THE DETERMINACY OF INFINITE GAMES WITH EVENTUAL PERFECT MONITORING

ERAN SHMAYA

ABSTRACT. An infinite two-player zero-sum game with a Borel winning set, in which the opponent's actions are monitored eventually but not necessarily immediately after they are played, admits a value. The proof relies on a representation of the game as a stochastic game with perfect information, in which Nature operates as a delegate for the players and performs the randomizations for them.

1. SETUP

Consider an infinite two-player zero-sum game that is given by a triple $(A, (P_n)_{n \in \mathbb{N}}, W)$ where A is a finite set of *actions*, P_n is a partition of A^n for every $n \in \mathbb{N}$, and $W \subseteq A^{\mathbb{N}}$ is a Borel set, the *winning set* of player 1. The game is played in stages: Player 1 chooses an action $a_0 \in A$; then player 2 chooses an action $a_1 \in A$; then player 1 chooses an action $a_2 \in A$, and so on, ad infinitum. Before choosing a_n, the player who plays at stage n receives some information about his opponent's actions at previous stages: Let $h = (a_0, a_1, \ldots, a_{n-1})$ be the *finite history* that consists of the actions played before stage n; then

before choosing a_n, the player who plays at stage n observes the atom of P_n that contains h. Player 1 wins the game if the *infinite history* $(a_0, a_1, \dots)$ is in W. When the action set and information partitions are fixed, I denote the game by $\Gamma(W)$.

A *behavioral strategy* $x = (x_n)_{n \in \mathbb{N}}$ of player 1 is a sequence $\{x_n : P_n \to \Delta(A)\}_{n=0,2,4,\dots}$ of functions: At stage n, after observing the finite history $h = (a_0, a_1, \dots, a_{n-1})$, player 1 randomizes his action according to $x_n(\pi_n(h))$, where $\pi_n(h)$ is the atom of P_n that contains h. Abusing notations, I sometimes write $x_n(h)$ instead of $x_n(\pi_n(h))$. Behavioral strategies y of player 2 are defined analogously.

Every pair x, y of strategies induces a probability distribution $\mu_{x,y}$ over the set $A^{\mathbb{N}}$ of infinite histories or *plays*: $\mu_{x,y}$ is the joint distribution of a sequence $\alpha_0, \alpha_1, \dots \dots$ of A-valued random variables such that

$$
(1) \quad \mathbb{P}\left(\alpha_n = a \,|\, \alpha_0, \dots, \alpha_{n-1}\right) =
\begin{cases}
x_n(\alpha_0, \dots, \alpha_{n-1})[a], & \text{if } n \text{ is even,} \\
y_n(\alpha_0, \dots, \alpha_{n-1})[a], & \text{if } n \text{ is odd.}
\end{cases}
$$

I call such a sequence of random variables an (x, y)-*random play*. If the players play according to the strategy profile (x, y), then the expected payoff for player 1 is given by

$$
(2) \qquad\qquad \mu_{x,y}(W) = \mathbb{P}\left((\alpha_0, \alpha_1, \dots) \in W\right),
$$

where $\alpha_0, \alpha_1, \dots$ is an (x, y)-random play.

The *lower value* $\underline{\mathrm{val}}\ \Gamma(W)$ and *upper value* $\overline{\mathrm{val}}\ \Gamma(W)$ of the game $\Gamma(W)$ are defined by:

$$\underline{\mathrm{val}}\ \Gamma(W) = \sup_x \inf_y \mu_{x,y}(W), \quad \text{and } \overline{\mathrm{val}}\ \Gamma(W) = \inf_y \sup_x \mu_{x,y}(W),$$

where the suprema are taken over all strategies x of player 1 and the infima over all strategies y of player 2. The game is *determined* if the lower and upper values are equal, $\underline{\mathrm{val}}\ \Gamma(W) = \overline{\mathrm{val}}\ \Gamma(W)$, in which case their common value is called the *value* of the game. For $\epsilon \geq 0$, a strategy x of player 1 is ϵ-*optimal* if $\mu_{x,y}(W) \geq \underline{\mathrm{val}}\ \Gamma(W) - \epsilon$ for every strategy y of player 2. We also say that player 1 can *guarantee* payoff of at least $\underline{\mathrm{val}}\ \Gamma(W) - \epsilon$ by playing such a strategy x. ϵ-optimal strategies of player 2 are defined analogously.

Let $\sim_n$ be the equivalence relation over infinite histories such that $u \sim_n u'$ whenever $u|_n$ and $u'|_n$ belong to the same atom of P_n, where $u|_n$ and $u'|_n$ are the initial segments of u and u' of length n. The interpretation is that if $u, u' \in A^{\mathbb{N}}$ and $u \sim_n u'$, then at stage n the player cannot distinguish between u and u'. Say that *at stage n the player observes the action of stage m* if, for every pair of infinite histories $u = (a_0, a_1, \dots)$ and $u' = (a'_0, a'_1, \dots)$, $u \sim_n u'$ implies $a_m = a'_m$.

1.1. **Definition.** The information partitions $(P_n)_{n \geq 0}$ satisfy *perfect recall* if the following conditions are satisfied:

(1) Players know their own actions: at stage n the player observes the action of stage $n - 2$.

(2) Players do not forget information: if $u, u' \in A^{\mathbb{N}}$ and $u \sim_{n+2} u'$ then $u \sim_n u'$.

The setup of infinite games with perfect recall is general enough to subsume two special cases which have been extensively studied:

Borel games. If, at every stage n, players observe previous actions of their opponents, then the game is called a *Borel game* or a *game with perfect information*. Gale and Stewart [3] proved that such games are determined if the winning set W is closed. In a seminal paper, Martin [7] proved that the game is determined for every Borel winning set W. Borel games admit pure 0-optimal strategies, and the value is 0 or 1. Moreover, Borel games with an infinite action set A are also determined.

Blackwell games. Assume that at even stages $n = 2k$, player 1 observes the actions of stages $0, 1, \ldots, 2k - 1$, and at odd stages $n = 2k + 1$, player 2 observes the actions of stages $0, 1, \ldots, 2k - 1$ (his own actions and all the previous actions of his opponent except for the last one), and that the information partitions are the roughest partitions that satisfy these conditions. This means essentially that the players play *simultaneously* at stages $2k$ and $2k + 1$ for every $k \in \mathbb{N}$, and then both actions are announced. Such games are called *Blackwell games*. Blackwell [1, 2] proved the determinacy of Blackwell games (which he called "infinite games with imperfect information") with a G_δ winning set, and conjectured that every Blackwell game with a Borel winning set is determined. Vervoort [11] advanced higher in the Borel hierarchy, proving determinacy of games with $G_{\delta\sigma}$ winning sets. Blackwell's conjecture was proved by Martin in 1998 [8].

Borel games and Blackwell games differ in the timing of *monitoring* – the observation of the opponent's actions: whereas in Borel games monitoring is immediate, in Blackwell games player 2's monitoring is delayed by one stage. Both setups satisfy a property that I call *eventual perfect monitoring*. This means that the entire history of the game is known to every player at infinity. One example of eventual perfect monitoring, of which Blackwell games are a special case, is *delayed monitoring*, when the action of stage m is monitored after some delay d_m. But the setup of games with eventual perfect monitoring is more general than the setup of games with delayed monitoring. First, the former setup allows the length of the delay to depend on the history of the games. Second, it allows the information to be revealed in pieces; for example, a player can observe some function of the previous actions of his opponent before he observes the actions themselves.

1.2. **Definition.** The information partitions $\{P_n\}_{n \in \mathbb{N}}$ satisfy *eventual perfect monitoring* if for every $u, u' \in A^{\mathbb{N}}$ such that $u \neq u'$, there exist an even n such that $u \nsim_n u'$ and an odd n such that $u \nsim_n u'$.

The purpose of this paper is to prove the following theorem.

1.3. **Theorem.** *Let $\Gamma = (A, (P_n)_{n \geq 0}, W)$ be an infinite game with a finite action set, a Borel winning set, perfect recall, and eventual perfect monitoring. Then Γ is determined.*

The proof of the theorem relies heavily on the stochastic extension of Martin's theorem about the determinacy of Blackwell games. However, except for the simple case in which the stages are divided into

blocks and previous actions are monitored at the end of each block,
I was unable to find an immediate reduction of the eventual perfect
monitoring setup to the Blackwell games setup, nor was able to adapt
Martin's proof to the eventual perfect monitoring setup.

Infinite games with Borel winning sets have recently been applied in
economics literature on testing the quality of probabilistic predictions.
Consider a forecaster who claims to know the probability distribution
that governs some stochastic process. To prove his claim, the forecaster
provides probabilistic predictions about the process. An inspector tests
the forecaster's reliability using the infinite sequence of predictions pro-
vided by the forecaster and the observed realization of the process.
Using Martin's Theorem about the determinacy of Blackwell games,
I proved [9] that any inspection which is based on predictions about
the next-day realization of the process is manipulable, i.e., it can be
strategically passed by a charlatan. Theorem 1.3 can be used to prove
that tests based on predictions about an arbitrarily long finite horizon
are also manipulable [9, Section 5].

In Section 2 I give some examples of games with and without eventual
perfect monitoring. In Section 3 I prove the determinacy of infinite
games with perfect recall and a compact winning set; this result is
used in the proof of Theorem 1.3. The proof of the theorem is in
Section 4. Martin's Theorem is reviewed in the appendix.

2. Examples

All the examples in this section have the same action set and the
same winning set. The action set is $A = \{S, L\}$. At every stage,

each player decides whether to Stay or Leave the game. Once a player leaves, his future actions do not affect the outcome of the game. For an infinite history $u = (a_0, a_1, \dots)$, let $n^1(u) = \min\{n \text{ even } | a_n = \mathrm{L}\}$ be the (possibly infinite) first stage in which player 1 left the game, and let $n^2(u)$ be the first stage in which player 2 left. Let

$$W = \left\{ u \in A^{\mathbb{N}} \,\middle|\, \left(n^2(u) < n^1(u) < \infty\right) \text{ or } \left(n^1(u) < \infty \text{ and } n^2(u) = \infty\right) \right\}$$

be the winning set of player 1. So player 1 wins if he leaves the game after player 2 leaves, or if he leaves the game at some point and player 2 never leaves. In Example 2.1 both players have eventual perfect monitoring. In Example 2.2 none of the players has eventual perfect monitoring. In Example 2.3 only player 1 has eventual perfect monitoring.

2.1. **Example.** Let k be a positive integer. Assume that at stage n each player observes his own actions and the actions of his opponent at stages smaller than $n - k$. Then the value of the game is 0. An optimal strategy for player 2 is to play S as long as he is not informed that player 1 has played L. When player 2 knows that player 1 played L at some point, player 2 then plays L.

Note that in the previous example, the number k need not be constant. It can depend on the stage number, and can differ between the players. As long as player 2 knows the actions of his opponent eventually, the game is determined and the value is 0. (The fact that the value in this example does not depend on the information partitions is not typical.)

2.2. **Example.** Assume that each player knows his own previous actions, but does not observe his opponent's actions. Then the game is not determined. In fact, $\underline{\text{val}}\ \Gamma = 0$ and $\overline{\text{val}}\ \Gamma = 1$.

2.3. **Example.** Assume that player 1 observes the past actions of player 2, but player 2 doesn't observe the past actions of player 1. Then the game is not determined. In fact, $\underline{\text{val}}\ \Gamma = 1/2$ and $\overline{\text{val}}\ \Gamma = 1$. An optimal strategy for player 1 is: At stage 0 play L or S with probability $1/2$, and, at stage $2k$ for $k \geq 1$, play the action of player 2 from stage $2k-1$.

3. Games with a compact winning set

The set $A^{\mathbb{N}}$ of plays is naturally endowed with the product topology. In this section I prove the special case of Theorem 1.3 for compact winning sets. The determinacy follows from perfect recall alone, even without eventual perfect monitoring. The proof relies on two standard results from game theory: the Minimax Theorem for normal form games and Kuhn's Theorem.

Recall that a *normal form game* is given by a triple (Σ, Θ, R) where Σ and Θ are Borel spaces of *pure strategies* for players 1 and 2, and $R : \Sigma \times \Theta \to [0,1]$ is the *payoff function*. A *mixed strategy* ξ of player 1 is a probability distribution over Σ. Mixed strategies τ of player 2 are defined analogously. Say that the mixed extension of the normal form game (Σ, Θ, r) is *determined* if

$$\sup_{\xi \in \Delta(\Sigma)} \inf_{\theta \in \Theta} \int R(\sigma, \theta) \xi(\mathrm{d}\sigma) = \inf_{\tau \in \Delta(\Theta)} \sup_{\sigma \in \Sigma} \int R(\sigma, \theta) \tau(\mathrm{d}\theta).$$

The Minimax Theorem [10, Proposition A.10] states that if Σ is a compact topological space and the function $R(\cdot, \theta)$ is upper semicontinuous for every $\theta \in \Theta$, then the mixed extension of the normal form game (Σ, Θ, R) is determined.

Let $\Gamma = (A, \{P_n\}_{n\in\mathbb{N}}, W)$ be an infinite game with perfect recall. The *normal form of* Γ is the normal form game $N(\Gamma) = (\Sigma, \Theta, R)$ defined as follows. A pure strategy $\sigma \in \Sigma$ of player 1 is a sequence $\{\sigma_n : P_n \to A\}_{n \text{ even}}$ of functions: at stage n, after the finite history $h = (h_0, h_1, \ldots, h_{n-1})$ was played, player 1 plays $\sigma_n(\pi_n(h))$, where $\pi_n(h)$ is the atom of P_n that contains h. Pure strategies θ of player 2 are defined analogously. Every pair σ, θ of pure strategies of players 1 and 2 determines an infinite history $u(\sigma, \theta) = (a_0, a_1, \ldots)$ that is given by

$$
a_n = \begin{cases}
\sigma_n\left(\pi_n\left(a_0, \ldots, a_{n-1}\right)\right), & \text{for even } n, \\
\theta_n\left(\pi_n\left(a_0, \ldots, a_{n-1}\right)\right), & \text{for odd } n.
\end{cases}
$$

The payoff function of $N(\Gamma)$ is $R(\sigma, \theta) = \mathbf{1}_W(u(\sigma, \theta))$. Kuhn's Theorem [10, Theorem D.1] states the equivalence between mixed strategies and behavioral strategies in games with perfect recall. In particular, the game Γ is determined if and only if its normal form game $N(\Gamma)$ is determined.

3.1. **Lemma.** *An infinite game with a finite action set, perfect recall, and a compact winning set is determined.*

Proof. Let $\Gamma = (A, \{P_n\}_{n\in\mathbb{N}}, W)$ be an infinite game with a finite action set A, perfect recall, and a compact winning set W. By Kuhn's

Theorem it is sufficient to prove that the normal form game $N(\Gamma) = (\Sigma, \Theta, R)$ of Γ is determined. This follows from the minimax theorem. Indeed, the set Σ of pure strategies of player 1 is compact in the product topology, and the payoff function R is upper semicontinuous as a composition of the continuous function $(\sigma, \theta) \mapsto u(\sigma, \theta)$ and the function $\mathbf{1}_W$, which is upper semi continuous because W is closed. $\quad\square$

4. Proof of Theorem 1.3

Overview of the proof. Roughly speaking, I am going to construct a stochastic game Γ^* with perfect information that mimics the original game Γ. In Γ^*, at every stage m, the player announces a mixture over A contingent on his information at that stage. So in Γ^*, instead of choosing an action which is not revealed his the opponent (as in Γ), the player announces how he intends to randomize his action. The actual randomization is performed by Nature at a future stage $k(m)$, in which the opponent should have observed the m-stage action in Γ, and the realization of that randomization is immediately made public. So in the game Γ^*, Nature performs the randomization for the player. By Martin's Theorem the game Γ^* is determined, and I prove that the value of Γ^* is also the value of the original game Γ. For this purpose I have to show that the fact that in Γ^* the player announces his randomization plan cannot be used by the opponent to change the payoff in the game. This step, which is the core of the proof, uses approximations of the winning set by compact sets, and the fact that by Lemma 3.1 the original game Γ is determined when the winning set is compact.

Since the sets of actions must be finite for Martin's Theorem to apply, I first prove that every behavioral strategy in Γ can be approximated by a behavioral strategy in which all the mixtures are taken from some finite sets. This is done in Lemma 4.2. Because of the approximation argument, the stochastic game Γ^* that is constructed in the proof depends on an additional parameter ϵ which corresponds to the level of approximation.

Preliminaries. Let $A^{<\mathbb{N}} = \bigcup_{n \in \mathbb{N}} A^n$ be the set of finite histories of the game. For a finite history $h \in A^n$, the *length of h* is given by $\text{length}(h) = n$. For an infinite history $u = (a_0, a_1, a_2, \dots) \in A^{\mathbb{N}}$ and $n \in \mathbb{N}$, let $u|_n = (a_0, \dots, a_{n-1}) \in A^{<\mathbb{N}}$ be the initial segment of u of length n. Similarly, for a finite history $h \in A^{<\mathbb{N}}$ and $n < \text{length}(h)$, let $h|_n$ be the initial segment of h of length n.

Eventual perfect monitoring entails that the action of stage m is known to the opponent at infinity. Lemma 4.1 below shows that in fact more is true: for every m there exists some finite stage $n > m$ at which the opponent knows the action of stage m.

4.1. **Lemma.** *If the game admits eventual perfect monitoring, then for every $m \in \mathbb{N}$ there exists an $n > m$ such that $n \neq m \mod 2$ and such that at stage n the opponent observes the action of stage m, i.e., for every pair $u = (a_0, a_1, \dots), u' = (a'_0, a'_1, \dots)$ of infinite histories, $u \sim_n u'$ implies $a_m = a'_m$.*

Proof. This is an application of König's Lemma. Assume without loss of generality that m is odd. Let $a \in A$, and let $C_a = \{u = (u_0, u_1, \dots) \in$

$A^{\mathbb{N}}|u_m = a\}$. Then C_a and C_a^c are compact. Let $T_a \subseteq A^{<\mathbb{N}}$ be the set of all finite histories h of even length n such that $\pi_n^{-1}(h) \cap C_a \neq \emptyset$ and $\pi_n^{-1}(h) \cap C_a^c \neq \emptyset$, where $\pi_n(h)$ is the atom of P_n that contains h.

It follows from the perfect recall assumption that T_a is a tree over A^2. I claim that T_a is well-founded. Indeed, if v is an infinite branch of T, then $\bigcap_{n \text{ even}} \pi_n^{-1}(v|_n) \cap C_a$ and $\bigcap_{n \text{ even}} \pi_n^{-1}(v|_n) \cap C_a^c$ are nonempty as the intersections of decreasing sequences of compact sets. Let $u = (u_0, u_1, \dots) \in \bigcap_{n \text{ even}} \pi_n^{-1}(v|_n) \cap C_a$ and $u' = (u_0', u_1', \dots) \in \bigcap_{n \text{ even}} \pi_n^{-1}(v|_n) \cap C_a^c$. Then $u_m = a \neq u_m'$ and therefore $u \neq u'$, but $u \sim_n u'$ for every even n, in contradiction to the eventual perfect monitoring assumption.

By König's Lemma, T_a is finite. Let n^a be the maximal length of elements of T_a, and let $n = \max\{n^a | a \in A\} + 2$. Then at stage n player 1 observes the action of stage m. $\square$

For two strategies x, x' of player 1, let $d(x, x')$, the *distance between x and x'*, be given by

$$d(x, x') = \sum_{n \text{ even}} \max_{p \in P_n} \|x_n(p) - x_n'(p)\|_1,$$

where the maximum is taken over all atoms p of P_n. The distance $d(y, y')$ between two behavioral strategies y, y' of player 2 is defined analogously.

4.2. **Lemma.** *Let x, x' be strategies of player 1 and y, y' be strategies of player 2. Then*

$$\|\mu_{x,y}(W) - \mu_{x',y'}(W)\| \leq \big(d(x, x') + d(y, y')\big)/2$$

for every Borel subset W of $A^{\mathbb{N}}$.

Proof. The idea is to join a (x, y)-random play and a (x', y')-random play such that the two random plays are equal with high probability. Let $z_n : P_n \to A$ be given by $z_n = x_n$ for even n's and $z_n = y_n$ for odd n's and $z'_n : P_n \to A$ be given by $z'_n = x'_n$ for even n's and $z'_n = y'_n$ for odd n's. Let $\alpha_0, \alpha'_0, \alpha_1, \alpha'_1, \dots$ be a sequence of A-valued random variables defined inductively such that the conditional joint distribution of the pair (α_n, α'_n) given the event $\{\alpha_i = a_i, \alpha'_i = a'_i \text{ for } 0 \le i < n\}$ satisfies

(3)
$$\mathbb{P}\left(\alpha_n = a \,|\, \alpha_i = a_i, \alpha'_i = a'_i \text{ for } 0 \le i < n\right) = z_n(a_0, \dots, a_{n-1})[a],$$

(4)
$$\mathbb{P}\left(\alpha'_n = a' \,|\, \alpha_i = a_i, \alpha'_i = a'_i \text{ for } 0 \le i < n\right) = z'_n(a'_0, \dots, a'_{n-1})[a'], \text{ and}$$

(5)
$$\mathbb{P}\left(\alpha'_n \ne \alpha_n \,|\, \alpha_i = a_i, \alpha'_i = a'_i \text{ for } 0 \le i < n\right) \le \|z_n(a_0, \dots, a_{n-1}) - z'_n(a'_0, \dots, a'_{n-1})\|_1/2,$$

for every n and every $a_0, a'_0, \dots, a_{n-1}, a'_{n-1} \in A$. The existence of random variables α_n, α'_n with the prescribed conditional distribution follows from a standard coupling argument [4, Theorem 5.2]. From (3) it follows that

$$\mathbb{P}\left(\alpha_n = a \,|\, \alpha_i = a_i \text{ for } 0 \le i < n\right) = z_n(a_0, \dots, a_{n-1})[a]$$

for every n and every $a_0, \dots, a_{n-1} \in A$, i.e., that $\alpha_0, \alpha_1, \dots$ is an (x, y)-random play of $\Gamma(W)$. Similarly, from (4) it follows that $\alpha'_0, \alpha'_1, \dots$ is

a (x', y')-random play of $\Gamma(W)$. From (5) it follows that

$$\mathbb{P}\left(\alpha_n \neq \alpha'_n \,|\, \alpha_i = \alpha'_i \text{ for } 0 \leq i < n\right) \leq \max_{p \in P_n} \|z_n(p) - z'_n(p)\|_1/2.$$

Therefore,

$$\mathbb{P}\left(\alpha_n \neq \alpha'_n \text{ for some } n\right) \leq \sum_{n \in \mathbb{N}} \mathbb{P}\left(\alpha_n \neq \alpha'_n \,|\, \alpha_i = \alpha'_i \text{ for } 0 \leq i < n\right)$$

$$\leq \sum_n \max_{p \in P_n} \|z_n(p) - z'_n(p)\|_1/2 = \left(d(x, x') + d(y, y')\right)/2.$$

The assertion follows from the last inequality and the fact that $\mu_{x,y}$ and $\mu_{x',y'}$ are the distributions of $\alpha_0, \alpha_1, \dots$ and $\alpha'_0, \alpha'_1, \dots$, respectively. $\square$

4.3. **Corollary.** *Let $\Delta_{\epsilon,n}$ be a finite set which is $\epsilon/2^n$-dense in $\Delta(A)$ endowed with $\|\|_1$, i.e., such that the $\epsilon/2^n$-balls around elements of $\Delta_{\epsilon,n}$ cover $\Delta(A)$. Then there exists an ϵ-optimal strategy y for player 2 in $\Gamma(W)$ such that $y_n(p) \in \Delta_{\epsilon,n}(A)$ for every odd n and every atom p of P_n.*

Proof. Let y' be an $\epsilon/2$-optimal strategy of player 2 in $\Gamma(W)$ and let y be a strategy of player 2 such that $\|y_n(p) - y'_n(p)\|_1 < \epsilon/2^n$ and $y_n(p) \in \Delta_{\epsilon,n}(A)$ for every odd n and every atom p of P_n. Then $d(y, y') < \epsilon$, and therefore,

$$\mu_{x,y}(W) \geq \mu_{x,y'}(W) - \epsilon/2 \geq \overline{\mathrm{val}}\Gamma(W) - \epsilon$$

for every strategy x of player 1, where the first inequality follows from Lemma 4.2, and the second inequality from the fact that y' is $\epsilon/2$-optimal. Therefore y is ϵ-optimal. $\square$

Nature as the players' randomization delegate. Let $\Gamma = (A, P_n, W)$ be an infinite game with perfect recall and eventual perfect monitoring. In this section, I define an auxiliary stochastic game $\Gamma^*_\epsilon = \Gamma^*_\epsilon(W)$ with perfect information, which mimics the original games Γ.

Fix $\epsilon > 0$ and, for every $n \in \mathbb{N}$, let $\Delta_{\epsilon,n}$ be a finite set which is $\epsilon/2^n$-dense in $\Delta(A)$ endowed with $\|\|_1$, i.e., such that the $\epsilon/2^n$-balls around elements of $\Delta_{\epsilon,n}$ cover $\Delta(A)$. For every $m \in \mathbb{N}$, fix $k(m) > m$ such that $m \neq k(m) \mod 2$, and such that at stage $k(m)$ the opponent observes the action of stage m, as in Lemma 4.1.

For every n, let $B_n = \{b : P_n \to \Delta_{\epsilon,n}\}$ be the set of *actions* of stage n in $\Gamma^*_\epsilon(W)$, so that an action is a function from P_n (viewed as a collection of atoms) to $\Delta_{\epsilon,n}$; and let $S_n = A^{K_n}$ be the set of *states* of stage n in $\Gamma^*_\epsilon(W)$, where $K_n = \{m | k(m) = n\}$. Let $f_n : A^n \to S_n$ be the projection over the corresponding coordinates $m \in K_n$, and let $F : S_0 \times S_1 \cdots \to A^{\mathbb{N}}$ be such that

$$(6) \qquad F(f_0(u|_0), f_1(u|_1), \dots) = u$$

for every $u \in A^{\mathbb{N}}$.

$\Gamma^*_\epsilon(W)$ is played as follows: Player 1 plays at even stages and player 2 at odd stages. At every stage n, Nature announces a state s_n in S_n, and then the player that play at that stage announces an action b_n in B_n. Nature chooses the state s_n of stage n from the distribution

$z\left(s_0, b_0, \ldots, s_{n-1}, b_{n-1}\right)$ that is given by

(7) $z\left(s_0, b_0, \ldots, s_{n-1}, b_{n-1}\right)[s] =$

$\quad \mathbb{P}\left(f_n\left(\bar{\alpha}_0, \ldots, \bar{\alpha}_{n-1}\right) = s \,|\, f_k(\bar{\alpha}_0, \ldots \bar{\alpha}_{k-1}) = s_k \text{ for } 0 \le k < n-1\right),$

where $\bar{\alpha}_0, \ldots, \bar{\alpha}_n$ is a sequence of A-valued random variables such that

(8) $\mathbb{P}\left(\bar{\alpha}_k = a \,|\, \bar{\alpha}_0, \ldots, \bar{\alpha}_{k-1}\right) = b_k\left(\pi_k(\bar{\alpha}_0, \ldots, \bar{\alpha}_{k-1})\right)[a],$

where $\pi_k(h)$ is the atom of P_k that contains h for every $h \in A^k$. Player 1 wins the game if $F\left(s_0, s_1, \ldots\right) \in W$.

A *pure strategy* of player 1 in $\Gamma_\epsilon^*(W)$ is a sequence $\{x_n^* : S_0 \times B_0 \times \cdots \times S_{n-1} \times B_{n-1} \times S_n \to B_n\}_{n=0,2,\ldots}$ of functions: at stage n, after observing the *finite history* $(s_0, b_0, \ldots, s_{n-1}, b_{n-1}, s_n)$, player 1 plays $x^*\left(s_0, b_0, \ldots, s_{n-1}, b_{n-1}, s_n\right)$. Pure strategies y^* of Player 2 are defined analogously. Let X^* and Y^* be the sets of pure strategies of players 1 and 2 respectively. The expected payoff for player 1 in the game $\Gamma_\epsilon^*(W)$ when the players play according to (x^*, y^*) is given by $R(x^*, y^*) = \mathbb{P}\left(F(\zeta_0, \zeta_1, \ldots) \in W\right)$ where $\beta_0, \zeta_0, \beta_1, \zeta_1, \ldots$ is a sequence of random variables, where the values of β_n are in B_n and the values of ζ_n are in S_n such that

$\quad \mathbb{P}\left(\zeta_n = s \,|\, \zeta_0, \beta_0, \ldots, \zeta_{n-1}, \beta_{n-1}\right) = z\left(\zeta_0, \beta_0, \ldots, \zeta_{n-1}, \beta_{n-1}\right)[s],$

$\quad \beta_n = x_n^*\left(\zeta_0, \beta_0, \ldots, \zeta_{n-1}, \beta_{n-1}, \zeta_n\right) \text{ for even } n, \text{ and}$

$\quad \beta_n = y_n^*\left(\zeta_0, \beta_0, \ldots, \zeta_{n-1}, \beta_{n-1}, \zeta_n\right) \text{ for odd } n.$

I call such a sequence $\zeta_0, \beta_0, \zeta_1, \beta_1, \ldots$ of random variables an (x^*, y^*)-*random play* of $\Gamma_\epsilon^*(W)$.

Identifying the game $\Gamma_\epsilon^*(W)$ with its normal form, say that $\Gamma_\epsilon^*(W)$ is *determined* if

$$\sup_{\xi \in \Delta(X^*)} \inf_{y^* \in Y^*} \int R(x^*, y^*) \xi(\mathrm{d}x^*) = \inf_{\tau \in \Delta(Y^*)} \sup_{x^* \in X^*} \int R(x^*, y^*) \tau(\mathrm{d}y^*).$$

In this case the common value of the two sides of the last equations is called the *value* of the game, and is denoted by $\mathrm{val}\,\Gamma_\epsilon^*(W)$.

4.4. **Lemma.** *Let $W \subseteq A^{\mathbb{N}}$ be a Borel set. Then the game $\Gamma_\epsilon^*(W)$ is determined, and $\mathrm{val}\,\Gamma_\epsilon^*(W_0) > \mathrm{val}\,\Gamma_\epsilon^*(W) - \epsilon$ for some compact subset W_0 of W.*

Proof. In the terminology of appendix A, the game $\Gamma_\epsilon^*(W)$ is the stochastic game with stochastic setup $\mathcal{S} = ((S_n, B_n)_{n \in \mathbb{N}}, z)$ and the winning set $\eta^{-1}(W)$, where $\eta : S_0 \times B_0 \times S_1 \times B_1 \times \cdots \to A^{\mathbb{N}}$ is the continuous map given by

(9) $$\eta(s_0, b_0, s_1, b_1, \ldots) = F(s_0, s_1, \ldots).$$

Thus $\eta^{-1}(W)$ is a Borel set and therefore by Proposition A.1 the game $(\mathcal{S}, \eta^{-1}(W))$ is determined. Moreover, there exists a compact set $C \subseteq S_0 \times B_0 \times S_1 \times B_1 \times \ldots$ such that $C \subseteq \eta^{-1}(W)$ and $\mathrm{val}(\mathcal{S}, C) > \mathrm{val}(\mathcal{S}, \eta^{-1}(W)) - \epsilon$. Let $W_0 = \eta(C)$. Then W_0 is a compact subset of W and $\mathrm{val}(\mathcal{S}, \eta^{-1}(W_0)) \geq \mathrm{val}(\mathcal{S}, C) > \mathrm{val}(\mathcal{S}, \eta^{-1}(W)) - \epsilon$, since $\eta^{-1}(W_0) \supseteq C$. The assertion follows from the fact that the games $(\mathcal{S}, \eta^{-1}(W_0))$ and $(\mathcal{S}, \eta^{-1}(W))$ are $\Gamma_\epsilon^*(W_0)$ and $\Gamma_\epsilon^*(W)$, respectively. $\square$

The following lemma says that, up to ϵ, player 2 can guarantee in Γ^*_ϵ the same amount he can guarentee in Γ. Intuitively, when player 2 computes the upper value of the game, he assumes that player 1 is going to play the best response to player 2's strategy; so the fact that in $\Gamma^*_\epsilon(W)$ player 2 has to declare his contingent mixed action does not reduce the upper value of the game.

4.5. **Lemma.** *For every Borel set W of $A^{\mathbb{N}}$,*

$$val\ \Gamma^*_\epsilon(W) - \epsilon \leq \overline{val}\ \Gamma(W).$$

Proof. Note first that by definition of K_n and from the perfect recall assumption, there exist functions $g_{n,k} : P_n \to S_k$ for every n and every $k \leq n$ such that

(10) $$g_{n,k}(\pi_n(a_0, \ldots, a_{n-1})) = f_k(a_0, \ldots, a_{k-1})$$

for every $h = (a_0, \ldots, a_{n-1}) \in A^n$ and where $\pi_n : A^n \to P_n$ is the natural projection.

Let y be an ϵ-optimal behavioral strategy for player 2 in $\Gamma(W)$ such that $y_n(p) \in \Delta_{\epsilon,n}(A)$ for every odd n and every atom p of P_n. The existence of such a strategy y follows from Corollary 4.3. Consider a pure strategy y^* of player 2 in $\Gamma^*_\epsilon(W)$ that is given by $y^*_n(s_0, b_0, \ldots, s_{n-1}, b_{n-1}, s_n) = y_n$ for every odd n and every partial history $(s_0, b_0, \ldots, s_{n-1}, b_{n-1}, s_n)$ of $\Gamma^*_\epsilon(W)$. (Thus, in every odd stage n, player 2's action is y_n, regardless of the history.) Let x^* be any strategy of player 1 in Γ^*. Let x be the

behavioral strategy of player 1 in $\Gamma(W)$ that is given by

$$x_n(p) = x_n^*(s_0, b_0, \ldots, s_{n-1}, b_{n-1}, s_n)(p),$$

where $(s_0, b_0, \ldots, s_{n-1}, b_{n-1}, s_n)$ is the finite history of $\Gamma_p^*(W)$ defined inductively by $b_k = x_k^*(s_0, b_0, \ldots, s_{k-1}, b_{k-1}, s_k)$ for even k, $b_k = y_k$ for odd k, and $s_k = g_{n,k}(p)$.

I am going to join an (x, y)-random play of $\Gamma(W)$ and an (x^*, y^*)-random play of $\Gamma_\epsilon^*(W)$ with equal payoffs. Let $\zeta_0, \beta_0, \alpha_0, \zeta_1, \beta_1, \alpha_1, \ldots$ be a sequence of random variables such that the values of ζ_n are in S_n, the values of β_n are in B_n, and the values of α_n are in A, and such that

(11) $\zeta_n = f_n(\alpha_0, \ldots, \alpha_{n-1})$,

(12) $\beta_n = x_n^*(\zeta_0, \beta_0, \ldots, \zeta_{n-1}, \beta_{n-1}, \zeta_n)$ for even n,

(13) $\beta_n = y_n^*(\zeta_0, \beta_0, \ldots, \zeta_{n-1}, \beta_{n-1}, \zeta_n)$ for odd n, and

(14) $\mathbb{P}(\alpha_n = a \,|\, \alpha_0, \ldots, \alpha_{n-1}) = \beta_n(\Pi_n)$.

From (10) and (11) it follows that

(15) $\zeta_k = g_{n,k}(\pi_n(\alpha_0, \ldots, \alpha_{n-1}))$

for every n and every $k \leq n$. From (13) and the definition of y^*, it follows that $\beta_n = y_n$ for every odd n. In particular,

(16) $y_n(\pi_n(\alpha_0, \ldots, \alpha_{n-1})) = \beta_n(\pi_n(\alpha_0, \ldots, \alpha_{n-1}))$

for every odd n. From (12), the definition of x, the fact that $\beta_k = y_k$ for every odd k, and (15), it follows that

$$(17) \qquad x_n \left(\pi_n \left(\alpha_0, \ldots, \alpha_{n-1} \right) \right) = \beta_n \left(\pi_n \left(\alpha_0, \ldots, \alpha_{n-1} \right) \right)$$

for every even n. From (14), (16), (17), it follows that

$$\mathbb{P} \left(\alpha_n = a \, | \alpha_0, \ldots, \alpha_{n-1} \right) = \begin{cases} x_n(\pi_n(\alpha_0, \ldots, \alpha_{n-1})), & n \text{ even,} \\ \\ y_n(\pi_n(\alpha_0, \ldots, \alpha_{n-1})) & n \text{ odd,} \end{cases}$$

i.e., that $\alpha_0, \alpha_1, \ldots$ is an (x, y)-random play of $\Gamma(W)$.

From (11),(12),(13),(14), and (7), it follows that

$$(18) \quad \mathbb{P}(\zeta_n = s_n | \zeta_0, \beta_0, \ldots, \zeta_{n-1}, \beta_{n-1}) = z \left(\zeta_0, \beta_0, \ldots, \zeta_{n-1}, \beta_{n-1} \right).$$

Indeed, given the event $\{ \zeta_0 = s_0, \beta_0 = b_0, \ldots, \zeta_{n-1} = s_{n-1}, \beta_{n-1} = b_{n-1} \}$, the conditional distribution of $\alpha_0, \ldots, \alpha_n$ is like the conditional distribution of a sequence $\bar{\alpha}_0, \ldots, \bar{\alpha}_n$ that satisfies (8) given that $f_k(\bar{\alpha}_0) = s_k$ for $k < n$. (Here I use the fact that β_n is measurable with respect to $\zeta_0, \ldots, \zeta_n$.)

From (18),(12), and (13) it follows that $\zeta_0, \beta_0, \zeta_1, \beta_1, \ldots$ is an (x^*, y^*)-random play of $\Gamma_\epsilon(W)$. Therefore, the expected payoff for player 1 in $\Gamma_\epsilon^*(W)$ under (x^*, y^*) is

$$\mathbb{P} \left(F(\zeta_0, \zeta_1, \ldots) \in W \right) = \mathbb{P} \left((\alpha_0, \alpha_1, \ldots) \in W \right) = \mu_{x,y}(W) \leq \overline{\text{val}} \, \Gamma(W) + \epsilon,$$

where the first equality follows from (6) and (11), the second equality from (2), and the inequality from the fact that y is ϵ-optimal.

Summing up, I have provided a pure strategy y^* of player 2 in $\Gamma_\epsilon^*(W)$ (namely, play $y_1, y_3, \ldots$) that gives expected payoff of at most $\overline{\text{val}}\,\Gamma(W) + \epsilon$ against any pure strategy x^* of player 1 in $\Gamma_\epsilon^*(W)$. Therefore, $\text{val}\,\Gamma_\epsilon^*(W) \leq \overline{\text{val}}\,\Gamma(W) + \epsilon$. $\square$

Proof of Theorem 1.3. Consider the stochastic game $\Gamma_\epsilon^*(W)$ defined above. Let W_0 be a compact subset of W such that $\text{val}\,\Gamma_\epsilon^*(W_0) > \text{val}\,\Gamma_\epsilon^*(W) - \epsilon$, and whose existence follows from Lemma 4.4. Then

$$\underline{\text{val}}\,\Gamma(W) \geq \underline{\text{val}}\,\Gamma(W_0) = \overline{\text{val}}\,\Gamma(W_0) \geq \text{val}\,\Gamma_\epsilon^*(W_0) - \epsilon > \text{val}\,\Gamma_\epsilon^*(W) - 2\epsilon,$$

where the first inequality follows from the fact that $W \supseteq W_0$, the first equality follows from Lemma 3.1, the second inequality follows from Lemma 4.5, and the third inequality follows from the choice of W_0.

Similarly, for player 2 we get $\overline{\text{val}}\,\Gamma(W) < \text{val}\,\Gamma_\epsilon^*(W) + 2\epsilon$. It follows that $\overline{\text{val}}\,\Gamma(W) < \underline{\text{val}}\,\Gamma(W) + 4\epsilon$. Since ϵ was arbitrary, it follows that $\overline{\text{val}}\,\Gamma(W) = \underline{\text{val}}\,\Gamma(W)$. $\square$

Appendix A. Martin's Theorem for Stochastic Games

In this section I formulate Martin's Theorem about the determinacy of stochastic games. The stochastic game used in this paper has complete information, while Martin studied a more general setup in which the players play simultaneously. Note, however, that Martin's older theorem about the determinacy of Borel games [7] is not sufficient for my purposes because of the presence of Nature.

A *stochastic game with perfect information* is given by $((S_n, B_n)_{n \in \mathbb{N}}, z, V)$, where $B_0, B_1, \ldots$ are finite sets of *actions*, $S_0, S_1, \ldots$ are finite sets of

states or *Nature's actions,* $z = \{z_n : S_0 \times B_0 \times \cdots \times S_{n-1} \times B_{n-1} \to \Delta(S_n)\}$ is *Nature's strategy,* and $V \subseteq S_0 \times B_0 \times S_1 \times B_1 \times \ldots$ is the *winning set* of Player 1.

The game is played as follows: Player 1 plays at even stages and player 2 at odd stages. At every stage n, Nature announces a state s_n in S_n, and then the player that play at that stage announces an action b_n in B_n. Nature chooses the state s_n of stage n from the distribution $z(s_0, b_0, \ldots, s_{n-1}, b_{n-1})$. Player 1 wins the game if $(s_0, b_0, s_1, b_1, \ldots) \in V$.

I call a triple $\mathcal{S} = ((S_n, B_n)_{n \in \mathbb{N}}, z,)$ of action sets, states sets, and Nature's strategy a *stochastic setup.* So the stochastic games that I use in this paper are given by a stochastic setup $\mathcal{S} = ((S_n, B_n)_{n \in \mathbb{N}}, z\})$ and a winning set $V \subseteq S_0 \times B_0 \times S_1 \times B_1 \times \ldots$.

The definitions of strategies of the players, and of determinacy and value of the game, are omitted. Note that since this is a game with perfect information (i.e., before a player chooses the action b_n of stage n, he observes the finite history of the game $(s_0, b_0, \ldots, s_{n-1}, b_{n-1}, s_n)$ up to that stage), Kuhn's Theorem [10, Theorem D.1] applies, so that behavioral strategies and mixed strategies are equivalent.

The following proposition was proved by Martin [8]. For the stochastic extension, see Maitra and Sudderth's paper [6]. The fact that the lower value of the game can be approximated by the value on some compact subset was proved earlier by Maitra et al. [5], using Choquet's Capacity Theorem.

A.1. **Proposition.** *Let $\mathcal{S} = ((S_n, B_n)_{n \in \mathbb{N}}, z\})$ be a stochastic setup, and let V be a Borel subset of $S_0 \times B_0 \times S_1 \times B_1 \times \ldots$. Then:*

(1) *The game $(\mathcal{S}, V)$ is determined.*

(2) *For every $\epsilon > 0$, there exists a compact subset C of V such that*

$$val(\mathcal{S}, C) > val(\mathcal{S}, V) - \epsilon.$$

References

[1] David Blackwell. Infinite G_δ games with imperfect information. *Zastosowania Matematyki (Applicationes Mathematicae)*, 10:99–101, 1969.

[2] David Blackwell. Operator solution of infinite G_δ games of imperfect information. In Theodore W. Anderson, Krishna B. Athreya, and Donald L. Igelhart, editors, *Probability, Statistics, and Mathematics*, pages 83–87. Academic Press, Boston, Massachusetts, 1989.

[3] David Gale and F.M. Stewart. Infinite games with perfect information. *Annals of Mathematical Studies*, 28, 1953.

[4] Torgny Lindvall. *Lectures on the Coupling Method*. Wiley, New York, 1992.

[5] Ashok P. Maitra, Roger A. Purves, and William D. Sudderth. Approximation theorems for gambling problems and stochastic games. *Lecture Notes in Economics and Mathematical Systems*, 389:114–132, 1992.

[6] Ashok P. Maitra and William D. Sudderth. Finitely additive stochastic games with Borel measurable payoffs. *International Journal of Game Theory*, 27(2):257–267, July 1998.

[7] Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.

[8] Donald A. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63(4):1565–1581, December 1998.

[9] Eran Shmaya. Many inspections are manipulable. *Theoretical Economics*, 3(3):376–382, 2008.

[10] Sylvain Sorin. *A First Course on Zero-Sum Repeated Games.* Springer, Heidelberg, 2002.

[11] Marco R. Vervoort. Blackwell games. In Thomas S. Ferguson, Lloyd S. Shapley, and James B. MacQueen, editors, *Statistics, Probability, and Game Theory*, pages 369–390. Institute of Mathematical Statistics, Hayward, California, 1996.

E-mail address: e-shmaya@kellogg.northwestern.edu

Kellogg School of Management, Northwestert University