposes a fictitious continuation payoff contingent on each possible stage 1 action profile, and Player II chooses the stage 1 action profile, etc. Player I wins if the sequence of continuation payoffs is “justified” by the actual payoff on a play produced by Player II. Ultimately the proof rests on the determinacy of the game $G_i(f_i, c)$, which follows by Martin [35].

The perfect information game $G_i(f_i, c)$ is a version of the games used in Martin [36]. The main difference is in the use of player i 's minmax value that constrains Player I's choice of fictitious continuation payoffs. The details of our proof are slightly closer to those in Maitra and Sudderth [34]. Like them we invoke martingale convergence and the Fatou lemma.

Finally, we state an approximation result for a Blackwell game with history-independent minmax values. Recall that a function $g : A^N \rightarrow \mathbb{R}$ is *upper semicontinuous* if, for each $r \in \mathbb{R}$, the set $\{p \in [H] : r \leq f(p)\}$ is a closed set, and g is *lower semicontinuous* if $-g$ is upper semicontinuous. When $g = 1_B$ for some $B \subseteq A^N$, g is upper semicontinuous (resp. lower semicontinuous) if and only if B is closed (resp. open).

Theorem 3.4. *Consider a Blackwell game. Suppose that player i 's payoff function f_i is bounded and Borel-measurable, and player i 's minmax values are history-independent. Then*

$$\begin{aligned} v_i(f_i) &= \sup\{v_i(g) : g \text{ is a bounded upper semicontinuous function and } g \leq f_i\} \\ &= \inf\{v_i(g) : g \text{ is a bounded lower semicontinuous function and } f_i \leq g\}. \end{aligned}$$

As the proof reveals, both the upper semicontinuous and the lower semicontinuous functions can be chosen to be two-valued. Recall that (Hausdorff [23]) an upper semicontinuous function and a lower semicontinuous function are both a limsup and a liminf function. Consequently, in comparison to Theorem 3.3, an additional assumption of history-independence of the minmax values in Theorem 3.4 leads to a stronger approximation result. The latter condition cannot be dropped; see Section 7 for an example of a game with a limsup payoff function such that the minmax value cannot be approximated from below by an upper semicontinuous function.

4 Existence of equilibria

In this section, we employ the results of the previous section to establish existence of ε -equilibria in two distinct classes of Blackwell games. Theorem 4.1 concerns n -player Blackwell games where each player has a finite action space at each history, her objective is represented by an analytic winning set, and the sum of the minmax values over the players exceeds $n - 1$. Theorem 4.3 concerns Blackwell games with bounded upper semi-analytic payoff functions, history-independent finite action spaces, and history-independent minmax values.

Consider a Blackwell game Γ and let $\varepsilon \geq 0$. A strategy profile $\sigma \in \Sigma$ is an ε -equilibrium of Γ if for each player $i \in I$ and each strategy $\eta_i \in \Sigma_i$ of player i ,

$$\mathbb{E}_{\sigma_{-i}, \eta_i}(f_i) \leq \mathbb{E}_{\sigma_{-i}, \sigma_i}(f_i) + \varepsilon.$$

We state our first existence result.

Theorem 4.1. Consider an n -player Blackwell game $\Gamma = (I, A, H, (1_{W_i})_{i \in I})$. Suppose that for each player $i \in I$ player i 's action set $A_i(h)$ at each history $h \in H$ is finite, and that her winning set W_i is analytic. If $v_1(W_1) + \dots + v_n(W_n) > n - 1$, then the set $W_1 \cap \dots \cap W_n$ is not empty. Consequently, Γ has a 0-equilibrium.

Note that any play $p \in W_1 \cap \dots \cap W_n$ induces a 0-equilibrium: any strategy profile that requires all the players to follow p is a 0-equilibrium, because it yields all players the maximal payoff 1.

The key step of the proof is the approximation of the minmax value of a player using a closed subset of her winning set. To prove Theorem 4.1 we need the following technical observation.

Lemma 4.2. Let (X, \mathcal{B}, P) be a probability space, and let $Q_1, \dots, Q_n \in \mathcal{B}$ be n events. Then

$$P(Q_1 \cap \dots \cap Q_n) \geq P(Q_1) + \dots + P(Q_n) - n + 1.$$

Proof. For $n = 1$ the statement is obvious, and for $n = 2$ we have

$$P(Q_1 \cap Q_2) = P(Q_1) + P(Q_2) - P(Q_1 \cup Q_2) \geq P(Q_1) + P(Q_2) - 1. \quad (2)$$

Assume that the statement holds for some $n - 1$. Then for n we have

$$\begin{aligned} P(Q_1 \cap \dots \cap Q_n) &= P((Q_1 \cap \dots \cap Q_{n-1}) \cap Q_n) \\ &\geq P(Q_1 \cap \dots \cap Q_{n-1}) + P(Q_n) - 1 \\ &\geq (P(Q_1) + \dots + P(Q_{n-1}) - n + 2) + P(Q_n) - 1 \\ &= P(Q_1) + \dots + P(Q_n) - n + 1, \end{aligned}$$

where the first inequality follows from Eq. (2) and the second by the induction hypothesis. \square

Proof of Theorem 4.1: We first establish the theorem in the special case of Borel winning sets, and then generalize it to analytic winning sets.

PART I: Suppose that for each $i \in I$ the set $W_i \subseteq [H]$ is Borel.

By Theorem 3.1 there are closed sets $C_1 \subseteq W_1, \dots, C_n \subseteq W_n$ such that $v_1(C_1) + \dots + v_n(C_n) > n - 1$. We show that the intersection $C_1 \cap \dots \cap C_n$ is not empty.

Given $m \in \mathbb{N}$ consider the n -player Blackwell game $\Gamma^m = (I, A, H, (1_{C_i^m})_{i \in I})$, where player i 's winning set is defined by

$$C_i^m := \bigcup \{O(h) : h \in H_m \text{ such that } O(h) \cap C_i \neq \emptyset\}.$$

The game Γ^m essentially ends after m stages: by stage m each player i knows whether the play is an element of her winning set C_i^m or not. In Γ^m , player i wins if after m stages there is a continuation play that leads to C_i . Note that this continuation play might be different for different players.

The set C_i^m is a clopen set. For each $m \in \mathbb{N}$ and $i \in I$ we have the inclusion $C_i^m \supseteq C_i^{m+1}$ (winning in Γ^{m+1} is more difficult than winning in Γ^m). Moreover, $\bigcap_{m \in \mathbb{N}} C_i^m = C_i$. Indeed, the inclusion $C_i^m \supseteq C_i$ is evident from the definition. Conversely, take an element q of the set

$[H] \setminus C_i$. Since $[H] \setminus C_i$ is an open set, there exists a history $h \in H$ such that $q \in O(h)$ and $O(h) \subseteq [H] \setminus C_i$. But then $q \in [H] \setminus C_i^m$, where m is the length of the history h .

Define $C^m := C_1^m \cap \dots \cap C_n^m$. Thus $\{C^m\}_{m \in \mathbb{N}}$ is a nested sequence of closed sets converging to $C_1 \cap \dots \cap C_n$. Note that, since by the assumption of the theorem H is a finitely branching tree, the space $[H]$ is compact. Thus C^m is a compact set. Consequently, to prove that $C_1 \cap \dots \cap C_n$ is not empty, we only need to argue that C^m is not empty for each $m \in \mathbb{N}$.

The game Γ^m being finite, it has a 0-equilibrium (Nash [41]), say σ^m . By the definition of 0-equilibrium, the equilibrium payoff is not less than the minmax value:

$$\mathbb{P}_{\sigma_{-i}, \sigma_i^m}(C_i^m) = \sup_{\sigma_i \in \Sigma_i} \mathbb{P}_{\sigma_{-i}, \sigma_i}(C_i^m) \geq \inf_{\sigma_{-i} \in \Sigma_{-i}} \sup_{\sigma_i \in \Sigma_i} \mathbb{P}_{\sigma_{-i}, \sigma_i}(C_i^m) = v_i(C_i^m).$$

Moreover, since $C_i^m \supseteq C_i$, it holds that $v_i(C_i^m) \geq v_i(C_i)$. We conclude that

$$\mathbb{P}_{\sigma^m}(C_1^m) + \dots + \mathbb{P}_{\sigma^m}(C_n^m) > n - 1.$$

Finally, we apply Lemma 4.2 to conclude that $\mathbb{P}_{\sigma^m}(C^m) > 0$, hence C^m is not empty.

PART II: Now let Γ be any game as in the statement of the theorem. Suppose by way of contradiction that $W_1 \cap \dots \cap W_n$ is empty. By Novikov's separation theorem (Kechris [25, Theorem 28.5]) there exist Borel sets B_1, \dots, B_n such that $W_i \subseteq B_i$ for each $i \in I$ and $B_1 \cap \dots \cap B_n = \emptyset$. But since $v_i(W_i) \leq v_i(B_i)$ for each $i \in I$, the game $\Gamma = (I, A, H, (1_{B_i})_{i \in I})$ satisfies the assumptions of the theorem, and Part I of the proof yields a contradiction. \square

We state our second and main existence result.

Theorem 4.3. *Consider a Blackwell game $\Gamma = (I, A, H, (f_i)_{i \in I})$. Suppose that for each player $i \in I$, player i 's action set $A_i(h)$ at each history $h \in H$ is finite, her payoff function f_i is bounded and upper semi-analytic, and her minmax value is history-independent. Then for every $\varepsilon > 0$ the game admits an ε -equilibrium.*

The key idea behind the proof is to consider an auxiliary Blackwell game with winning sets, the winning set of a player consisting of that player's ε -individually rational plays. We show that in the thus-defined auxiliary Blackwell game each player's minmax value equals 1, and apply Theorem 4.1.

Given $\varepsilon > 0$ we define the set of *player i 's ε -individually rational plays*:

$$Q_{i,\varepsilon}(f_i) := \{p \in [H] : f_i(p) \geq v_i(f_i) - \varepsilon\}.$$

Also define the set

$$U_{i,\varepsilon}(f_i) := \{p \in [H] : f_i(p) \geq v_i(f_i) + \varepsilon\}.$$

Note that under the assumptions of Theorem 4.3 both sets are analytic.

Proposition 4.4. *Consider a Blackwell game $\Gamma = (I, A, H, (f_i)_{i \in I})$ and a player $i \in I$. Suppose that player i 's payoff function f_i is bounded and upper semi-analytic, and that her minmax values are history-independent. Let $\varepsilon > 0$. Then*

1. $v_i(Q_{i,\varepsilon}(f_i)) = 1$. In fact, for each strategy profile $\sigma_{-i} \in \Sigma_{-i}$ of players $-i$ there is a strategy $\sigma_i \in \Sigma_i$ for player i such that $\mathbb{P}_{\sigma_{-i}, \sigma_i}(Q_{i,\varepsilon}(f_i)) = 1$.

2. $v_i(U_{i,\varepsilon}(f_i)) = 0$. In fact, there exists a strategy profile $\sigma_{-i} \in \Sigma_{-i}$ of players $-i$ such that for each strategy $\sigma_i \in \Sigma_i$ for player i it holds that $\mathbb{P}_{\sigma_{-i}, \sigma_i}(U_{i,\varepsilon}(f_i)) = 0$.

Proof. CLAIM 1: It suffices to prove the second statement. Take a strategy profile σ_{-i} of players $-i$. It is known that player i has a strategy σ_i that is an $\varepsilon/2$ -best response to σ_{-i} in each subgame (see, for example, Mashiah-Yaakovi [38, Proposition 11], or Flesch, Herings, Maes, and Predtetchinski [16, Theorem 5.7]), and therefore

$$\mathbb{E}_{\sigma_{-i}, \sigma_i}(f_i | h) \geq v_i(f_{i,h}) - \varepsilon/2 = v_i(f_i) - \varepsilon/2,$$

for each history $h \in H$. Since the payoff function f_i is bounded, it follows that there is $d > 0$ such that

$$\mathbb{P}_{\sigma_{-i}, \sigma_i}(Q_{i,\varepsilon}(f_i) | h) \geq d$$

for each $h \in H$. Indeed, it is easy to verify that one can choose

$$d = \frac{\varepsilon}{2(\sup_{p \in [H]} f_i(p) - v_i(f_i) + \varepsilon)}.$$

Since $Q_{i,\varepsilon}(f_i)$ is an analytic set, there is a Borel set B such that $\mathbb{P}_{\sigma_{-i}, \sigma_i}(Q_{i,\varepsilon}(f_i) \Delta B) = 0$, where Δ stands for the symmetric difference of two sets. It follows that $\mathbb{P}_{\sigma_{-i}, \sigma_i}(Q_{i,\varepsilon}(f_i) \Delta B | h) = 0$, and consequently $\mathbb{P}_{\sigma_{-i}, \sigma_i}(B | h) \geq d$ for each history $h \in H$ that is reached under $\mathbb{P}_{\sigma_{-i}, \sigma_i}$ with positive probability. Lévy's zero-one law implies that $\mathbb{P}_{\sigma_{-i}, \sigma_i}(B) = 1$, and hence $\mathbb{P}_{\sigma_{-i}, \sigma_i}(Q_{i,\varepsilon}(f_i)) = 1$.

CLAIM 2: There is a strategy profile $\sigma_{-i} \in \Sigma_{-i}$ such that

$$\mathbb{E}_{\sigma_{-i}, \sigma_i}(f_i | h) \leq v_i(f_{i,h}) + \varepsilon/2 = v_i(f_i) + \varepsilon/2,$$

for each history $h \in H$ and each strategy $\sigma_i \in \Sigma_i$. We sketch the main idea behind the construction of σ_{-i} , following closely the arguments in Mashiah-Yaakovi [38, Proposition 11] and Flesch, Herings, Maes, and Predtetchinski [16, Theorem 5.7].

By the definition of the minmax value, for each history $h \in H$, there is a strategy profile $\sigma_{-i,h} \in \Sigma_{-i}$ such that for each strategy $\sigma_i \in \Sigma_i$

$$\mathbb{E}_{\sigma_{-i,h}, \sigma_i}(f_i | h) \leq v_i(f_{i,h}) + \varepsilon/4 = v_i(f_i) + \varepsilon/4. \quad (3)$$

Now the strategy profile σ_{-i} prescribes to play as follows: let $t_0 = 0$ and $h_0 = \emptyset$ be the empty history at stage 0. Take a $k \in \mathbb{N}$ for which a stage t_k and a history h_k (called the k th switching history) have been defined. As of stage t_k , the players $-i$ are required to start following the strategy profile σ_{-i,h_k} , and do so until a history h_{k+1} arises at some stage t_{k+1} such that σ_{-i,h_k} is no longer good enough in the corresponding subgame, in the sense that

$$\sup_{\sigma_i \in \Sigma_i} \mathbb{E}_{\sigma_{-i,h_k}, \sigma_i}(f_i | h_{k+1}) > v_i(f_i) + \varepsilon/2.$$

As of stage t_{k+1} , the players are required to switch to $\sigma_{-i,h_{k+1}}$, and follow it until stage t_{k+2} such that $\sigma_{-i,h_{k+1}}$ is no longer good enough, and so on.

The crucial point of the construction is that, given a strategy σ_i of player i , at each history h_{k+1} prescribing a switch, the corresponding expected payoff exceeds that at the previous

switching history h_k by $\epsilon/4$. Payoffs, however, are bounded. It follows that $\mathbb{P}_{\sigma_{-i}, \sigma_i}$ -almost surely, at most finitely many switches can occur. And hence, starting from some finite k , the strategy profile σ_{-i} coincides with σ_{-i, h_k} . And this implies that the expected payoff under (σ_{-i}, σ_i) is at most $v_i(f_i) + \epsilon/2$, as desired.

With σ_{-i} defined, the rest of the proof of the claim is similar to that of Claim 1. \square

The proof of Theorem 4.3: Fix an $\epsilon > 0$. By Proposition 4.4, $v_i(Q_{i,\epsilon}(f_i)) = 1$.

Let $\Gamma^\epsilon = (I, A, H, (1_{Q_{i,\epsilon}(f_i)})_{i \in I})$ be an auxiliary Blackwell game where player i 's winning set is $Q_{i,\epsilon}(f_i)$, the set of player i 's ϵ -individually rational plays in Γ . Each player's minmax value in the game Γ^ϵ equals 1. Therefore, the auxiliary game Γ^ϵ satisfies the hypothesis of Theorem 4.1. We conclude that the intersection $\bigcap_{i \in I} Q_{i,\epsilon}(f_i)$ is not empty, and hence there is a play $p^* \in [H]$ such that $f_i(p^*) \geq v_i(f_i) - \epsilon$, for every $i \in I$.

The following strategy profile is a 2ϵ -equilibrium of Γ (see also Aumann and Shapley [5]):

- The players follow the play p^* , until the first stage in which one of the players deviates from this play. Denote by i the minimal index of a player who deviates from p^* at that stage.
- From the next stage and on, the players in $-i$ switch to a strategy profile that reduces player i 's payoff to $v_i(f_i) + \epsilon$. A strategy profile with this property does exist by the assumption of history-independence of the minmax values.

This completes the proof of the theorem. \square

We illustrate the construction of the ϵ -equilibrium with the following example.

Example 4.5. We consider a two-player Blackwell game with history-independent action sets where the same stage game is being played at each stage, and a player's objective is to maximize the long-term frequency of the stages she wins. Specifically, $\Gamma = (\{1, 2\}, A_1, A_2, f_1, f_2)$, where A_1 and A_2 are finite, and

$$f_i(a_0, a_1, \dots) = \limsup_{t \rightarrow \infty} \frac{1}{t} \cdot \#\{k < t : a_k \in U_i\},$$

for each $(a_0, a_1, \dots) \in A^\mathbb{N}$. Here U_i is player i 's stage winning set. We assume that U_1 and U_2 are disjoint, and let d_i denote player i 's minmax value in the stage game.

Note that f_i is a tail function (see Section 6), and it is a limsup function (in the sense of the definition in Section 3). We have $d_i = v_i(f_i)$, i.e., player i 's minmax value in the stage game is also player i 's minmax value in the Blackwell game.

Take any Nash equilibrium $x \in \prod_{i \in I} \Delta(A_i)$ of the stage game. Playing x at each stage is certainly a 0-equilibrium of the Blackwell game Γ , but typically it is **not** of the type that appears in the proof of Theorem 4.3. An important feature of the ϵ -equilibrium constructed in the proof is that the equilibrium play is pure; only off the equilibrium path might a player be requested to play a mixed action. In this particular example we can even choose the equilibrium play to be periodic. This can be done as follows.

First note that $d_1 + d_2 \leq 1$. This follows since $\mathbb{P}_x(U_1) + \mathbb{P}_x(U_2) \leq 1$ (because U_1 and U_2 are disjoint by supposition) and since $d_i \leq \mathbb{P}_x(U_i)$ for $i = 1, 2$ (because the Nash equilibrium payoff is at least the minmax value). Let $\epsilon > 0$. Choose natural numbers m, m_1 , and m_2 such

that $d_i - \varepsilon \leq \frac{m_i}{m} \leq d_i$, for $i = 1, 2$. Note that $m_1 + m_2 \leq m$. Pick a point $a_1 \in U_1$ and a point $a_2 \in U_2$, and let p^* be the periodic play with period $m_1 + m_2$ obtained by repeating a_1 for the first m_1 stages, and repeating a_2 for the next m_2 stages. We have

$$d_i - \varepsilon \leq \frac{m_i}{m} \leq \frac{m_i}{m_1 + m_2} = f_i(p^*),$$

for $i \in \{1, 2\}$. One can support p^* as an ε -equilibrium play by a threat of punishment: in case of a deviation by player 1, player 2 will switch to playing the minmax action profile from the stage game for the rest of the game, thus reducing i 's payoff to d_i . A symmetric punishment is imposed on player 2 in case of a deviation.

Under the periodic play p^* , the sum of the players' payoffs is 1. There are alternative plays where the payoff to *both* players is 1, which can support 0-equilibria. For example, consider the non-periodic play p that is played in blocks of increasing size: for each $k \in \mathbb{N}$, the length of block k is 2^{2^k} . In even (resp. odd) blocks the players play the action profile a_1 (resp. a_2). The reader can verify that since the ratio between the length of block k and the total length of the first k blocks goes to ∞ , the payoff to both players at p is 1.

5 Regularity and the folk theorem

A payoff vector $w \in \mathbb{R}^{|I|}$, assigning a payoff to each player, is called *an equilibrium payoff* of the Blackwell game Γ if for every $\varepsilon > 0$ there exists an ε -equilibrium σ^ε of Γ such that $\|w - \mathbb{E}_{\sigma^\varepsilon}(f)\|_\infty \leq \varepsilon$. In other words, an equilibrium payoff is an accumulation point of ε -equilibrium payoff vectors as ε goes to 0. We let \mathcal{E} denote the set of equilibrium payoffs. Our goal here is to provide a description of \mathcal{E} .

In repeated games with stage payoffs, where the total payoff is some average (discounted average with low discounting, liminf of average, limsup of average, etc.) of the stage game payoffs, the folk theorem states that the set of equilibrium payoffs coincides with the set of all individually rational vectors that are in the convex hull of the feasible payoff vectors, see, e.g., Aumann and Shapley [5], Sorin [50], and Mailath and Samuelson [31]. As we will see, when the payoff functions are general, the set of equilibrium payoffs is the convex hull of the set of feasible payoff vectors that are individually rational. The reason for the difference is that in repeated games with stage payoffs, getting a low payoff in one stage can be compensated by getting a high payoff in the following stage; when the payoff is obtained only at the end of the game, there is no opportunity to compensate low payoffs.

Define

$$Q^\varepsilon(f) := \bigcap_{i \in I} Q_{i,\varepsilon}(f_i),$$

$$W^\varepsilon(f) := \{f(p) : p \in Q^\varepsilon(f)\}.$$

The set $Q^\varepsilon(f)$ is the set of ε -individually rational plays, and $W^\varepsilon(f)$ is the set of feasible and ε -individually rational payoffs vectors. Whenever convenient, we write simply Q^ε and W^ε . For every set X in a Euclidean space we denote its closure by $\text{cl}(X)$ and its convex hull by $\text{conv}(X)$.

Theorem 5.1. *Consider a Blackwell game $\Gamma = (I, A, H, (f_i)_{i \in I})$. Suppose that for each player $i \in I$, player i 's action set $A_i(h)$ at each history $h \in H$ is finite, her payoff function f_i is bounded and upper*

semi-analytic, and her minmax value is history-independent. Then

$$\mathcal{E} = \bigcap_{\varepsilon > 0} \text{conv}(\text{cl}(W^\varepsilon(f))).$$

To prove Theorem 5.1 we need the following result, which states that every ε -equilibrium assigns high probability to plays in $Q^{\varepsilon^{1/3}}(f)$.

Lemma 5.2. *Consider an n -player Blackwell game $\Gamma = (I, A, H, (f_i)_{i \in I})$. Suppose that for each player $i \in I$, player i 's action set $A_i(h)$ at each history $h \in H$ is finite, her payoff function f_i is bounded and upper semi-analytic, and her minmax value is history-independent. Let $\varepsilon > 0$ be sufficiently small, and let σ^ε be an ε -equilibrium. Then*

$$\mathbb{P}_{\sigma^\varepsilon}(A^\mathbb{N} \setminus Q^{\varepsilon^{1/3}}(f)) < n\varepsilon^{1/3}.$$

Proof. Set $\eta := \varepsilon^{1/3}$. It suffices to show that for every $i \in I$,

$$\mathbb{P}_{\sigma^\varepsilon}(A^\mathbb{N} \setminus Q_{i,\eta}(f_i)) < \eta. \quad (4)$$

Fix a player $i \in I$ and suppose to the contrary that Eq. (4) does not hold. We derive a contradiction by showing that player i has a deviation from σ^ε that yields her a gain higher than ε .

For $t \in \mathbb{N}$, denote by $X_t := \mathbb{P}_{\sigma^\varepsilon}(Q_{i,\eta}(f_i) | \mathcal{F}_t)$ the conditional probability of the event $Q_{i,\eta}(f_i)$ under the strategy profile σ^ε given the sigma-algebra \mathcal{F}_t . By Doob's martingale convergence theorem, $(X_t)_{t \in \mathbb{N}}$ converges to the indicator function of the event $Q_{i,\eta}(f_i)$, almost surely under $\mathbb{P}_{\sigma^\varepsilon}$. Since by supposition $\mathbb{P}_{\sigma^\varepsilon}(A^\mathbb{N} \setminus Q_{i,\eta}(f_i)) > \eta$, we know that $\mathbb{P}_{\sigma^\varepsilon}(X_t \rightarrow 0) > \eta$.

Let K be a bound on the game's payoffs, and let $\rho := \varepsilon^2/K$. Let us call a history $h \in H_t$ a *deviation history* if under h , stage t is the first one such that $X_t < \rho$. On the event $\{X_t \rightarrow 0\}$, a deviation history arises at some point during play. Consequently, under $\mathbb{P}_{\sigma^\varepsilon}$, a deviation history arises with probability of at least η .

Consider the following strategy σ'_i of player i : play according to σ_i^ε until a deviation history, say h , occurs (and forever if a deviation history never occurs). At h , switch to playing a strategy which guarantees player i a payoff of at least $v_i(f_i) - \varepsilon$ against σ_{-i}^ε in Γ_h . Such a strategy exists by our supposition of history-independence of the minmax values. To conclude the argument, we compute the gain from the deviation to σ'_i .

For every deviation history $h \in H_t$,

$$\mathbb{E}_{\sigma_{-i}^\varepsilon, \sigma'_i}(f_i | h) \geq v_i(f_i) - \varepsilon, \quad (5)$$

$$\mathbb{E}_{\sigma_{-i}^\varepsilon, \sigma_i^\varepsilon}(f_i | h) \leq \rho K + (1 - \rho)(v_i(f_i) - \eta). \quad (6)$$

Eq. (5) holds by the choice of σ'_i . To derive Eq. (6), suppose that, following the history h , player i conforms to σ_i^ε . Then, conditional on h , with probability at most ρ the play belongs to $Q_{i,\eta}(f_i)$, and player i 's payoff is at most K , and with probability at least $1 - \rho$ the play does not belong to $Q_{i,\eta}(f_i)$, and player i 's payoff is at most $v_i(f_i) - \eta$.

We can now compute the gain from the deviation to σ'_i : If a deviation history never arises, σ'_i recommends the same actions as σ_i^ε , and therefore the gain is 0. A deviation history occurs with a probability of at least η , and thus

$$\mathbb{E}_{\sigma_{-i}^\varepsilon, \sigma'_i}(f_i) - \mathbb{E}_{\sigma_{-i}^\varepsilon, \sigma_i^\varepsilon}(f_i) \geq \eta(v_i(f_i) - \varepsilon - \rho K - (1 - \rho)(v_i(f_i) - \eta))$$

$$\begin{aligned}
&= \eta(-\varepsilon - \varepsilon^2 + \rho v_i(f_i) + \eta - \rho\eta) \\
&= \varepsilon^{\frac{2}{3}}(1 - \varepsilon^{\frac{2}{3}} - \varepsilon^{\frac{5}{3}} + \frac{v_i(f_i)}{K} \varepsilon^{\frac{5}{3}} - \frac{1}{K} \varepsilon^2),
\end{aligned}$$

which behaves like $\varepsilon^{\frac{2}{3}}$ when ε is small, and therefore exceeds ε . \square

Proof of Theorem 5.1: Let $|I| = n$. Let $w \in \mathbb{R}^n$ be an equilibrium payoff. Assume by contradiction that there is an $\alpha > 0$ such that $w \notin \text{conv}(\text{cl}(W^\alpha))$. For a vector $z \in \mathbb{R}^n$ write $\text{dist}(z)$ to denote the distance from z to the set $\text{conv}(\text{cl}(W^\alpha))$ under the $\|\cdot\|_\infty$ metric on \mathbb{R}^n . By assumption, $\delta := \frac{1}{4}\text{dist}(w) > 0$. Denote $\varepsilon := \min(\delta, \alpha^3, (\frac{\delta}{Kn})^3) > 0$, where K is a bound on the game payoff.

From w being an equilibrium payoff, there exists an ε -equilibrium, say σ^ε , such that $\|w - \mathbb{E}_{\sigma^\varepsilon}(f)\|_\infty \leq \varepsilon \leq \delta$. We have the following chain of inequalities:

$$\text{dist}(\mathbb{E}_{\sigma^\varepsilon}(f)) \leq \mathbb{E}_{\sigma^\varepsilon}(\text{dist}(f)) \leq 2K \cdot \mathbb{P}_{\sigma^\varepsilon}(A^N \setminus Q^\alpha(f)) \leq 2K \cdot n \cdot \varepsilon^{\frac{1}{3}} \leq 2\delta,$$

where the first inequality follows from the fact that $\text{dist} : \mathbb{R}^n \rightarrow \mathbb{R}$ is a convex function, the second from the fact that $f(p) \in W^\alpha$ whenever $p \in Q^\alpha(f)$, the third follows since $Q^{\varepsilon^{1/3}}(f) \subseteq Q^\alpha(f)$ and by Lemma 5.2, and the last holds by the choice of ε . But then

$$\text{dist}(w) \leq \|w - \mathbb{E}_{\sigma^\varepsilon}(f)\|_\infty + \text{dist}(\mathbb{E}_{\sigma^\varepsilon}(f)) \leq 3\delta,$$

contradicting the choice of δ .

We turn to prove the other direction. Let $w \in \bigcap_{\varepsilon > 0} \text{conv}(\text{cl}(W^\varepsilon))$. We need to show that w is an equilibrium payoff. Fix an $\varepsilon > 0$.

Carathéodory's Theorem (Carathéodory, [11]) implies that $\text{cl}(\text{conv}(W^\varepsilon)) = \text{conv}(\text{cl}(W^\varepsilon))$, hence w is an element of $\text{cl}(\text{conv}(W^\varepsilon))$, and thus we can choose a vector $w_\varepsilon \in \text{conv}(W^\varepsilon)$ such that $\|w - w_\varepsilon\|_\infty \leq \varepsilon$. We argue that w_ε is a vector of expected payoffs in some 3ε -equilibrium.

The payoff w_ε can be presented as a convex combination of $n + 1$ vector payoffs, say $f(p^1), \dots, f(p^{n+1})$, with each p^k an element of $Q^\varepsilon(f)$. Using jointly controlled lotteries as done, e.g., in Forges [18], Lehrer [28], or Lehrer and Sorin [29], the players can generate the required randomization over the plays p^1, \dots, p^{n+1} during the first stages of the game. Once a specific play p^k has been chosen, the construction of the 3ε -equilibrium is standard: the players play p^k , and if player i deviates, her opponents revert to playing a strategy profile that gives player i at most $v_i(f_i) + \varepsilon$. Such a strategy exists by the assumption of history-independence of the minmax values. \square

Example 5.3. Consider the two-player Blackwell game $\Gamma = (\{1, 2\}, A_1, A_2, f_1, f_2)$, where the action sets are $A_1 = \{T, M, B\}$ and $A_2 = \{L, C, R\}$, and for a play $p = (a_0, a_1, \dots)$ the payoffs are

$$(f_1(p), f_2(p)) = \begin{cases} (1, 1) & \text{if } \liminf_{n \rightarrow \infty} \frac{1}{t} \cdot \#\{k < t : a_k = (T, L) \text{ or } a_k = (M, C)\} > \frac{1}{2}, \\ (4, -1) & \text{if } \liminf_{n \rightarrow \infty} \frac{1}{t} \cdot \#\{k < t : a_k = (B, L)\} = 1, \\ (-1, 4) & \text{if } \liminf_{n \rightarrow \infty} \frac{1}{t} \cdot \#\{k < t : a_k = (T, R)\} = 1, \\ (0, 0) & \text{otherwise.} \end{cases}$$

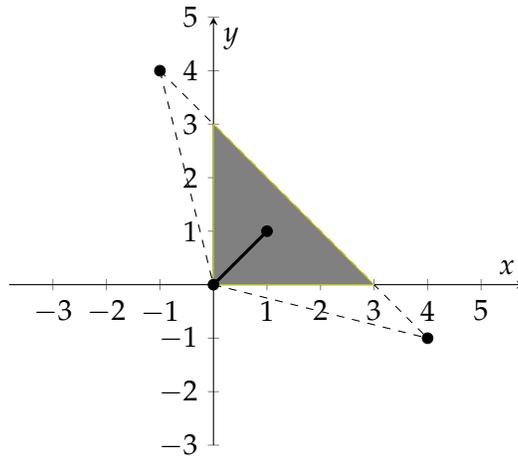


Figure 1: The set of equilibrium payoffs (the segment connecting $(0,0)$ and $(1,1)$) vs. the set of convex combinations of feasible payoffs that are individually rational (the dark triangle).

Thus the payoff is $(1,1)$ if the (liminf) frequency of the stages where either (T,L) or (M,C) is played is larger than $\frac{1}{2}$. It is $(4,-1)$ if (B,L) is played with frequency of 1, and $(-1,4)$ if (T,R) is played with the frequency of 1. All other cases result in a payoff of $(0,0)$.

Observe that when player 1 plays B repeatedly, the maximal payoff that player 2 can achieve is 0, and this is player 2's minmax value. Similarly, player 1's minmax value is 0. For each $\varepsilon \in (0,1)$, the set $W^\varepsilon(f)$ consists of the two points $(0,0)$ and $(1,1)$. By Theorem 5.1, the set of equilibrium payoffs \mathcal{E} is the line segment connecting $(0,0)$ and $(1,1)$, see Figure 1.

Naturally, all equilibrium payoffs w are (a) convex combinations of the feasible payoff vectors $(0,0)$, $(1,1)$, $(4,-1)$, and $(-1,4)$, and (b) individually rational, i.e., they satisfy $w_1 \geq 0$ and $w_2 \geq 0$. The set of all payoff vectors satisfying (a) and (b) is represented in Figure 1 by the shaded triangle. The point we wish to make here is that the properties (a) and (b) are not sufficient for a payoff vector to be an equilibrium payoff.

Take for concreteness the point $(3,0)$. This payoff vector is in the convex hull of the feasible payoff vectors and is individually rational. Yet, for $\varepsilon < \frac{2}{3}$, there is no ε -equilibrium with the payoff (close to) the vector $(3,0)$. We give a heuristic argument.

Suppose to the contrary that σ is such an ε -equilibrium. The strategy profile σ necessarily assigns a probability of at least $\frac{2}{3}$ to the set of plays that yield the payoff vector $(4,-1)$. But this implies that Player 2 has a deviation that would improve her payoff over the candidate ε -equilibrium by at least $\frac{2}{3}$. Player 2 needs to deviate to playing R forever (for example), at any history of the game where her conditional expected payoff under σ is close enough to -1 . Since playing R would yield at least 0, by such a deviation, she would improve her conditional expected payoff by at least 1. Levy's zero-one law guarantees that the histories where player 2 is called to deviate in this way arise with a probability close to $\frac{2}{3}$, so that the expected gain from the deviation is also close to $\frac{2}{3}$.

The above discussion of Example 5.3 leads to a slightly more general conclusion: if the set of feasible payoffs is finite, then the set of equilibrium payoffs is the convex-hull of the feasible payoffs that are individually rational (equal or larger than the minmax). For each player the

minmax value is within the finite set of feasible payoffs, and placing any probability on a payoff that is not individually rational enables profitable deviations.

6 Blackwell games with tail-measurable payoffs

An important class of games with history-independent minmax values are those where the payoff functions are tail-measurable. In this section we concentrate on games with tail-measurable payoffs.

Consider a Blackwell game with history-independent action sets, $\Gamma = (I, (A_i, f_i)_{i \in I})$. A set $Q \subseteq A^{\mathbb{N}}$ is said to be a *tail set* if whenever a play $p = (a_0, a_1, \dots)$ is an element of Q and $q = (b_0, b_1, \dots)$ is such that $a_t = b_t$ for all $t \in \mathbb{N}$ sufficiently large, then q is also an element of Q . Let \mathcal{T} denote the sigma-algebra of the tail subsets of $A^{\mathbb{N}}$. We note that the tail sigma-algebra \mathcal{T} and the Borel sigma-algebra \mathcal{B} are not nested. For constructions of tail sets that are not Borel, see Rosenthal [42] and Blackwell and Diaconis [8].

Examples of tail sets are: (1) the winning sets of Example 3.2, (2) the set of plays in which a certain action profile $a \in A$ is played with limsup-frequency at most $\frac{1}{2}$, and (3) the set of plays in which a certain action profile $a^* \in A$ is played at most finitely many times at even stages (with no restriction at odd stages).

An important class of tail sets are the shift invariant sets. A set $Q \subseteq A^{\mathbb{N}}$ is a *shift invariant set* if for each play $p = (a_0, a_1, \dots)$, $p \in Q$ if and only if $(a_1, a_2, \dots) \in Q$. Equivalently, shift invariant sets are the sets that are invariant under the backward shift operator on $A^{\mathbb{N}}$. Shift invariant sets are tail sets. The converse is not true: while the sets in examples (1) and (2) above are shift invariant, that of example (3) is not.

A function $f : A^{\mathbb{N}} \rightarrow \mathbb{R}$ is called tail-measurable if, for each $r \in \mathbb{R}$, the set $\{p \in A^{\mathbb{N}} : r \leq f(p)\}$ is an element of \mathcal{T} . Intuitively, a payoff function is tail measurable if an action taken in any particular stage of the game has no impact on the payoff. The payoff function in Example 4.5 is tail-measurable.

Remark 6.1. The assumption that the set of actions of each player is history-independent is required so that the tail-measurability of the payoff functions has a bite. If the sets of actions were history-dependent, then by having a different set of actions at each history, any function could be turned into tail-measurable.

We now state one key implication of tail-measurability, namely the history-independence of minmax values.

Proposition 6.2. *Let $\Gamma = (I, (A_i, f_i)_{i \in I})$ be a Blackwell game with history-independent action sets, and let $i \in I$ be a player. If player i 's payoff function is bounded, upper semi-analytic, and tail-measurable, then her minmax value is history-independent.*

Proof. It suffices to show that $v_i(f_{i,a}) = v_i(f_i)$ for each $a \in A$, where, with a slight abuse of notation, we write a for a history in stage 1. Since f_i is tail-measurable, all the functions $f_{i,a}$ for $a \in A$ are identical to each other. Hence, fixing any particular action profile $\bar{a} \in A$, letting $X_i := \Delta(A_i)$ and $X_{-i} := \prod_{j \in -i} X_j$, we have

$$v_i(f_i) = \inf_{x_{-i} \in X_{-i}} \sup_{x_i \in X_i} \sum_{a \in A} \prod_{j \in I} x_j(a_j) \cdot \left(\inf_{\sigma_{-i} \in \Sigma_{-i}} \sup_{\sigma_i \in \Sigma_i} \mathbb{E}_{\sigma_{-i}, \sigma_i}(f_{i,a}) \right)$$

$$= \inf_{x_{-i} \in X_{-i}} \sup_{x_i \in X_i} \sum_{a \in A} \prod_{j \in I} x_j(a_j) \cdot v_i(f_{i,\bar{a}}) = v_i(f_{i,\bar{a}}).$$

□

If the payoff functions of all the players in a game Γ are tail-measurable, then, for each fixed stage $t \in \mathbb{N}$, all the subgames of Γ starting at stage t are identical. On the other hand, the subgames starting, say, at stage 1, are not identical to the game itself (see example (3) of a tail-measurable payoff function above). Nonetheless, as Proposition 6.2 implies, the players' minmax values *are* the same in every subgame.

The condition of history-independence of the minmax values is more inclusive than that of tail-measurability of the payoffs; the examples that follow illustrate the point.

Example 6.3. Consider a one-player Blackwell game where the player's payoff function is 1_S , the indicator of a set $S \subseteq [H]$. If S is dense in $[H]$, then the minmax value of the player is 1 in each subgame. A dense set may or may not be a tail set.

Example 6.4. We consider a Blackwell game similar to that of Example 3.2, but where the stage game may depend on the history, as long as each player's stage minmax value is the same.

Specifically, let $\Gamma = (I, A, H, (1_{W_i})_{i \in I})$. Suppose that at each history $h \in H$, each player $i \in I$ has a stage winning set $U_i(h) \subseteq A(h)$, and her winning set in the Blackwell game Γ is

$$W_i = \{(a_0, a_1, \dots) \in [H] : a_t \in U_i(a_0, \dots, a_{t-1}) \text{ for infinitely many } t \in \mathbb{N}\}.$$

Assume that the stage minmax value of player i is the same at each history: there is a number d_i such that

$$d_i = \inf_{x_{-i} \in \Delta(A_{-i}(h))} \sup_{x_i \in \Delta(A_i(h))} \mathbb{P}_{x_{-i}, x_i}(U_i(h))$$

for every $h \in H$. Then player i 's minmax value in each subgame of Γ is 0 if $d_i = 0$, and is 1 if $d_i > 0$. Thus player i 's minmax value is history-independent.

Note that the game Γ need not have history-independent action sets. Even when the action sets *are* history-independent, the winning sets need not necessarily be tail-measurable.

To illustrate the last claim, suppose that there are two players playing matching pennies at each stage. At stage 0, player 1 wants to match the choice of player 2 (and player 2 wants to mismatch the choice of player 1). Subsequently the roles of the two players swap as follows: the player to win stage t wants to match her opponent's action at stage $t + 1$, while the loser at stage t wants to mismatch the action of her opponent at stage $t + 1$. Formally, we let $\Gamma = (\{1, 2\}, A_1, A_2, 1_{W_1}, 1_{W_2})$ be the two-player Blackwell game with history-independent action sets, where $A_1 = A_2 = \{H, T\}$, the winning sets W_1 and W_2 are as above, and the stage winning sets are defined recursively as follows:

$$U_1(\emptyset) = \{(H, H), (T, T)\} \quad \text{and} \quad U_2(\emptyset) = \{(H, T), (T, H)\},$$

and

$$U_1(h, a) = \begin{cases} U_1(\emptyset) & \text{if } a \in U_1(h), \\ U_2(\emptyset) & \text{if } a \in U_2(h), \end{cases} \quad \text{and} \quad U_2(h, a) = \begin{cases} U_2(\emptyset) & \text{if } a \in U_1(h), \\ U_1(\emptyset) & \text{if } a \in U_2(h), \end{cases}$$

for each $h \in H$ and $a \in A$. The sets W_1 and W_2 are not tail: out of the two plays

$$((H, H), (H, H), (H, H), \dots) \text{ and}$$

$$((H, T), (H, H), (H, H), \dots),$$

the first is an element of $W_1 \setminus W_2$, while the second is an element of $W_2 \setminus W_1$.

Example 6.5. Consider a Blackwell game $\Gamma = (I, A, H, (f_i)_{i \in I})$, where player i 's objective is (as in Example 4.5) to maximize the long-term frequency of the stages she wins:

$$f_i(a_0, a_1, \dots) = \limsup_{t \rightarrow \infty} \frac{1}{t} \cdot \#\{k < t : a_k \in U_i(a_0, \dots, a_{k-1})\}.$$

As in the previous example, $U_i(h) \subseteq A(h)$ is player i 's stage winning set at history $h \in H$. Assume, as above, that player i 's minmax value in each stage game is d_i . Then also her minmax value in each subgame of Γ is d_i .

Example 6.6. Start with a Blackwell game with tail-measurable payoff functions. Suppose that the minmax values of all the players in the game are 0. Take any history h , and redefine the payoff functions so that any play having h as a prefix has a payoff of 0. In the resulting game, the minmax value of each player in each subgame remains 0, but the payoff functions are no longer tail-measurable (unless the original payoff functions are constant). A similar modification can be performed with any subset of histories, not just one.

From the results above we now deduce a number of implications for Blackwell games with tail-measurable payoffs.

Corollary 6.7. *Consider a Blackwell game $\Gamma = (I, (A_i, 1_{W_i})_{i \in I})$ with history-independent action sets. If player i 's winning set W_i is an analytic tail set, then $v_i(W_i)$ is either 0 or 1.*

Proof. Suppose that $v_i(W_i) > 0$. Let $\varepsilon := v_i(W_i)/2$. In view of Proposition 6.2, player i 's minmax value in Γ is history-independent. Applying Proposition 4.4, we conclude that $v_i(Q_{i,\varepsilon}(1_{W_i})) = 1$. But $Q_{i,\varepsilon}(1_{W_i}) = W_i$ by the choice of ε . \square

The following conclusion follows directly from Proposition 6.2 and Theorem 4.3.

Corollary 6.8. *Suppose that the game $\Gamma = (I, (A_i, 1_{W_i})_{i \in I})$ has history-independent action sets. Suppose, furthermore, that for each player $i \in I$, player i 's action set A_i is finite and her payoff function f_i is bounded, upper semi-analytic, and tail-measurable. Then for every $\varepsilon > 0$ the game admits an ε -equilibrium.*

7 Concluding remarks

Approximations by compact sets. Any Borel probability measure on $[H]$ (recall that $[H]$ is Polish under the maintained assumptions), is not merely regular, but is tight: the probability of a Borel set $B \subseteq [H]$ can be approximated from below by the probability of a compact subset $K \subseteq B$ (Kechris [25, Theorem 17.11]). The minmax value is not tight in this sense. To see this, consider any two-player Blackwell game where player 1's winning set W_1 is the entire set of plays $[H]$, so that $v_1(W_1) = 1$, and where $A_2(\emptyset)$, player 2's action set at the beginning of the game, is \mathbb{N} . We argue that $v_1(K) = 0$ for every compact set $K \subseteq W_1$. Indeed, the projection of a compact set $K \subseteq W_1$ on $A_2(\emptyset)$ is a compact, and hence a finite set. Therefore, player 2 can

guarantee that the realized play is outside K by choosing a sufficiently large action at stage 0. Thus $v_1(K) = 0$, as claimed.

Approximations by semicontinuous functions. The conclusion of Theorem 3.4 would no longer be true without the assumption of history-independence of the minmax values. Here we give an example of a game with a limsup payoff function where the minmax value cannot be approximated from below by an upper semicontinuous function.

Consider a zero-sum game Γ where $A_1 = A_2 = \{0, 1\}$, and player 1's payoff function is

$$f(a_0, a_1, \dots) = \begin{cases} \limsup_{t \rightarrow \infty} \frac{1}{t} \#\{k < t : a_{2,k} = 0\}, & \text{if } \tau = \infty, \\ 2, & \text{if } \tau < \infty \text{ and } a_{2,\tau} = 1, \\ 0, & \text{if } \tau < \infty \text{ and } a_{2,\tau} = 0, \end{cases}$$

where $\tau = \tau(a_0, a_1, \dots) \in \mathbb{N} \cup \{\infty\}$ is the first stage where player 1 chooses action 1. The game was analyzed in Sorin [49], who showed that $v_1(f) = 2/3$.

Let $g \leq f$ be a bounded upper semicontinuous function. We argue that $v_1(g) \leq 1/2$. For $t \in \mathbb{N}$, let S_t denote the set of plays p such that $t \leq \tau(p)$. Note that S_t is closed. We argue that

$$\inf_{t \in \mathbb{N}} \sup\{g(p) : p \in S_t\} \leq 1.$$

Suppose this is not the case. Take an $\varepsilon > 0$ such that $1 + \varepsilon < \sup\{g(p) : p \in S_t\}$ for each $t \in \mathbb{N}$. Let $U_0 := \{1 + \varepsilon \leq g\}$, and for each $t \geq 1$ let $U_t := U_0 \cap S_t$. The set U_t is not empty for each $t \in \mathbb{N}$. Moreover, it is a closed, and hence a compact subset of $A^\mathbb{N}$. Thus $U_0 \supseteq U_1 \supseteq \dots$ is a nested sequence of non-empty compact sets. Therefore, there is a play $p \in \bigcap_{t \in \mathbb{N}} U_t$. It holds that $\tau(p) = \infty$, and consequently $f(p) \leq 1 < g(p)$, a contradiction.

Take an $\varepsilon > 0$. Find a $t \in \mathbb{N}$ such that $\sup\{g(p) : p \in S_t\} \leq 1 + \varepsilon$. Suppose that player 2 plays 0 for the first t stages, and thereafter plays 0 with probability $1/2$ at each stage. This guarantees that the payoff under the function g is at most $(1 + \varepsilon)/2$.

On the assumption of finiteness of the action sets. The hypothesis of Theorem 4.1 requires that the action sets at each history be finite, and its conclusion is not true without this assumption. Indeed, consider the two-player Blackwell game $\Gamma = (\{1, 2\}, A_1, A_2, W_1, W_2)$ with history-independent action sets $A_1 = A_2 = \mathbb{N}$. Player 1's winning set W_1 consists of all plays $(a_{1,t}, a_{2,t})_{t \in \mathbb{N}}$ such that $a_{1,t} > a_{2,t}$ holds for all sufficiently large $t \in \mathbb{N}$, and player 2's winning set W_2 consists of all plays $(a_{1,t}, a_{2,t})_{t \in \mathbb{N}}$ such that $a_{1,t} < a_{2,t}$ holds for all sufficiently large $t \in \mathbb{N}$. Then W_1 and W_2 are Borel-measurable and tail-measurable, and $v_1(W_1) = v_2(W_2) = 1$, but $W_1 \cap W_2 = \emptyset$. Hence, the game has no ε -equilibrium for any $\varepsilon < 1/2$. Indeed, an ε -equilibrium σ would need to satisfy $\mathbb{P}_\sigma(W_i) \geq v_i(W_i) - \varepsilon > 1/2$ for both $i \in \{1, 2\}$.

As discussed above, the assumption that the sets of actions are history-dependent is intertwined with the assumption that the payoffs are tail-measurable.

Continuity of the minmax. Unlike Borel probability measures, the minmax value is in general not continuous in the following sense: there is an increasing sequence of Borel sets $C_0 \subseteq C_1 \subseteq \dots$ such that $\lim_{n \rightarrow \infty} v_i(C_n) < v_i(\bigcup_{n \in \mathbb{N}} C_n)$. In fact, one can construct an example of this kind where C_n is both a G_δ and an F_σ set, as follows. Consider a two-player Blackwell game with history-independent action sets where A_1 is a singleton (player 1 is a dummy) while A_2 contains at least two distinct elements. Let $\{p_0, p_1, \dots\}$ be a converging (with respect to any

compatible metric on $A^{\mathbb{N}}$ sequence of plays, no two members of which are the same. Let $C_n := A^{\mathbb{N}} \setminus \{p_n, p_{n+1}, \dots\}$. Then $v_1(C_n) = 0$ for each $n \in \mathbb{N}$ while $v_1(\bigcup_{n \in \mathbb{N}} C_n) = v_1(A^{\mathbb{N}}) = 1$.

Maxmin value. Consider a Blackwell game Γ , and suppose that player i 's payoff function f_i is bounded and upper semi-analytic. Player i 's *maxmin value* is defined as

$$z_i(f_i) = \sup_{\sigma_i \in \Sigma_i} \inf_{\sigma_{-i} \in \Sigma_{-i}} \mathbb{E}_{\sigma_{-i}, \sigma_i}(f_i).$$

The minmax value is not smaller than the maxmin value: $z_i(f_i) \leq v_i(f_i)$. If $I = \{1, 2\}$, player 1's payoff function f_1 is bounded and Borel-measurable, and for every $h \in H$ either the set $A_1(h)$ of player 1's actions or the set $A_2(h)$ of player 2's actions at h is finite, then in fact $z_1(f_1) = v_1(f_1)$, as follows from the determinacy of zero-sum Blackwell games (Martin [36]). Strict inequality might arise for at least two reasons.

The first is the failure of determinacy. The results of Section 3 are established under the assumption that the action sets are countable, an assumption that is insufficient to guarantee determinacy of a two-player zero-sum Blackwell game even if player 1's winning set is clopen. Wald's game provides an illustration. Suppose that each of the two players chooses a natural number; player 1 wins provided that his choice is at least as large as player 2's. Formally, consider a Blackwell game with $I = \{1, 2\}$, where the action sets at \emptyset are $A_1(\emptyset) = A_2(\emptyset) = \mathbb{N}$, and player 1's winning set W_1 consists of plays such that player 1's stage 0 action is at least as large as player 2's stage 0 action: $a_{1,0} \geq a_{2,0}$. Then player 1's minmax value is $v_1(W_1) = 1$ while her maxmin value is $z_1(W_1) = 0$.

The second possibility for a maxmin and the minmax values to be different arises in games with three or more players. The reason is that the definitions of both the maxmin and the minmax values impose that the opponents of player i choose their actions independently after each history. The point is illustrated by Maschler, Solan, and Zamir [37, Example 5.41], which can be seen as a three-player Blackwell game with binary action sets, where the player's payoff function only depends on the stage 0 action profile.

Analogues of Theorems 3.1, 3.3, and 3.4 could be established for the maxmin values using the same approach.

Open problems. Existence of an ε -equilibrium in dynamic games with general (Borel-measurable) payoffs has been, and still is, one of the Holy Grails of game theory. A more modest approach, also pursued in this paper, is to establish existence in some special classes of games. Blackwell games, as they are defined here, do not include moves of nature. An interesting avenue for a follow up research is to extend the methods developed in this paper to the context of stochastic games with general Borel-measurable payoff functions.

Theorems 3.1 and 3.3 provide two distinct approximation results, and neither seems to be a consequence of the other. This raises the question of whether there is a natural single generalization that would encompass both these results as two special cases.

8 Appendix: The proof of Theorems 3.1, 3.3, and 3.4

The proofs of Theorems 3.1 and 3.3 are adaptations of the corresponding arguments in Maitra and Sudderth [34] and in Martin [36] and are provided here for completeness. Theorem 3.4 follows easily from Theorem 3.1 and Proposition 4.4.

Consider a Blackwell game $\Gamma = (I, A, H, (f_i)_{i \in I})$, fix a player $i \in I$, and suppose that player i 's payoff function f_i is bounded and Borel-measurable. Also assume w.l.o.g. that $0 \leq f_i \leq 1$. When we will consider Theorem 3.1 we will substitute $f_i = 1_{W_i}$.

Given $h \in H$, let $R(h)$ denote the set of one-shot payoff functions $r : A(h) \rightarrow [0, 1]$. Let $X_i(h) := \Delta(A_i(h))$ denote player i 's set of mixed actions at history h , and let $X_{-i}(h) := \prod_{j \in -i} X_j(h)$. For $x \in \prod_{i \in I} X_i(h)$ we write $r(x)$ to denote $\mathbb{E}_x(r)$, the expectation of r with respect to x . Player i 's minmax value of the function $r \in R(h)$ is

$$d_i(r) := \inf_{x_{-i} \in X_{-i}(h)} \sup_{x_i \in X_i(h)} r(x_{-i}, x_i).$$

We next introduce the main tool of the proof, an auxiliary two-player game of perfect information denoted by $G_i(f_i, c)$. This is a variation of the games G_v and G'_v in Martin [36, pp. 1575].

Given $c \in (0, 1]$ and a Borel measurable function $f_i : [H] \rightarrow [0, 1]$, define the game $G_i(f_i, c)$ as follows:

- Let $h_0 := \emptyset$. Player I chooses a one-shot payoff function $r_0 : A(h_0) \rightarrow [0, 1]$ such that $d_i(r_0) \geq c$.
- Player II chooses an action profile $a_0 \in A(h_0)$ such that $r_0(a_0) > 0$.
- Let $h_1 := (a_0)$. Player I chooses a one-shot payoff function $r_1 : A(h_1) \rightarrow [0, 1]$ such that $d_i(r_1) \geq r_0(a_0)$.
- Player II chooses an action profile $a_1 \in A(h_1)$ such that $r_1(a_1) > 0$.
- Let $h_2 := (a_0, a_1)$. Player I chooses a one-shot payoff function $r_2 : A(h_2) \rightarrow [0, 1]$ such that $d_i(r_2) \geq r_1(a_1)$. And so on.

This results in a run¹ $(r_0, a_0, r_1, a_1, \dots)$. Player I wins the run if

$$\limsup_{t \rightarrow \infty} r_t(a_t) \leq f_i(a_0, a_1, \dots) \quad \text{and} \quad 0 < f_i(a_0, a_1, \dots).$$

Let T be the set of all legal positions in the game $G_i(f_i, c)$. This is a tree on the set $R \cup A$ where $R := \cup_{h \in H} R(h)$. Sequences of even (odd) length in T are Player I's (Player II's) positions. The tree T is pruned: an active player has a legal move at each legal position of the game. Indeed, consider Player I's legal position in the game $G_i(f_i, c)$ and let h_t denote, as above, the sequence of action profiles produced, to date, by Player II. Then the function r_t which is identically equal to 1 on the set $A(h_t)$ is a legal move for Player I. Consider now Player II's legal position in $G_i(f_i, c)$, let h_t denote the sequence of action profiles produced to date by Player II, and let r_t be Player I's latest move. Then $d_i(r_t) > 0$. Therefore, there exists an action profile $a_t \in A(h_t)$ such that $r_t(a_t) > 0$, and thus a_t is Player II's legal move at the given position.

The set $[T]$ is the set of all runs of the game $G_i(f_i, c)$, a subset of $(R \cup A)^{\mathbb{N}}$.

¹To distinguish histories and plays of Γ from those of $G_i(f_i, c)$, we refer to the latter as *positions* and *runs*. To distinguish the players of Γ from those of $G_i(f_i, c)$, we refer to the latter as Player I and Player II, using the initial capital letters.

A run is *consistent* with a pure strategy σ_I of Player I if it is generated by the pair (σ_I, σ_{II}) , for some pure strategy σ_{II} of Player II. Runs that are consistent with pure strategies of Player II are defined analogously.

Player I's pure strategy σ_I in $G_i(f_i, c)$ is said to be *winning* if Player I wins all runs of the game that are consistent with σ_I .

Proposition 8.1. *Let $c \in (0, 1]$ and let $f_i : [H] \rightarrow [0, 1]$ be a Borel-measurable function. If Player I has a winning strategy in the game $G_i(f_i, c)$, then there exists a closed set $C \subseteq [H]$ and a limsup function $g : [H] \rightarrow [0, 1]$ such that $g \leq f_i$, $\{g > 0\} \subseteq C \subseteq \{f_i > 0\}$, and $c \leq v_i(g)$. In particular, $c \leq v_i(C)$; and if $f_i = 1_{W_i}$, then $C \subseteq W_i$.*

Proof. Fix Player I's winning strategy σ_I in $G_i(f_i, c)$.

STEP 1: Defining $C \subseteq [H]$ and $g : [H] \rightarrow [0, 1]$.

Let $T_I \subseteq T$ denote the set of positions in the game $G_i(f_i, c)$ of even length (i.e., Player I's positions) that are consistent with σ_I , i.e., those positions that can be reached under a strategy profile (σ_I, σ_{II}) for some pure strategy of σ_{II} of Player II. Let $\pi_I : T_I \rightarrow H$ be the projection that maps a position of length $2t$ in $G_i(f_i, c)$ to a history of length t in Γ : Formally, $\pi_I(\emptyset) := \emptyset$, $\pi_I(r_0, a_0) := (a_0)$, etc. Let $H_I \subseteq H$ be the image of T_I under π_I . Since in the tree T_I Player I's moves are uniquely determined by σ_I , the map π_I is in fact a bijection between T_I and H_I . We write $\phi : H_I \rightarrow T_I$ for the inverse of π_I . The map ϕ induces a continuous bijection $[H_I] \rightarrow [T_I]$, which we also denote by ϕ . We say that positions in H_I are σ_I -*acceptable*, and define C to be the set $[H_I]$.

For each $t \in \mathbb{N}$, define the function $\rho_t : H_t \rightarrow \mathbb{R}$ as follows: $\rho_0(\emptyset) := c$. Let $t \in \mathbb{N}$ and consider a history $h_t \in H_t$. If h_t is not σ_I -acceptable, we define $\rho_{t+1}(h_t, a_t) := 0$ for each $a_t \in A(h_t)$. Suppose that h_t is σ_I -acceptable, and let $r_t := \sigma_I(\phi(h_t))$. For each $a_t \in A(h_t)$ define $\rho_{t+1}(h_t, a_t) := r_t(a_t)$. Note that if h_t is σ_I -acceptable while (h_t, a_t) is not, we have $\rho_{t+1}(h_t, a_t) = r_t(a_t) = 0$.

Also define $g : [H] \rightarrow [0, 1]$ by letting

$$g(a_0, a_1, \dots) := \limsup_{t \rightarrow \infty} \rho_t(a_0, \dots, a_{t-1}).$$

STEP 2: Verifying that $g \leq f_i$ and $\{g > 0\} \subseteq C \subseteq \{f_i > 0\}$.

Since σ_I is Player I's winning strategy in $G_i(f_i, c)$, all runs in $[T_I]$ are won by Player I, and hence $[H_I] \subseteq \{f_i > 0\}$. For a play $p = (a_0, a_1, \dots)$ in $[H_I]$, if $\phi(p) = (r_0, a_0, r_1, a_1, \dots)$, then $g(p)$ equals $\limsup_{t \rightarrow \infty} r_t(a_t)$. Since the run $\phi(p)$ is won by Player I, we conclude that $g(p) \leq f_i(p)$. Thus $g \leq f_i$ on $[H]$.

STEP 3: Verifying that $c \leq v_i(g)$. Since $g \leq 1_C$ it will then follow that $c \leq v_i(C)$.

Fix a strategy profile $\sigma_{-i} \in \Sigma_{-i}$ for the players in $-i$ in the game Γ . Take any $\epsilon > 0$. We define a strategy σ_i for player i in the game Γ with the property that $\mathbb{E}_{\sigma_{-i}, \sigma_i}(g) \geq c - 2\epsilon$.

STEP 3.1: Defining player i 's strategy σ_i .

Let $r_0 := \sigma_I(\emptyset)$, Player I's first move in $G_i(f_i, c)$ according to her strategy σ_I . Define $\sigma_i(\emptyset)$ to be a mixed action on $A_i(\emptyset)$ such that

$$r_0(\sigma_{-i}(\emptyset), \sigma_i(\emptyset)) \geq c - \epsilon.$$

Let $t \geq 1$ and consider a history $h_t = (a_0, \dots, a_{t-1}) \in H_t$ of Γ . If h_t is not σ_I -acceptable, then $\sigma_i(h_t)$ is arbitrary. If h_t is σ_I -acceptable, let $\phi(h_t) := (r_0, a_0, \dots, r_{t-1}, a_{t-1})$ and $r_t := \sigma_I(\phi(h_t))$. Define $\sigma_i(h_t)$ to be a mixed action on $A_i(h_t)$ such that

$$r_t(\sigma_{-i}(h_t), \sigma_i(h_t)) \geq r_{t-1}(a_{t-1}) - \epsilon \cdot 2^{-t}.$$

STEP 3.2: Verifying that $\mathbb{E}_{\sigma_{-i}, \sigma_i}(g) \geq c - 2\epsilon$.

For each $t \in \mathbb{N}$ let us define $\rho_t^\epsilon := \rho_t - \epsilon \cdot 2^{-t+1}$. One can think of the functions $\rho_0^\epsilon, \rho_1^\epsilon, \dots$ as a stochastic process on $[H]$ that is measurable with respect to the filtration $\{\mathcal{F}_t\}_{t \in \mathbb{N}}$. We now argue that this process is a submartingale with respect to the measure $\mathbb{P}_{\sigma_{-i}, \sigma_i}$.

Letting $r_0 := \sigma_I(\emptyset)$ we have

$$\mathbb{E}_{\sigma_{-i}, \sigma_i}(\rho_1^\epsilon) = \mathbb{E}_{\sigma_{-i}, \sigma_i}(r_0(a_0)) - \epsilon = r_0(\sigma_{-i}(\emptyset), \sigma_i(\emptyset)) - \epsilon \geq c - 2\epsilon = \rho_0^\epsilon(\emptyset).$$

Consider a σ_I -acceptable history $h_t = (a_0, \dots, a_{t-1}) \in H_t$ of length $t \geq 1$. Let $(r_0, a_0, \dots, r_{t-1}, a_{t-1}) := \phi(h_t)$ and $r_t := \sigma_I(\phi(h_t))$. We have

$$\begin{aligned} \mathbb{E}_{\sigma_{-i}, \sigma_i}(\rho_{t+1}^\epsilon | h_t) &= \mathbb{E}_{\sigma_{-i}, \sigma_i}(r_t(a_t) | h_t) - \epsilon \cdot 2^{-t} \\ &= r_t(\sigma_{-i}(h_t), \sigma_i(h_t)) - \epsilon \cdot 2^{-t} \\ &\geq r_{t-1}(a_{t-1}) - \epsilon \cdot 2^{-t} - \epsilon \cdot 2^{-t} \\ &= \rho_t^\epsilon(h_t). \end{aligned}$$

On the other hand, if h_t is not σ_I -acceptable, then

$$\mathbb{E}_{\sigma_{-i}, \sigma_i}(\rho_{t+1}^\epsilon | h_t) = -\epsilon \cdot 2^{-t} > -\epsilon \cdot 2^{-t+1} = \rho_t^\epsilon(h_t).$$

This establishes the submartingale property for $\rho_0^\epsilon, \rho_1^\epsilon, \dots$.

The submartingale property implies that $\mathbb{E}_{\sigma_{-i}, \sigma_i}(\rho_t^\epsilon) \geq \rho_0^\epsilon(\emptyset) = c - 2\epsilon$ for each $t \in \mathbb{N}$. Using Fatou lemma we thus obtain

$$\mathbb{E}_{\sigma_{-i}, \sigma_i}(g) = \mathbb{E}_{\sigma_{-i}, \sigma_i}(\limsup_{t \rightarrow \infty} \rho_t) \geq \mathbb{E}_{\sigma_{-i}, \sigma_i}(\limsup_{t \rightarrow \infty} \rho_t^\epsilon) \geq \limsup_{t \rightarrow \infty} \mathbb{E}_{\sigma_{-i}, \sigma_i}(\rho_t^\epsilon) \geq c - 2\epsilon,$$

as desired. □

Proposition 8.2. *Let $c \in (0, 1]$ and let $f_i : [H] \rightarrow [0, 1]$ be a Borel-measurable function. If Player II has a winning strategy in the game $G_i(f_i, c)$, then for every $\epsilon > 0$ there exists an open set $O \subseteq [H]$ and a limsup function $g : [H] \rightarrow [0, 1]$ such that $f_i \leq g$, $\{f_i = 1\} \subseteq O \subseteq \{g = 1\}$, and $v_i(g) \leq c + \epsilon$. In particular, $v_i(O) \leq c + \epsilon$; and if $f_i = 1_{W_i}$, then $W_i \subseteq O$.*

Proof. Fix Player II's winning strategy σ_{II} in $G_i(f_i, c)$.

STEP 1: Defining $O \subseteq [H]$ and $g : [H] \rightarrow [0, 1]$.

We recursively define (a) the notion of a σ_{II} -acceptable history in the game Γ , (b) for each σ_{II} -acceptable history h in Γ , Player I's position $\psi(h)$ in the game $G_i(f_i, c)$, and (c) for each σ_{II} -acceptable history h of Γ , a function $u_h : A(h) \rightarrow [0, 1]$.

The empty history \emptyset of Γ is σ_{II} -acceptable. We define $\psi(\emptyset) := \emptyset$, the empty history in $G_i(f_i, c)$. Let $t \in \mathbb{N}$ and consider a history $h_t \in H_t$ of the game Γ . If h_t is not σ_{II} -acceptable, so is the history (h_t, a_t) for each $a_t \in A(h_t)$. Suppose that h_t is σ_{II} -acceptable and that Player I's

position $\psi(h_t)$ in $G_i(f_i, c)$ has been defined. Take $a_t \in A(h_t)$. Let $R^*(h_t, a_t)$ denote the set of Player I's legal moves at position $\psi(h_t)$ to which σ_{II} responds with a_t :

$$R^*(h_t, a_t) := \{r_t \in R(h_t) : (\psi(h_t), r_t) \in T \text{ and } \sigma_{\text{II}}(\psi(h_t), r_t) = a_t\}.$$

The history (h_t, a_t) is defined to be σ_{II} -acceptable if $R(h_t, a_t)$ is not empty. In this case we define

$$u_{h_t}(a_t) := \inf\{r_t(a_t) : r_t \in R^*(h_t, a_t)\}.$$

Choose $r_t \in R^*(h_t, a_t)$ with the property that

$$u_{h_t}(a_t) \leq r_t(a_t) \leq u_{h_t}(a_t) + \epsilon \cdot 3^{-t-2}, \quad (7)$$

and define $\psi(h_t, a_t) := (\psi(h_t), r_t, a_t)$.

Finally, extend the definition of u_h to all histories h of Γ by setting $u_h(a) := 1$ whenever (h, a) is not σ_{II} -acceptable.

Let H_{II} be the set of σ_{II} -acceptable histories of Γ . We define the set O to be the complement of $[H_{\text{II}}]$, that is $O := [H] \setminus [H_{\text{II}}]$. Since $[H_{\text{II}}]$ is a closed subset of $[H]$ (e.g. Kechris [25, Proposition 2.4]), O is an open subset of $[H]$. Let $T_{\text{II}} \subseteq T$ be the image of H_{II} under ψ . The function $\psi_{\text{II}} : H_{\text{II}} \rightarrow T_{\text{II}}$ induces a continuous function $\psi_{\text{II}} : [H_{\text{II}}] \rightarrow [T_{\text{II}}]$. Note that all runs in $[T_{\text{II}}]$ are consistent with Player II's winning strategy σ_{II} .

For $t \in \mathbb{N}$ define a function $v_t : H_t \rightarrow \mathbb{R}$ by letting $v_0(\emptyset) := c$; and for each $t \in \mathbb{N}$ and each history $(h_t, a_t) \in H_{t+1}$, by letting $v_{t+1}(h_t, a_t) := u_{h_t}(a_t)$. Note that, for $t \in \mathbb{N}$ and $h_t \in H_t$, we have $v_t(h_t) = 1$ whenever h_t is not σ_{II} -acceptable.

Also define $g : [H] \rightarrow [0, 1]$ by letting

$$g(a_0, a_1, \dots) := \limsup_{t \rightarrow \infty} v_t(a_0, \dots, a_{t-1}).$$

STEP 2: Verifying that $f_i \leq g \leq 1$ and $\{f_i = 1\} \subseteq O \subseteq \{g = 1\}$.

The function g is equal to 1 on the set O ; thus $O \subseteq \{g = 1\}$. Consider a play $p = (a_0, a_1, \dots) \in [H_{\text{II}}]$, and let $\psi(p) := (r_0, a_0, r_1, a_1, \dots)$. It follows by (7) that

$$g(p) := \limsup_{t \rightarrow \infty} r_t(a_t).$$

Since the run $\psi(p)$ is won by Player II, it must hold that either $f_i(p) < g(p)$ or $0 = f_i(p)$; in either case $f_i(p) < 1$ and $f_i(p) \leq g(p)$. We conclude that $[H_{\text{II}}] \subseteq \{f_i < 1\}$, or equivalently that $\{f_i = 1\} \subseteq O$, and that $f_i \leq g$ on $[H]$.

STEP 3: Verifying that $v_i(g) \leq c + \epsilon$. Since $1_O \leq g$, it then follows that $v_i(O) \leq c + \epsilon$.

STEP 3.1: Defining a strategy profile for player i 's opponents.

First we argue that

$$d_i(u_{\emptyset}) \leq c. \quad (8)$$

Suppose to the contrary that $c \leq d_i(u_{\emptyset}) - \lambda$ for some $\lambda > 0$. Define $r_0 \in R(\emptyset)$ by letting $r_0(a) := \max\{u_{\emptyset}(a) - \lambda, 0\}$. Since $u_{\emptyset} - \lambda \leq r_0$, it holds that $c \leq d_i(u_{\emptyset}) - \lambda \leq d_i(r_0)$. Consequently, r_0 is a legal move of Player I in the game $G_i(f_i, c)$ at position \emptyset . Denote $a_0 := \sigma_{\text{II}}(r_0)$. As a_0 is Player II's legal move in $G_i(f_i, c)$ at position (r_0) , it must be the case that $r_0(a_0) > 0$, and

hence $r_0(a_0) = u_\emptyset(a_0) - \lambda$. On the other hand, $r_0 \in R^*(\emptyset, a_0)$, so the definition of u_\emptyset implies that $u_\emptyset(a_0) \leq r_0(a_0)$, a contradiction.

Take $t \geq 1$, let $h_t := (h_{t-1}, a_{t-1}) \in H_t$ be a σ_{II} -acceptable history, and let r_{t-1} be such that $\psi(h_t) = (\psi(h_{t-1}), r_{t-1}, a_{t-1})$. Then

$$d_i(u_{h_t}) \leq r_{t-1}(a_{t-1}). \quad (9)$$

Indeed, suppose to the contrary that $r_{t-1}(a_{t-1}) \leq d_i(u_{h_t}) - \lambda$ for some $\lambda > 0$. Define $r_t \in R(h_t)$ by letting $r_t(a) := \max\{u_{h_t}(a) - \lambda, 0\}$. Since $u_{h_t} - \lambda \leq r_t$, it holds that $r_{t-1}(a_{t-1}) \leq d_i(u_{h_t}) - \lambda \leq d_i(r_t)$. Consequently, r_t is a legal move of Player I at position $\psi(h_t)$. Let $a_t := \sigma_{\text{II}}(\psi(h_t), r_t)$. As a_t is Player II's legal move at position $(\psi(h_t), r_t)$, it must be the case that $r_t(a_t) > 0$, and hence $r_t(a_t) = u_{h_t}(a_t) - \lambda$. On the other hand, $r_t \in R^*(h_t, a_t)$, so the definition of u_{h_t} implies that $u_{h_t}(a_t) \leq r_t(a_t)$, a contradiction.

We now define a strategy profile σ_{-i} of i 's opponents in Γ as follows: For a history $h_t \in H_t$ of Γ let $\sigma_{-i}(h_t) \in X_{-i}(h)$ be such that

$$u_{h_t}(\sigma_{-i}(h_t), x_i) \leq d_i(u_{h_t}) + \epsilon \cdot 3^{-t-1} \text{ for each } x_i \in \Delta(A_i(h_t)). \quad (10)$$

STEP 3.2: Verifying that $\mathbb{E}_{\sigma_{-i}, \sigma_i}(g) \leq c + \epsilon$ for each strategy $\sigma_i \in \Sigma_i$ of player i in Γ .

Fix a strategy $\sigma_i \in \Sigma_i$. For $t \in \mathbb{N}$ define a function $v_t^\epsilon := v_t + \epsilon \cdot 3^{-t}$. The sequence $v_0^\epsilon, v_1^\epsilon, \dots$ could be thought of as a process on $[H]$, measurable with respect to the filtration $\{\mathcal{F}_t\}_{t \in \mathbb{N}}$. We next show that the process is a supermartingale w.r.t $\mathbb{P}_{\sigma_{-i}, \sigma_i}$.

By Eqs. (10) and (8),

$$\begin{aligned} \mathbb{E}_{\sigma_{-i}, \sigma_i}(v_1^\epsilon) &= \mathbb{E}_{\sigma_{-i}, \sigma_i}(u_\emptyset(a_0)) + \epsilon \cdot 3^{-1} \\ &= u_\emptyset(\sigma_{-i}(\emptyset), \sigma_i(\emptyset)) + \epsilon \cdot 3^{-1} \\ &\leq d_i(u_\emptyset) + \epsilon \cdot 2 \cdot 3^{-1} \\ &\leq c + \epsilon = v_0^\epsilon(\emptyset). \end{aligned}$$

Take $t \geq 1$, let $h_t = (h_{t-1}, a_{t-1}) \in H_t$ be a σ_{II} -acceptable history, and let r_{t-1} be such that $\psi(h_t) = (\psi(h_{t-1}), r_{t-1}, a_{t-1})$. We have by Eqs. (10), (9), and (7):

$$\begin{aligned} \mathbb{E}_{\sigma_{-i}, \sigma_i}(v_{t+1}^\epsilon \mid h_t) &= \mathbb{E}_{\sigma_{-i}, \sigma_i}(u_{h_t}(a_t) \mid h_t) + \epsilon \cdot 3^{-t-1} \\ &= u_{h_t}(\sigma_{-i}(h_t), \sigma_i(h_t)) + \epsilon \cdot 3^{-t-1} \\ &\leq d_i(u_{h_t}) + \epsilon \cdot 2 \cdot 3^{-t-1} \\ &\leq r_{t-1}(a_{t-1}) + \epsilon \cdot 2 \cdot 3^{-t-1} \\ &\leq u_{h_{t-1}}(a_{t-1}) + \epsilon \cdot 3 \cdot 3^{-t-1} \\ &= v_t^\epsilon(h_{t-1}, a_{t-1}) = v_t^\epsilon(h_t). \end{aligned}$$

If, on the other hand, the history h_t is not σ_{II} -acceptable, then

$$\mathbb{E}_{\sigma_{-i}, \sigma_i}(v_{t+1}^\epsilon \mid h_t) = 1 + \epsilon \cdot 3^{-t-1} \leq 1 + \epsilon \cdot 3^{-t} = v_t^\epsilon(h_t).$$

Since the process $v_0^\epsilon, v_1^\epsilon, \dots$ is bounded below (by 0), by the Martingale Convergence Theorem, it converges pointwise $\mathbb{P}_{\sigma_{-i}, \sigma_i}$ -almost surely; whenever the process converges, its limit is g . Hence $\mathbb{E}_{\sigma_{-i}, \sigma_i}(g) = \mathbb{E}_{\sigma_{-i}, \sigma_i}(\lim_{t \rightarrow \infty} v_t^\epsilon) \leq v_0^\epsilon(\emptyset) = c + \epsilon$, as desired. \square

We now invoke the result of Martin [35] on Borel determinacy of perfect information games. To do so, we endow $[T]$ with its relative topology as a subspace of the product space $(R \cup A)^{\mathbb{N}}$, where $R \cup A$ is given its discrete topology. One can then check that Player I's winning set in $G_i(f_i, c)$ is a Borel subset of $[T]$. It follows that for each $c \in (0, 1]$ the game $G_i(f_i, c)$ is determined: either Player I has a winning strategy in the game or Player II does. We arrive at the following conclusion.

Proposition 8.3. *If $v_i(f_i) < c$, then Player II has a winning strategy in $G_i(f_i, c)$. If $c < v_i(f_i)$, then Player I has a winning strategy in $G_i(f_i, c)$.*

Theorems 3.1 and 3.3 follow from Propositions 8.1, 8.2, and 8.3.

Proof of Theorem 3.4: Take an $\varepsilon > 0$. Without loss of generality, suppose that f_i takes values in $[0, 1]$.

By Proposition 4.4 we know that $v_i(Q_{i,\varepsilon}(f_i)) = 1$. To obtain an approximation from below, use Theorem 3.1 to choose a closed set $C \subseteq Q_{i,\varepsilon}(f_i)$ such that $1 - \varepsilon \leq v_i(C)$, and define the function $g := (v_i(f_i) - \varepsilon) \cdot 1_C$. Then $g \leq f_i$ and $v_i(f_i) - 2\varepsilon \leq (v_i(f_i) - \varepsilon) \cdot (1 - \varepsilon) \leq v_i(g)$. Since C is closed, g is upper semicontinuous.

By Proposition 4.4 we know that $v_i(U_{i,\varepsilon}(f_i)) = 0$. To obtain an approximation from above, use Theorem 3.1 to choose an open set $O \supseteq U_{i,\varepsilon}(f_i)$ such that $v_i(O) \leq \varepsilon$, and define the function $g := v_i(f_i) + \varepsilon + (1 - v_i(f_i) - \varepsilon) \cdot 1_O$. Then $f_i \leq g \leq 1$ and $v_i(g) \leq v_i(f_i) + 2\varepsilon$. Since O is open, g is lower semicontinuous. \square

References

- [1] Apt KR, Grädel E, eds. (2011): *Lectures in game theory for computer scientists*. Cambridge University Press.
- [2] Arieli I, Levy YJ (2015): Determinacy of games with stochastic eventual perfect monitoring. *Games and Economic Behavior* 91(1), 166–185.
- [3] Ashkenazi-Golan G, Flesch J, Predtetchinski A, Solan E (2022): Existence of equilibria in repeated games with long-run payoffs. *Proceedings of the National Academy of Sciences of the USA* 119(11).
- [4] Aumann RJ, Maschler M (1972): Some thoughts on the minimax principle. *Management Science* 18, 54–63.
- [5] Aumann RJ, Shapley L (1994): Long-term competition – a game-theoretic analysis. *Essays in game theory* 1–15. Springer, New York, NY.
- [6] Bertsekas DP, Shreve SE (1996): *Stochastic optimal control: the discrete-time case* (Vol. 5). Athena Scientific.
- [7] Blackwell D (1969): Infinite G_δ -games with imperfect information. *Applicationes Mathematicae* 1(10), 99–101.
- [8] Blackwell D, Diaconis P (1996): A non-measurable tail set. *Statistics, Probability and Game Theory IMS Lecture Notes - Monograph Series Volume 30*.

- [9] Bruyère V (2017): Computer aided synthesis: a game-theoretic approach. In *International Conference on Developments in Language Theory*, 3–35. Springer, Cham.
- [10] Bruyère V (2021): Synthesis of equilibria in infinite-duration: games on graphs. *ACM SIGLOG News* 8(2), 4–29.
- [11] Carathéodory C (1907): Über den Variabilitätsbereich der Koeffizienten von Potenzreihen, die gegebene Werte nicht annehmen. *Mathematische Annalen* 64(1), 95–115.
- [12] Chatterjee K (2005): Two-player nonzero-sum ω -regular games. In *CONCUR'05*, 413–427. LNCS 3653, Springer, 2005.
- [13] Chatterjee K (2007): Concurrent games with tail objectives. *Theoretical Computer Science* 388, 181–198.
- [14] Chatterjee K, Henzinger TA (2012): A survey of stochastic ω -regular games. *Journal of Computer and System Sciences* 78(2), 394–413.
- [15] Chatterjee K, Majumdar R, Jurdziński M (2004): On Nash equilibria in stochastic games. In *CSL'04*, 26–40. LNCS 3210, Springer.
- [16] Flesch J, Herings PJJ, Maes J, and Predtetchinski A (2021): Subgame maxmin strategies in zero-sum stochastic games with tolerance levels. *Dynamic Games and Applications* 1–34.
- [17] Flesch J, Solan E (2022): Equilibrium in two-player stochastic games with shift-invariant payoffs. arXiv:2203.14492.
- [18] Forges F (1995): Non-zero sum repeated games and information transmission, in *Essays in Game Theory: In Honor of Michael Maschler*, edited by Nimrod Megiddo, Springer Verlag, 65–95.
- [19] Fudenberg D, Maskin E (1986): The folk theorem in repeated games with discounting or with incomplete information. *Econometrica* 54, 533–554.
- [20] Gale D, Stewart FM (1953) Infinite games with perfect information. *Contributions to the Theory of Games* 2(245-266), 2–16.
- [21] Gillette D (1957): Stochastic games with zero stop probabilities. *Contributions to the Theory of Games* 3, 179–187.
- [22] Harsanyi JC (1986): *Rational behaviour and bargaining equilibrium in games and social situations*. Cambridge University Press.
- [23] Hausdorff F (2005): *Set theory*. American Mathematical Society, volume 119.
- [24] Hörner J, Sugaya T, Takahashi S, Vieille N (2011): Recursive methods in discounted stochastic games: An algorithm for $\delta \rightarrow 1$ and a folk theorem. *Econometrica* 79, 1277–318.
- [25] Kechris AS (1995): *Classical descriptive set theory*, Springer.
- [26] Kuipers J, Flesch J, Schoenmakers G, and Vrieze K (2021): Subgame perfection in recursive perfect information games. *Economic Theory* 71, 603–662.

- [27] Jaśkiewicz A, Nowak AS (2018): Non-zero-sum stochastic games. *Handbook of Dynamic Game Theory*, 1–64.
- [28] Lehrer E (1996): Mediated talk, *International Journal of Game Theory* 25, 177–188.
- [29] Lehrer E, Sorin S (1997): One-shot public mediated talk, *Games and Economic Behavior* 20, 131–148.
- [30] Levy YJ, Solan E (2020): Stochastic Games. In: Sotomayor M., Pérez-Castrillo D., Castiglione F. (eds) *Complex Social and Behavioral Systems. Encyclopedia of Complexity and Systems Science Series*. Springer, New York, NY.
- [31] Mailath GJ, Samuelson L (2006): *Repeated games and reputations: long-run relationships*. Oxford University Press.
- [32] Maitra A, Purves R, and Sudderth W (1992): Approximation theorems for gambling problems and stochastic games. *Lecture Notes in Economics and Mathematical Systems* 389, 114–132, Springer, Berlin.
- [33] Maitra A, Sudderth W (1993): Borel stochastic games with lim sup payoff. *The Annals of Probability* 21, 861–885.
- [34] Maitra A, Sudderth W (1998): Finitely additive stochastic games with Borel measurable payoffs. *International Journal of Game Theory* 27, 257–267.
- [35] Martin DA (1975): Borel determinacy. *Annals of Mathematics* 102, 363–371.
- [36] Martin DA (1998): The determinacy of Blackwell games. *The Journal of Symbolic Logic* 63, 1565–1581.
- [37] Maschler M, Solan E, Zamir S (2013): *Game theory*. Cambridge University Press.
- [38] Mashiah-Yaakovi A (2015): Correlated equilibria in stochastic games with Borel measurable payoffs. *Dynamic Games and Applications* 5, 120–135.
- [39] Mertens JF (1987): Repeated games. *Proceedings of the International Congress of Mathematicians*, 1986 Providence, RI. American Mathematical Society.
- [40] Mertens JF, Neyman A (1981): Stochastic games. *International Journal of Game Theory* 10, 53–66.
- [41] Nash JF (1950): Equilibrium points in n -person games. *Proceedings of the National Academy of Sciences of the USA* 36, 48–49.
- [42] Rosenthal J (1975): Nonmeasurable invariant sets. *The American Mathematical Monthly* 82, 488–491.
- [43] Rubinstein A (1994): Equilibrium in supergames. *Essays in Game Theory*, 17–27. Springer, New York, NY.
- [44] Secchi P, Sudderth WD (2002): Stay-in-a-set games. *International Journal of Game Theory* 30, 479–490.

- [45] Shmaya E (2011): The determinacy of infinite games with eventual perfect monitoring. *Proceedings of the American Mathematical Society* 139, 3665–3678.
- [46] Simon RS (2007): The structure of non-zero-sum stochastic games. *Advances in Applied Mathematics* 38(1), 1–26.
- [47] Simon, RS, Tomkowicz G (2018): A Bayesian Game without ϵ -equilibria. *Israel Journal of Mathematics* 227(1), 215–231.
- [48] Solan E (2001): Characterization of correlated equilibria in stochastic games. *International Journal of Game Theory* 30, 259–277.
- [49] Sorin S (1986) Asymptotic properties of a non-zero sum stochastic game. *International Journal of Game Theory* 15, 101–107.
- [50] Sorin S (1992): Repeated games with complete information. *Handbook of game theory with economic applications*, Volume I, 71–107.
- [51] Ummels M, Markey N, Brenguier R, Bouyer P (2015): Pure Nash equilibria in concurrent deterministic games. *Logical Methods in Computer Science* 11, 1–72.
- [52] Vieille N (2000): Two-player stochastic games I: A reduction. *Israel Journal of Mathematics* 119, 55–91.
- [53] Vieille N (2000): Two-player stochastic games II: The case of recursive games. *Israel Journal of Mathematics* 119, 93–126.