THE DETERMINACY OF INFINITE GAMES WITH EVENTUAL PERFECT MONITORING

ERAN SHMAYA

(Communicated by Ken Ono)

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, is determined. The proof relies on a representation of the game as a stochastic game with perfect information, in which Chance operates as a delegate for the players and performs the randomizations for them, and on Martin's Theorem about the determinacy of such games.

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 $n \in \mathbb{N}$, the *information partition of stage 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 the actions of previous stages: Let $h = (a_0, a_1, \dots, a_{n-1})$ 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, we denote the game by $\Gamma(W)$.

A *behavioral strategy* $x = (x_n)_{n \in \mathbb{N}}$ of player 1 is a sequence $\{x_n : P_n \rightarrow \Delta(A)\}_{n=0,2,4,\dots}$ 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 notation, we sometimes write $x_n(h)$ instead of $x_n(\pi_n(h))$. The 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 $\alpha_0, \alpha_1, \dots$

Received by the editors July 5, 2010.

2010 *Mathematics Subject Classification.* Primary 91A15, 03E75; Secondary 03E15, 91A60.

Key words and phrases. Infinite games, determinacy, stochastic games, imperfect monitoring.

I thank the anonymous referee for helpful suggestions and comments, and also thank Chris Chambers, Tzachi Gilboa, John Levy, Ehud Lehrer, Wojciech Olszewski, Phil Reny, Eilon Solan, Bill Sudderth and Rakesh Vohra.

©2011 American Mathematical Society
 Reverts to public domain 28 years from publication

of A -valued random variables such that

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

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

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

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

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

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

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, $\underline{\text{val}} \Gamma(W) = \overline{\text{val}} \Gamma(W)$, 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 $\mu_{x,y}(W) \geq \underline{\text{val}} \Gamma(W) - \epsilon$ for every strategy y of player 2. We also say that player 1 can *guarantee* payoff of at least $\underline{\text{val}} \Gamma(W) - \epsilon$ by playing such a strategy x . The 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 cases which have been extensively studied:

Gale-Stewart games. If, at every stage n , players observe the previous actions of their opponents, then the game is called a *Gale-Stewart game* or a *game with perfect information*. Gale and Stewart [4] proved that such games are determined if the winning set W is closed (or open) and asked whether the result can be extended to more complex winning sets. In a seminal paper, Martin [8] proved that the game is determined for every Borel winning set W . Gale-Stewart games admit pure 0-optimal strategies, and the value is 0 or 1. Moreover, Gale-Stewart games with an infinite action set A and Borel winning set are also determined.

Blackwell games. Assume that at even stages $n = 2k$, player 1 observes the actions of stages $0, 1, \dots, 2k - 1$, and at odd stages $n = 2k + 1$, player 2 observes the actions of stages $0, 1, \dots, 2k - 1$ (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 essentially means that the players choose actions simultaneously and independently at stages $2k$ and $2k + 1$, and then both actions are made public. 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 [15] 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 [9].

Gale-Stewart games and Blackwell games differ in the timing of *monitoring* – the observation of the opponent’s actions: whereas in Gale-Stewart 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*, which 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*, introduced by Scarf and Shapley [12], when the action of stage m is monitored after some lag 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 lag 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 exists an n such that $u \approx_k u'$ for every $k > n$.

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 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 I able to adapt Martin’s proof to the eventual perfect monitoring setup. The difficulty stems from the fact that Martin’s proof uses the existence of *proper subgames*, i.e., occurrences along the play path where the current partial history is commonly known. However, even under simple monitoring structures, such as when one player observes the opponent’s action with two periods delay, the game does not have proper subgames.

Infinite games with Borel winning sets have recently been used in the 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 provided by the forecaster and the observed realization of the process. Using Martin’s Theorem about the determinacy of Blackwell games, I proved [13] 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 of this paper can be used to prove that inspections based on daily predictions about an arbitrarily long finite horizon are also manipulable [13, Section 5]. On the other hand, Olszewski and Sandroni [10] gave an example for a non-manipulable inspection which is based on a prediction about the infinite future. Such an inspection relies on a game without eventual perfect monitoring. In Section 2 are some examples of games with and without eventual perfect monitoring. In Section 3 we 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. Section 5 discusses the role of eventual perfect monitoring in the proof. Section 6 discusses some open questions. Martin's Theorem is reviewed in the appendix.

2. ENDURANCE GAMES

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 $n^1(u) = \min \{n \text{ even} \mid a_n = L\}$ 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

$$(3) \quad W = \{u \in A^{\mathbb{N}} \mid (n^2(u) < n^1(u) < \infty) \text{ or } (n^1(u) < \infty \text{ and } n^2(u) = \infty)\}$$

be the winning set of player 1. The game models a 'last man standing' endurance contest, as in a college drinking game or religious conflicts: Player 1 wants to leave the game at some point, but he doesn't want to be the first to leave.

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 independence of the value on the information partitions in this example 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. The proof uses 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 $R : \Sigma \times \Theta \rightarrow [0, 1]$ is a Borel *payoff function*. A *mixed strategy* ξ of player 1 is a probability distribution over Σ . The mixed strategies τ of player 2 are defined analogously. The mixed extension of the normal form game (Σ, Θ, R) is *determined* if

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

It follows from Fan's Minimax Theorem [3, Theorem 2] 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 $\{\sigma_n : P_n \rightarrow A\}_{n \text{ even}}$ of functions: at stage n , after the finite history $h = (h_0, h_1, \dots, h_{n-1})$ was played, player 1 plays $\sigma_n(\pi_n(h))$, where $\pi_n(h)$ is the atom of P_n that contains h . The 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 $u(\sigma, \theta) = (a_0, a_1, \dots)$ that is given by

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

The payoff function of $N(\Gamma)$ is $R(\sigma, \theta) = \mathbf{1}_W(u(\sigma, \theta))$. Kuhn's Theorem [14, 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 semicontinuous because W is closed. $\square$

4. PROOF OF THEOREM 1.3

Throughout this section, we fix a game Γ that satisfies the conditions of Theorem 1.3.

Overview of the proof. 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, regardless of the play path, the opponent knows the action of stage m . This means that at such a stage n the action taken by the player at stage m becomes *common knowledge*: The opponent knows the action played at stage m , the player that played the action knows that the opponent knows, the opponent knows that the player that played the action knows that the opponent knows, ad infinitum.

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 the set A of actions contingent on his information at that stage. So in Γ^* , instead of choosing an action which is not revealed to his opponent (as in Γ), the player announces how he intends to randomize his action. The actual randomization is performed by Chance at a future stage $k(m)$, in which the action should become commonly known, and the realization of that randomization is immediately made public. So in the game Γ^* , Chance performs the randomization for the player. The determinacy of Γ^* follows from Martin's Theorem, 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 .

4.1. Lemma. *For every $m \in \mathbb{N}$ there exists an $n > m$ such that $n \not\equiv m \pmod{2}$ 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. Assume w.l.o.g. that m is odd. Let $a \in A$, and let $C_a = \{u = (u_0, u_1, \dots) \in A^{\mathbb{N}} \mid u_m = a\}$. 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 $\pi_k(h)$ is the atom of P_k that contains h .

It follows from the perfect recall assumption that T_a is a tree over A^2 . We claim that T_a is well-founded. Indeed, if v is an infinite branch of T_a , then

$\bigcap_{n \geq m} \pi_n^{-1}(v|_n) \cap C_a$ and $\bigcap_{n \geq m} \pi_n^{-1}(v|_n) \cap C_a^c$ are nonempty as the intersections of decreasing sequences of compact sets. Let $u = (u_0, u_1, \dots) \in \bigcap_{n \geq m} \pi_n^{-1}(v|_n) \cap C_a$ and $u' = (u'_0, u'_1, \dots) \in \bigcap_{n \geq m} \pi_n^{-1}(v|_n) \cap C_a^c$. Then $u_m = a \neq u'_m$ and therefore $u \neq u'$, but $u \sim_n u'$ for every 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 $n = \max\{n^a | a \in A\} + 1$. Then from stage n onwards the opponent observes the action of stage m . $\square$

Approximating strategies. 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*

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

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

Proof. The idea is to join an (x, y) -random play and an (x', y') -random play such that the two random plays are equal with high probability. Let $z_n : P_n \rightarrow A$ be given by $z_n = x_n$ for even n 's and $z_n = y_n$ for odd n 's and $z'_n : P_n \rightarrow A$ 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 $\{\alpha_i = a_i, \alpha'_i = a'_i \text{ for } 0 \leq i < n\}$ satisfies

- (4) $\mathbb{P}(\alpha_n = a | \alpha_i = a_i, \alpha'_i = a'_i \text{ for } 0 \leq i < n) = z_n(a_0, \dots, a_{n-1})[a],$
- (5) $\mathbb{P}(\alpha'_n = a' | \alpha_i = a_i, \alpha'_i = a'_i \text{ for } 0 \leq i < n) = z'_n(a'_0, \dots, a'_{n-1})[a'], \text{ and}$
- (6) $\mathbb{P}(\alpha'_n \neq \alpha_n | \alpha_i = a_i, \alpha'_i = a'_i \text{ for } 0 \leq i < n)$
 $\leq \|z_n(a_0, \dots, a_{n-1}) - z'_n(a'_0, \dots, a'_{n-1})\|_1/2,$

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 [5, Theorem 5.2]. From (4) it follows that

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

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 (5) it follows that $\alpha'_0, \alpha'_1, \dots$ is an (x', y') -random play of $\Gamma(W)$. From (6) it follows that

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

Therefore,

$$\begin{aligned} \mathbb{P}(\alpha_n \neq \alpha'_n \text{ for some } n) &\leq \sum_{n \in \mathbb{N}} \mathbb{P}(\alpha_n \neq \alpha'_n \mid \alpha_i = \alpha'_i \text{ for } 0 \leq i < n) \\ &\leq \sum_n \max_{p \in P_n} \|z_n(p) - z'_n(p)\|_1 / 2 = (d(x, x') + d(y, y')) / 2. \end{aligned}$$

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 $\|\cdot\|_1$, 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,

$$\mu_{x,y}(W) \leq \mu_{x,y'}(W) + \epsilon/2 \leq \overline{\text{val}}\Gamma(W) + \epsilon$$

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. $\square$

Chance 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 game Γ .

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

For every n , let $B_n = \{b : P_n \rightarrow \Delta_{\epsilon,n}\}$ 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 $K_n = \{m \mid k(m) = n\}$. Let $f_n : A^n \rightarrow S_n$ be the projection over the corresponding coordinates $m \in K_n$, and let $F : S_0 \times S_1 \cdots \rightarrow A^\mathbb{N}$ be such that

$$(7) \quad F(f_0(u|_0), f_1(u|_1), \dots) = u$$

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 , Chance announces a state s_n in S_n , and then the player that plays at that stage announces an action b_n in B_n . Chance chooses the state s_n of stage n from the distribution $z(s_0, b_0, \dots, s_{n-1}, b_{n-1})$ that is given by

$$(8) \quad \begin{aligned} &z(s_0, b_0, \dots, s_{n-1}, b_{n-1})[s] \\ &= \mathbb{P}(f_n(\bar{\alpha}_0, \dots, \bar{\alpha}_{n-1}) = s \mid f_k(\bar{\alpha}_0, \dots, \bar{\alpha}_{k-1}) = s_k \text{ for } 0 \leq k < n-1), \end{aligned}$$

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

$$(9) \quad \mathbb{P}(\bar{\alpha}_k = a \mid \bar{\alpha}_0, \dots, \bar{\alpha}_{k-1}) = b_k(\pi_k(\bar{\alpha}_0, \dots, \bar{\alpha}_{k-1}))[a].$$

A *pure strategy* of player 1 in $\Gamma_\epsilon^*(W)$ is a sequence $\{x_n^* : S_0 \times B_0 \times \cdots \times S_{n-1} \times B_{n-1} \times S_n \rightarrow B_n\}_{n=0,2,\dots}$ of functions: at stage n , after observing the *finite history* $(s_0, b_0, \dots, s_{n-1}, b_{n-1}, s_n)$, player 1 plays $x^*(s_0, b_0, \dots, s_{n-1}, b_{n-1}, s_n)$. 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 $R(x^*, y^*) = \mathbb{P}(F(\zeta_0, \zeta_1, \dots) \in W)$, where $\zeta_0, \beta_0, \zeta_1, \beta_1, \dots$ 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

$$\begin{aligned} \mathbb{P}(\zeta_n = s \mid \zeta_0, \beta_0, \dots, \zeta_{n-1}, \beta_{n-1}) &= z(\zeta_0, \beta_0, \dots, \zeta_{n-1}, \beta_{n-1})[s], \\ \beta_n &= x_n^*(\zeta_0, \beta_0, \dots, \zeta_{n-1}, \beta_{n-1}, \zeta_n) \text{ for even } n, \text{ and} \\ \beta_n &= y_n^*(\zeta_0, \beta_0, \dots, \zeta_{n-1}, \beta_{n-1}, \zeta_n) \text{ for odd } n. \end{aligned}$$

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

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

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

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 $\text{val } \Gamma_\epsilon^*(W)$.

4.4. Lemma. *Let $W \subseteq A^\mathbb{N}$ be a Borel set. Then the game $\Gamma_\epsilon^*(W)$ is determined, and $\text{val } \Gamma_\epsilon^*(W_0) > \text{val } \Gamma_\epsilon^*(W) - \epsilon$ 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 $\eta : S_0 \times B_0 \times S_1 \times B_1 \times \cdots \rightarrow A^\mathbb{N}$ is the continuous map given by

$$(10) \quad \eta(s_0, b_0, s_1, b_1, \dots) = F(s_0, s_1, \dots).$$

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 $C \subseteq S_0 \times B_0 \times S_1 \times B_1 \times \cdots$ such that $C \subseteq \eta^{-1}(W)$ and $\text{val}(\mathcal{S}, C) > \text{val}(\mathcal{S}, \eta^{-1}(W)) - \epsilon$. Let $W_0 = \eta(C)$. Then W_0 is a compact subset of W and $\text{val}(\mathcal{S}, \eta^{-1}(W_0)) \geq \text{val}(\mathcal{S}, C) > \text{val}(\mathcal{S}, \eta^{-1}(W)) - \epsilon$, 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 guarantee in Γ .

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

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

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

$$(11) \quad g_{n,k}(\pi_n(a_0, \dots, a_{n-1})) = f_k(a_0, \dots, a_{k-1})$$

for every $h = (a_0, \dots, a_{n-1}) \in A^n$ and where $\pi_n : A^n \rightarrow P_n$ 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 $y_n^*(s_0, b_0, \dots, s_{n-1}, b_{n-1}, s_n) = y_n$ for every odd n and every partial history $(s_0, b_0, \dots, s_{n-1}, b_{n-1}, s_n)$ 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_ϵ^* . Let x be the behavioral strategy of player 1 in $\Gamma(W)$ that is given by

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

where $(s_0, b_0, \dots, s_{n-1}, b_{n-1}, s_n)$ is the finite history of $\Gamma_\epsilon^*(W)$ defined inductively by $b_k = x_k^*(s_0, b_0, \dots, s_{k-1}, b_{k-1}, s_k)$ 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 $\Pi_0, \zeta_0, \beta_0, \alpha_0, \Pi_1, \zeta_1, \beta_1, \alpha_1, \dots$ 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

$$(12) \quad \zeta_n = f_n(\alpha_0, \dots, \alpha_{n-1}),$$

$$(13) \quad \beta_n = x_n^*(\zeta_0, \beta_0, \dots, \zeta_{n-1}, \beta_{n-1}, \zeta_n) \text{ for even } n,$$

$$(14) \quad \beta_n = y_n^*(\zeta_0, \beta_0, \dots, \zeta_{n-1}, \beta_{n-1}, \zeta_n) \text{ for odd } n, \text{ and}$$

$$(15) \quad \mathbb{P}(\alpha_n = a | \alpha_0, \dots, \alpha_{n-1}) = \beta_n(\Pi_n)[a].$$

From (11) and (12) it follows that

$$(16) \quad \zeta_k = g_{n,k}(\pi_n(\alpha_0, \dots, \alpha_{n-1}))$$

for every n and every $k \leq n$. From (14) and the definition of y^* , it follows that $\beta_n = y_n$ for every odd n . In particular,

$$(17) \quad y_n(\pi_n(\alpha_0, \dots, \alpha_{n-1})) = \beta_n(\pi_n(\alpha_0, \dots, \alpha_{n-1}))$$

for every odd n . From (13), the definition of x , (16), and the fact that $\beta_k = y_k$ for every odd k , it follows that

$$(18) \quad x_n(\pi_n(\alpha_0, \dots, \alpha_{n-1})) = \beta_n(\pi_n(\alpha_0, \dots, \alpha_{n-1}))$$

for every even n . From (15), (17), (18), it follows that

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

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

From (12), (13), (14), (15), and (8), it follows that

$$(19) \quad \mathbb{P}(\zeta_n = s_n | \zeta_0, \beta_0, \dots, \zeta_{n-1}, \beta_{n-1}) = z(\zeta_0, \beta_0, \dots, \zeta_{n-1}, \beta_{n-1})[s_n].$$

Indeed, given the event $\{\zeta_0 = s_0, \beta_0 = b_0, \dots, \zeta_{n-1} = s_{n-1}, \beta_{n-1} = b_{n-1}\}$, the conditional distribution of $\alpha_0, \dots, \alpha_{n-1}$ is like the conditional distribution of a sequence $\bar{\alpha}_0, \dots, \bar{\alpha}_{n-1}$ that satisfies (9) given that $f_k(\bar{\alpha}_0) = s_k$ for $k < n$. (Here we use the fact that β_n is measurable with respect to $\zeta_0, \dots, \zeta_n$.)

From (19), (13), and (14) it follows that $\zeta_0, \beta_0, \zeta_1, \beta_1, \dots$ 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

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

where the first equality follows from (7) and (12), 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 $y_1, y_3, \dots$) 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 $\underline{\text{val}} \Gamma(W) < \underline{\text{val}} \Gamma(W) + 4\epsilon$. Since ϵ was arbitrary, it follows that $\underline{\text{val}} \Gamma(W) = \overline{\text{val}} \Gamma(W)$. $\square$

5. THE ROLE OF EVENTUAL PERFECT MONITORING

In the auxiliary game that I use in the proof of Theorem 1.3, Chance plays the role of the players ‘randomization delegate’, and all the actions of the players are perfectly monitored. What happens if we try to apply the same argument to a game without eventual perfect monitoring? Consider the following game, which is a modification of the game in Example 2.3 with Chance as player 1’s randomization delegate.

5.1. Example. Let $A = \{S, L\}$ and consider the game with perfect monitoring that is played in stages $n = 0, 1, \dots$ as follows: Player 1 plays at even stages and player 2 plays at odd stages. At every even stage n , player 1 declares a mixed action $p_n \in [0, 1]$, viewed as a die over A (p_n is the probability that player 1 will leave the game at stage n if he did not stop already). At every odd stage n , player 2 declares an action $A = \{S, L\}$. At infinity, Chance tosses the dice that player 1 declared one after another, thus obtaining an infinite history $h \in A^\infty$. Player 1 wins the game if $h \in W$, where W is given by (3).

Unlike the original game of Example 2.3, the game of Example 5.1 is determined and has value $1/2$. An optimal strategy for player 2 is to leave the game at the first stage n for which $(1 - p_0) \cdot (1 - p_2) \cdot \dots \cdot (1 - p_{n-1}) < 1/2$. However, as we saw, in the game of Example 2.3, where player 1’s action is not monitored by player 2, player 2 cannot guarantee to pay a payoff smaller than 1.

So, the transition from imperfect monitoring to Chance acting as the player’s randomization delegate might fundamentally change the strategic potential of the game in the absence of eventual perfect monitoring. On the other hand, in games with eventual perfect monitoring, as with finite games, a player does not lose anything if he declares his mixed contingency plan and lets Chance do the randomization for him.

The idea that a player declares his contingency mixed plan and lets Chance randomize for him can be traced back to von Neumann and Morgenstern’s discussion of mixed actions. They argue that a player who fears his plan will be found out or deduced by the opponent may want to use randomization. When the player allows Chance to choose his action, the opponent cannot possibly find out what the action is going to be, since the player does not know it himself. “Ignorance is obviously

a very good safeguard against disclosing information directly or indirectly” [16, Section 17.2.1]. The fact that the game is determined shows that the player who follows the Founders’ advice does not lose anything if his plan is indeed found out, so he might as well declare it. In the proof of Theorem 1.3 we walked the reverse path: Determinacy follows from the fact that a player does not lose anything if he declares his contingency plan. Example 5.1 shows that the argument relies on the assumption of eventual perfect monitoring.

6. OPEN QUESTIONS

The following questions about possible extensions of Theorem 1.3 are interesting for their own sake, and also because they highlight the assumptions on the game that were used in the proof of Theorem 1.3.

The first question is about relaxing eventual perfect monitoring. Say that an infinite game $\Gamma = (A, (P_n)_{n \geq 0}, W)$ admits *eventual outcome monitoring* if both players know the identity of the winner (that is, the outcome of the game) at infinity: for every $u, u' \in A^\mathbb{N}$ such that $u \in W$ and $u' \notin W$, there exists an n such that $u \approx_k u'$ for every $k > n$.

Question. Is every game with perfect recall and eventual outcome monitoring determined?

The construction of the auxiliary game Γ^* in the proof of Theorem 1.3 relies on Lemma 4.1, that if both players know the play path at infinity, then each action becomes commonly known at some finite time, and so the play path in Γ can be reconstructed from Chance’s actions in Γ^* . However, when both players know the occurrence of W at infinity, it still needs not be the case that W depends only on things that are commonly known at some finite times.

The second question relates to general payoff functions. Assume that the game is played as in Section 1, and at infinity, player 2 pays player 1 the amount $f(a_0, a_1, \dots)$, where $f : A^\mathbb{N} \rightarrow [0, 1]$ is a Borel function, the *payoff function* of the game. The games studied in this paper are a special case, where $f = \mathbf{1}_W$ is an indicator function. The definition of determinacy for games with Borel payoff function is a straightforward extension of the definition in Section 1.

Question. Is every game with eventual perfect monitoring and general payoff function determined?

In the proof of Theorem 1.3 we approximated the winning set W from below by a closed set, and we used Lemma 3.1 that games with closed winning sets are determined. Lemma 3.1 and its proof hold for upper semicontinuous payoff functions. However, upper semicontinuous functions are not sufficient to approximate arbitrary payoff functions in Martin’s Theorem about Blackwell Games (one needs to use G_δ -functions [9, p. 1576, Remark (c)]).

One can also combine the two questions: What about game with general payoff function, perfect recall and eventual payoff monitoring (i.e., for all histories $u, u' \in A^\mathbb{N}$ such that $f(u) \neq f(u')$, there exists an n such that $u \approx_k u'$ for every $k > n$)? In this case the game need not be determined. Indeed, Rosenberg and Vieille give an example of a nondetermined recursive game with incomplete information on one side [11, Section 4.1, henceforth RV]. To modify RV’s game to an infinite game with eventual payoff monitoring, replace Nature’s randomization in RV’s game with two

stages in which player 1 and then player 2 choose a bit from $\{0, 1\}$ and the sum of these bits modulo 2 is the state of nature. Player 1's initial choice is not monitored by player 2, but player 2's choice is monitored by player 1 so that player 1 knows the state of nature. Then the game proceeds as in RV's game until there is an absorption, in which case player 2 is notified of the initial bit chosen by player 1. If there is no absorption, then at infinity player 2 will not know the initial bit of player 1, but will know that the payoff is 0. Since by choosing his initial bit uniformly every player can guarantee that the sum modulo 2 is uniform, the indeterminacy of RV's game implies the indeterminacy of the modified game.

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 Gale-Stewart games [8] is not sufficient for my purposes because of the presence of Chance.

A *stochastic game with perfect information* is given by $((S_n, B_n)_{n \in \mathbb{N}}, z, V)$, where $B_0, B_1, \dots$ are finite sets of *actions*, $S_0, S_1, \dots$ are finite sets of *states* or *Chance's actions*, $z = \{z_n : S_0 \times B_0 \times \dots \times S_{n-1} \times B_{n-1} \rightarrow \Delta(S_n)\}$ is *Chance's strategy*, and $V \subseteq S_0 \times B_0 \times S_1 \times B_1 \times \dots$ 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 , Chance announces a state s_n in S_n , and then the player that plays at that stage announces an action b_n in B_n . Chance chooses the state s_n of stage n from the distribution $z(s_0, b_0, \dots, s_{n-1}, b_{n-1})$. Player 1 wins the game if $(s_0, b_0, s_1, b_1, \dots) \in V$.

I call a triple $\mathcal{S} = ((S_n, B_n)_{n \in \mathbb{N}}, z)$ of action sets *states sets* and Chance'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 $V \subseteq S_0 \times B_0 \times S_1 \times B_1 \times \dots$.

The definitions of strategies of the players, and of determinacy and value of the game, are omitted.

The following proposition was proved by Martin [9]. For the stochastic extension, see Maitra and Sudderth's paper [7]. The fact that the lower value of the game can be approximated by the value on some compact subset was proved earlier by Maitra, Purves and Sudderth [6].

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 $S_0 \times B_0 \times S_1 \times B_1 \times \dots$. Then:*

- (1) *The game $(\mathcal{S}, V)$ is determined.*
- (2) *For every $\epsilon > 0$, there exists a compact subset C of V such that*

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

REFERENCES

- [1] David Blackwell. Infinite G_δ games with imperfect information. *Zastosowania Matematyki (Applicationes Mathematicae)*, 10:99–101, 1969. MR0243837 (39:5158)
- [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, MA, 1989. MR1031279 (91d:90107)

- [3] Ky Fan. Minimax theorems. *Proceedings of the National Academy of Sciences U.S.A.*, 39:42–47, 1953. MR0055678 (14:1109f)
- [4] David Gale and F.M. Stewart. *Infinite Games with Perfect Information*. Annals of Mathematics Studies, 28, Princeton Univ. Press, Princeton, NJ, 1953. MR0054922 (14:999b)
- [5] Torgny Lindvall. *Lectures on the Coupling Method*. Wiley, New York, 1992. MR1180522 (94c:60002)
- [6] 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, Springer, Berlin, 1992. MR1228669 (94g:90189)
- [7] 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. MR1645524 (2000d:91020)
- [8] Donald A. Martin. Borel determinacy. *Annals of Mathematics (2)*, 102(2):363–371, 1975. MR0403976 (53:7785)
- [9] Donald A. Martin. The determinacy of Blackwell games. *Journal of Symbolic Logic*, 63(4):1565–1581, December 1998. MR1665779 (2000d:03118)
- [10] Wojciech Olszewski and Alvaro Sandroni. A nonmanipulable test. *Ann. Statist.*, 37(2):1013–1039, 2009. MR2502659 (2010c:62011)
- [11] Dinah Rosenberg, Eilon Solan, and Nicolas Vieille. The MaxMin value of stochastic games with imperfect monitoring. *Internat. J. Game Theory*, 32(1):133–150, 2003. Special anniversary issue. Part 2. MR2046684 (2005f:91026)
- [12] Herbert Scarf and L.S. Shapley. *Games with Partial Information*. Annals of Mathematics Studies, 39:213–229, Princeton, Univ. Press, Princeton, NJ, 1957. MR0094260 (20:779)
- [13] Eran Shmaya. Many inspections are manipulable. *Theoretical Economics*, 3(3):376–382, 2008.
- [14] Sylvain Sorin. *A First Course on Zero-Sum Repeated Games*. Springer, Berlin-Heidelberg, 2002. MR1890574 (2002m:91001)
- [15] 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. MR1481790 (98j:90100)
- [16] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton Univ. Press, Princeton, NJ, 2004. Reprint of the 1980 edition, Princeton Univ. Press, Princeton, NJ. MR0565457.

KELLOGG SCHOOL OF MANAGEMENT, NORTHWESTERN UNIVERSITY, EVANSTON, ILLINOIS 60208
 E-mail address: e-shmaya@kellogg.northwestern.edu