

האוניברסיטה העברית בירושלים
THE HEBREW UNIVERSITY OF JERUSALEM

**DETERMINACY OF GAMES WITH
STOCHASTIC EVENTUAL PERFECT
MONITORING**

By

ITAI ARIELI and YEHUDA (JOHN) LEVY

Discussion Paper # 658 January 2014

מרכז לחקר הרציונליות

**CENTER FOR THE STUDY
OF RATIONALITY**

Feldman Building, Givat-Ram, 91904 Jerusalem, Israel
PHONE: [972]-2-6584135 FAX: [972]-2-6513681
E-MAIL: ratio@math.huji.ac.il
URL: <http://www.ratio.huji.ac.il/>

Determinacy of Games with Stochastic Eventual Perfect Monitoring*

Itai Arieli[†] and Yehuda John Levy^{‡§}

January 20, 2014

Abstract

We consider an infinite two-player stochastic zero-sum game with a Borel winning set, in which the opponent's actions are monitored via stochastic private signals. We introduce two conditions of the signalling structure: Stochastic Eventual Perfect Monitoring (SEPM) and Weak Stochastic Eventual Perfect Monitoring (WSEPM). When signals are deterministic these two conditions coincide and by a recent result due to [Shmaya (2011)] entail determinacy of the game. We generalize [Shmaya (2011)]'s result and show that in the stochastic learning environment SEPM implies determinacy while WSEPM does not.

1 Introduction

The issue of the existence of the value in zero-sum games of infinite duration has a long and rich history. In such games, sometimes called Gale-Stewart games, players play sequentially, one after the other, back and forth forever. Early models considered a perfect information monitoring structure. [Gale and Stewart (1953)] began this line by showing that if the eventual winning set W - the set Player 1 strives to have the infinite play of the game belong to, while Player 2 strives that the play not belong - is either open or closed, then one player can force a win. [Wolfe (1956)] (also [Blackwell (1969)]) extended this result to the case where W is G_δ (or, symmetrically, F_σ). Eventually, [Martin (1975)] demonstrated that if W is a Borel set, then the game is determined.

A natural generalization is the case in which the monitoring structure is not perfect. In fact, [Blackwell (1969)] already incorporated this result by allowing

*This work was carried out primarily when the authors were both Ph.D. students at the Center for the Study of Rationality, Hebrew University, Jerusalem.

[†]Faculty of Industrial Engineering and Management, Technion- Israel Institute of Technology, iarieli@tx.technion.ac.il

[‡]Research supported in part by Israel Science Foundation grants 1596/10.

[§]Nuffield College and Department of Economics, University of Oxford, john.calculus@gmail.com

the actions to be simultaneous - equivalent to a one-round information delay on the part of Player 2. In such cases, one cannot hope that one player or the other can force a win. Nonetheless, we can hope that the game - in which a win for a player is interpreted as the gain of a unit from another player - possesses a value. This weakened concept has earned, counter-intuitively or not, the title of determinacy as well, and has enjoyed generalization to more general payoff functions as well; see [Martin (1998)]. An application of such games to manipulability of inspections can be found in [Shmaya (2008)]; for pure mathematical applications, see, e.g., [Kechris(1995)].

Building on these results and motivations, [Shmaya (2011)] made a significant step forward when considering very general delays in information. [Shmaya (2011)] required only that each player have, at each stage, a partition over his opponent's possible histories of play, and learns to refine this partitions over time to the extent that he can eventually differentiate between any two different plays. This condition is termed by [Shmaya (2011)] as *Eventual Perfect Monitoring* (henceforth, EPM), and it is shown that it is sufficient to guarantee determinacy.

Our work focuses on a generalization of the EPM setup to games in which, as well, information is learned at a delay, but not by deterministic methods such as the partitions used in EPM, but rather by stochastic signalling. The first pressing question is, then, what should be the natural generalization of EPM? The key, it seems, is to observe the transition kernel (for each player) from infinite sequences of plays of the game to infinite sequences of his own signals. The natural generalization of partitions being disjoint to the non-deterministic case is the condition of measures being mutually singular. As such, two natural conditions on the monitoring structure have arisen:

One condition, to which we give the title of *Stochastic Eventual Perfect Monitoring* (henceforth, SEPM), requires that any two profiles of strategies which induce mutually orthogonal distributions on the space of plays of the game should induce, for each player, mutually orthogonal distributions on the space of sequences of that player. A weaker condition, however, which we call *Weak Stochastic Eventual Perfect Monitoring* (henceforth, WSEPM), requires only that for any two different infinite histories of play, the induced measures on the space of signals of either player should be mutually orthogonal. These conditions coincide in the case of the deterministic signalling of [Shmaya (2011)].

The purpose of this paper is then two-fold: Our main result is that SEPM is sufficient to imply determinacy. Our technique generalizes the techniques of [Shmaya (2011)], and like that work includes a reduction to a stochastic game with Borel winning set. In this framework, we also generalize [Shmaya (2011)] by allowing for stochastic states of Nature to be chosen at each period. Our other main result is to show, by example - and via development of some techniques that we hope are of independent interest - is to show that WSEPM is

not sufficient to guarantee determinacy. Properties and equivalent reformulations of the SEPM condition, as well as other applications, can be found in [Arieli and Levy (In Preperation.)].

Epistemically, we note that our determinacy result is fundamentally different than that of [Shmaya (2011)]. The EPM condition implies that every action played eventually becomes common knowledge among the players. Our condition, it turns out, in addition to only implying *p-belief* (as coined in [Monderer and Samet (1989)]) and not knowledge, only implies mutual belief and not common belief. (See Proposition 3.1.) That is, for every order k (and each $p < 1$), any action played eventually becomes mutual p -belief up to k levels among the players: Each player p -believes it, each player p -believes that each p -believes it, and so on up to k levels, but it need never become common p -belief, i.e., the chain of mutual p -belief may never, at any finite time, be continued *ad infinitum*. This difference is discussed further in [Arieli and Levy (In Preperation.)]. Our determinacy result shows that the eventual common learning - more precisely, common p -belief - is not what is required for determinacy - but rather only mutual learning in the appropriate sense.

2 Model and Results

2.1 Preliminary Notation

For a Borel space S , let $\Delta(S)$ denote the space of regular Borel probability measure on S , endowed with the topology of narrow convergence. For a finite set A , we denote by $|A|$ or $\#A$ the cardinality of A . For two sets A, B , $A\Delta B$ denotes the symmetric difference. For $j \in \mathbb{N}$, let $[j] \in \{1, 2\}$ be such that $j = [j] \bmod 2$.

2.2 Definition

Definition 2.1. *A two-player zero-sum sequential game with signals is given by a quadrupole $\Gamma(W) = ((A_j)_{j \in \mathbb{N}}, q, \Theta, (\eta_j)_{j \in \mathbb{N}}, W)$ where:*

- A_j is the finite action space used at stage j , respectively.
- W is a subset of $H_\infty := \prod_{j \in \mathbb{N}} A_j$.
- Θ is a standard Borel space of signals.¹
- For each $n \in \mathbb{N}$, $\eta_n : \prod_{j < n} \Theta^2 \times A_j \rightarrow \Delta(\Theta^2)$ is the transition kernel² of

¹One could also allow for time-dependent signals - as we will later - by taking Θ to be a disjoint union of all signalling spaces.

²Given Borel spaces X, Y , recall that a transition kernel is a Borel-measurable mapping from X to $\Delta(Y)$, where the latter is endowed with the measurable structure associated with the topology of narrow convergence; equivalently, it is a function $\mu(\cdot | \cdot)$ such that for each $x \in X$, $\mu(\cdot | x) \in \Delta(Y)$, and for each Borel $B \subseteq Y$, $\mu(B | \cdot)$ is measurable.

signals.

Denote $H_n = \prod_n A_j$, $H_* = \cup_n H_n$. We will treat the transition kernel of signals as a single function $\eta : H_* \rightarrow \Delta(\Theta^2)$.

The dynamics of the game are as follows: Player 1 (resp. 2) plays at odd (resp. even) stages. Before stage n , a signal is revealed to each player³ - denote the signal to Player j before stage n by θ_n^j ; the pair (θ_n^1, θ_n^2) is chosen by Nature according to the distribution $\eta(h)$; we will denote the marginal on each coordinate by η^j for $j = 1, 2$. Following this, Player $[n]$ chooses an action in A_n .

Player 1 wins if the resulting infinite history $h \in H_\infty$ is in W (and receives a payoff of one unit from player 2); Player 2 wins (and receives one unit from Player 1) if $h \notin W$.

2.3 The Signalling Transition Kernels

We define the mappings η , on H_* - specifically, each element of H_n defines a distribution on $(\Theta^2)^n$ - by

$$\eta(a_1, \dots, a_{n-1})[\theta_1^1, \theta_1^2, \dots, \theta_n^1, \theta_n^2] = \prod_{i \leq n} \eta(\theta_1^1, \theta_1^2, a_1, \dots, \theta_{i-1}^1, \theta_{i-1}^2, a_{i-1})[\theta_i^1, \theta_i^2]$$

That is, given a finite history of play, η gives the distribution induced on the signals via Bayesian inference.

Let $\eta^1(h), \eta^2(h)$ be the marginals on the signals for Player 1, 2, respectively. We shall make the following assumption throughout:

Assumption 2.2. (Perfect Recall) Let $j \in \{1, 2\}$, $n \in \mathbb{N}$, and let $\pi^j : H_n \rightarrow \prod_{[s]=j, s \leq n} A_s$ be the projection of j 's actions. Then, for any two $\rho, \lambda \in \Delta(H_n)$ which satisfy $\pi_*^j(\rho) \perp \pi_*^j(\lambda)$, we have $\eta^j(\rho) \perp \eta^j(\lambda)$, where $\nu^j(\rho)(\cdot) = \int_{H_n} \nu^j(h)(\cdot) d\rho(h)$ and similarly for λ .

Hence (since there are only finitely many actions), each player can almost surely deduce his own previous actions from the signals he has received, and hence when defining strategies below, we may assume each player makes decisions depending only his signals.

As such, we have two transition kernels η_∞^i , $i = 1, 2$, from H_∞ to Θ^∞ ; each infinite history $h \in H_\infty$ induces probability distributions $\eta^i(h)$ on Θ^∞ for $j = 1, 2$ - that is, probability distributions on the sequence of each players' signals defined for cylindrical sets by

$$\eta_\infty^j(h)(\{\bar{\theta} \in \Theta^\infty \mid \bar{\theta}_n = p\}) = \eta^j(h_n)(\bar{\theta}_n = p), \quad \forall p \in \Theta^n$$

³Since Player 1 plays only at odd states and Player 2 only at even states, its actually not necessary that each player receive a signal at each stage; it would have sufficed had they received before they play. It is simpler for our notation, however, to assume they each receive a signal at each stage.

2.4 Strategies

A behavioral strategy for Player 1 is a sequence of functions $\sigma = \{\sigma_n\}_{n=1,3,5,\dots}$, where σ_n assigns to each sequence $(\theta_1^1, \dots, \theta_{n-1}^1, \theta_n^1) \in \Theta^n$ a mixed action in $\Delta(A_n)$, and similarly for Player 2 at the even stages. By the assumption we made above, players are able to choose from these families of behavioral strategies have perfect recall.

Each pair of behavioral strategies (σ, τ) induces a probability distribution $P_{\sigma, \tau}$ on $\tilde{H}_\infty := H_\infty \times \Theta^\infty \times \Theta^\infty$, the space, of infinite plays of the game including the sequences of signals the players receive.

2.5 Deterministic Signalling and Determinacy

The concept defined by Shmaya (2011), [Shmaya (2011)] can be found in our context in the following manner:

Definition 2.3. *The signalling structure of a game, with notation as Definition 2.1, is said to be deterministic if $\eta_n(h)$ is a Dirac measure for each $n \in \mathbb{N}$, $h \in H_n$. In this case, let $\eta_\infty^1, \eta_\infty^2$ be the functions⁴ defined on H_∞ , which assign to each infinite history the resulting infinite sequence of signals for Player 1, 2, respectively. The game is said to have eventual perfect monitoring (EPM) if $\eta_\infty^1, \eta_\infty^2$ are injective.*

$\Gamma(W)$ is said to be *determined* if it possesses a value, that is, if

$$\sup_{\sigma \in \Sigma^1} \inf_{\tau \in \Sigma^2} P_{\sigma, \tau}(W) = \inf_{\tau \in \Sigma^2} \sup_{\sigma \in \Sigma^1} P_{\sigma, \tau}(W)$$

where Σ^j is the space of behavioral strategies for Player j . The result of Shmaya(2011), [Shmaya (2011)], is:

Theorem 2.4. *If a game has a Borel winning set, and deterministic signalling which satisfies the EPM condition, then the game is determined.*

2.6 SEPM & The Main Result for Sequential Games

Definition 2.5. *Let $\pi^H : \tilde{H}_\infty \rightarrow H_t$, (resp. $\pi^1, \pi^2 : \tilde{H}_\infty \rightarrow \Theta^\infty$) be projections on the space of plays (resp. sequences of signals for Players 1, 2). The game is said to possess Stochastic Eventual Perfect Monitoring (henceforth, SEPM) if for any pair of profile strategies (σ, τ) and (σ', τ') such that $\pi_*^H(P_{\sigma, \tau}) \perp \pi_*^H(P_{\sigma', \tau'})$, it holds that $\pi_*^j(P_{\sigma, \tau}) \perp \pi_*^j(P_{\sigma', \tau'})$ for $j = 1, 2$.*

The main result for sequential games is:

⁴For any Borel space Y , the set of Dirac measures is a closed subspace of $\Delta(Y)$ and the mapping $y \rightarrow \delta_y$ is a homeomorphism onto it; e.g., [Bertsekas and Shreve (1996), Cor. 7.21.1]. Hence, η^1, η^2 are Borel.

Theorem 2.6. *If a sequential game has a Borel winning set, and signalling which satisfies the SEPM condition, then the game is determined.*

The other result for sequential games of this paper is to show, by examples, that the following condition does not guarantee determinacy:

Definition 2.7. *The game is said to possess Weak Stochastic Eventual Perfect Monitoring (henceforth, WSEPM) if for any two $h, h' \in H_\infty$, $\eta_\infty^j(h) \perp \eta_\infty^j(h')$ for each $j \in \{1, 2\}$.*

Remark 2.8. These conditions coincide in the case of deterministic signalling.

Section 3 is dedicated to the proof of Theorem 2.12. We mimic the technique of [Shmaya (2011)], where a reduction was made to a stochastic game with standard signalling. A game satisfying WSEPM but which is not determined is given in Section 4.

2.7 Generalization: Stochastic Games

Definition 2.9. *A two-player zero-sum (sequential⁵) stochastic game with signals is given by a sextuple $\Gamma(W) = ((S_j)_{j \in \mathbb{N}}, (A_j)_{j \in \mathbb{N}}, q, \Theta, (\eta_j)_{j \in \mathbb{N}}, W)$ where:*

- S_j, A_j are the finite state and action spaces used at stage j , respectively.
- W is a subset of $H_\infty := \prod_{j \in \mathbb{N}} (S_j \times A_j)$.
- Θ is a standard Borel space of signals.
- For each $n \in \mathbb{N}$, $\eta_n : \tilde{H}_n \rightarrow \Delta(\Theta^2)$, where $\tilde{H}_n = \prod_{j < n} (S_j \times \Theta^2 \times A_j) \times S_n$, is the transition kernel of signals.
- For each $n \in \mathbb{N}$, $q_n : H_{n-1} \rightarrow \Delta(S_n)$ is the transition kernel of states, where $H_{n-1} = \prod_{j < n} S_j \times A_j$ ($H_0 = \{\emptyset\}$).

We will also denote $\tilde{H}_* = \cup_n \tilde{H}_n$, $H_n^\diamond = H_{n-1} \times S_n$, $H_*^\diamond = \cup_n H_n^\diamond$, $H_* = \cup_n H_n$. We will treat the transition kernel of signals as a single function $\eta : H_*^\diamond \rightarrow \Delta(\Theta^2)$, and similarly we will view the state transition kernel as a single function $q : H_* \rightarrow \Delta(\cup_n S_n)$ with $\text{supp}(q(h)) \subseteq S_n$ for $h \in H_n$.

The dynamics of the game are as follows: The initial state z_1 is chosen by Nature according to the distribution $q(\emptyset)$. Suppose at some stage n , we are in state $z_n \in S_n$, the history of the game up to that point being $h = (z_1, \theta_1^1, \theta_1^2, a_1, \dots, z_{n-1}, \theta_{n-1}^1, \theta_{n-1}^2, a_{n-1}, z_n) \in \tilde{H}_n$. A signal is revealed to each player - denote the signal to Player j by θ_n^j ; the pair (θ_n^1, θ_n^2) is chosen by Nature according to the distribution $\eta(h)$; we will denote the marginal on each coordinate by η^j for $j = 1, 2$. Following this, Player $[n]$ chooses an action in A_n . The next state z_{n+1} is chosen according to the distribution $q(z_1, a_1, \dots, z_n, a_n)$, and

⁵To differentiate from the standard stochastic game models in which players play simultaneously.

the process repeats.

Player 1 wins if the resulting infinite history $h \in H_\infty$ is in W (and receives a payoff of one unit from player 2); Player 2 wins if $h \notin W$, and receives one unit from Player 1).

The transition kernels η^1, η^2 (resp. $\eta_\infty^1, \eta_\infty^2$) are now defined on H_*^\diamond (resp. H_∞), and the notions of perfect recall (with H_t^\diamond replacing H_t in Definition 2.2), strategies (note that strategies still depend only on signals received, the players do not directly observe the states), and determinacy are defined in the same way as they were for sequential games.

Definition 2.10. $\hat{q} = (\hat{q}_1, \hat{q}_2, \dots)$ with $\hat{q}_n : H_{n-1} \rightarrow \Delta(S_n)$ is said to be a belief on Nature if for each $n \in \mathbb{N}$ and each $h \in H_n^\diamond$, there is $S_{n,h} \subseteq S_n$ such that $q(h)(S_{n,h}) > 0$ and $\hat{q}(h)(\cdot) = q(h)(\cdot | S_{n,h})$.

Just like above, given a belief on Nature \hat{q} and a strategy profile σ, τ , $P_{\hat{q}, \sigma, \tau}$ is the induced probability measure on \tilde{H}_∞ when the original transition kernel q is replaced with \hat{q} .

Definition 2.11. Let $\pi^H : \tilde{H}_\infty \rightarrow \prod_{n \in \mathbb{N}} S_t \times A_t$, (resp. $\pi^1, \pi^2 : \tilde{H}_\infty \rightarrow \Theta^\infty$) be projections on the space of plays (resp. sequences of signals for Players 1, 2). The game is said to possess Stochastic Eventual Perfect Monitoring (henceforth, SEPM) if for any pair of profile strategies (σ, τ) and (σ', τ') and any pair of beliefs on Nature \hat{q}, \hat{q}' such that $\pi_*^H(P_{\hat{q}, \sigma, \tau}) \perp \pi_*^H(P_{\hat{q}', \sigma', \tau'})$, it holds that $\pi_*^j(P_{\hat{q}, \sigma, \tau}) \perp \pi_*^j(P_{\hat{q}', \sigma', \tau'})$ for $j = 1, 2$.

The main theorem of this paper is:

Theorem 2.12. *Theorem 2.6 holds for stochastic games with SEPM as well.*

Remark 2.13. [Shmaya (2011)] works under the assumption that there are no states; i.e., S_n is trivial for all $n \in \mathbb{N}$. However, the proof there can be modified easily to incorporate states. This also follows from our main Theorem 2.12 below.

Remark 2.14. If one wants to define WSEMP for stochastic games in such a way that SEMP implies WSEMP, one should phrase Definition 2.7 as holding for those $h, h' \in H_\infty$ such that for every $n \in \mathbb{N}$, there exists strategy profiles $\sigma, \tau, \sigma', \tau'$ for which the projections $h|_n, h'|_n$ satisfy $P_{\sigma, \tau}(h|_n) > 0, P_{\sigma', \tau'}(h'|_n) > 0$ (i.e., those histories which are assigned positive probabilities for some strategy profile.) If one would generalize the definition without this modification - i.e., requiring it to hold for all $h, h' \in H_\infty$ - it is not clear if SEMP implies WSEMP.

Remark 2.15. We remark that the addition of the beliefs of Nature to the definition of SEPM helps facilitate the learning: Otherwise, we could have a situation where each player has trivial action spaces (only one option) and yet non-trivial state spaces. In that case, a notion of SEPM not allowing for beliefs on Nature would hold trivially - since there is only one action profile - and yet

players would not necessarily learn anything about the states. However, in such an example, determinacy would follow trivially so it is not clear to what extent this stringency of the SEPM condition is needed for Theorem 2.12 to hold.

Remark 2.16. Like in [Shmaya (2011)], it is still unknown if determinacy continues to hold if the payoff is given not by a Borel winning set but by a more general bounded Borel payoff function.

3 Proof of Theorem 2.12

Throughout this section, to keep the notation less cumbersome, we will use the same $P_{\sigma,\tau}$ to denote the marginals on either coordinate. In the course of probabilistic calculations, we will also identify freely between subsets $B \subseteq H_\infty$ and $B \times (\Theta^\infty)^2 \subseteq \tilde{H}_\infty$, and similarly for subsets of $(\Theta^\infty)^2$, or subsets Θ^∞ when it is clear which player's signals we are referring to.

3.1 Preliminaries

First, we make a reduction to the case in which the signal spaces are finite (but stage-dependent). This reduction is relegated to Appendix B. Henceforth, assume that at each n , the set of signals for each player is a finite set Θ_n .

Proposition 3.1. *Let $h \in H_*^\circ$, and let (σ, τ) be a strategy profile. If SEPM holds, then for each $\varepsilon > 0$, there is $N \in \mathbb{N}$ such that for $j = 1, 2$ and all $n \geq N$ with $P_{\sigma,\tau}(h) > 0$,*

$$P_{\sigma,\tau}(P_{\sigma,\tau}(h|\theta_1^j, \dots, \theta_n^j) > 1 - \varepsilon \mid h) > 1 - \varepsilon \quad (3.1)$$

When there are no states, this condition was shown to be equivalent to SEPM, [Arieli and Levy (In Preparation.)] and was termed Eventual Learning, but we only require one direction, and we provide in Appendix C here a proof more direct than that in [Arieli and Levy (In Preparation.)]; in that work, other equivalent conditions are also discussed.

The following is essentially Corollary 4.3 of [Shmaya (2011)] and we do not repeat the proof.

Lemma 3.2. *Let $\varepsilon > 0$, and assume the signal spaces $\Theta_1, \Theta_2, \dots$ are finite.⁶ There exists a sequence of finite sets $(\Delta_j^\varepsilon)_{j \in \mathbb{N}}$, with $\Delta_j^\varepsilon \subseteq \Delta(A_j)$, such that for any pair of behavioral strategies σ, τ , there is a pair σ', τ' which choose, at each stage j , mixed actions in Δ_j^ε , and such that $\|P_{\sigma, \hat{\tau}} - P_{\sigma', \hat{\tau}}\| < \varepsilon$ and $\|P_{\hat{\sigma}, \tau} - P_{\hat{\sigma}, \tau'}\| < \varepsilon$ for any strategies $\hat{\sigma}, \hat{\tau}$, the norm being the total-variation norm.*

⁶This lemma remains correct even if the state spaces are general; but the proof is straightforward the case of finite signal spaces.

For each $\varepsilon > 0$, fix $(\Delta_j^\varepsilon)_{j \in \mathbb{N}}$, let Σ, Υ denote the set of behavioural strategies for Player 1, 2, respectively, and let $\Sigma^\varepsilon, \Upsilon^\varepsilon$ be those strategies taking mixed actions in (Δ_j^ε) .

Lemma 3.3. *Assume SEPM. For each $h \in H_*^\diamond$, $\varepsilon > 0$, and $0 < \varrho < 1$, there is $N = N(h, \varepsilon, \varrho)$, such that for all strategy profiles (σ, τ) in $\Sigma^\varepsilon \times \Upsilon^\varepsilon$ with $P_{\sigma, \tau}(h) \geq \varrho$, each $j = 1, 2$, and all $n \geq N$, (3.1) holds.*

Proof. The space of Σ of behavioral strategies for Player 1 can be identified as $\prod_{[n]=j} (\Delta(A_n))^{\Theta^n}$, and similarly the space of behavioral strategies Υ of Player 2. Endow these spaces with the product topology. The strategy spaces $\Sigma^\varepsilon, \Upsilon^\varepsilon$ are identified as compact subsets of these spaces. Fix $0 < \varrho < 1$, $\varepsilon > 0$, $h \in H_*^\diamond$. The collection Ω of strategy pairs $(\sigma, \tau) \in \Sigma^\varepsilon \times \Upsilon^\varepsilon$ which satisfy $P_{\sigma, \tau}(h) \geq \varrho$ is compact, since h is finitary.⁷

For each n , let Ξ_n denote the collection of strategy pairs $(\sigma, \tau) \in \Omega$ which satisfy (3.1); clearly, (3.1) is also determined by the first n coordinates so Ξ_n is both open and closed, and by Proposition 3.1, each $(\sigma, \tau) \in \Omega$ belongs to all but finitely many Ξ_n . Hence, an appeal to Lemma 5.3 of the Appendix A completes the proof. \square

Remark 3.4. We introduce several notations, based on Lemma 3.3. Fix $\varepsilon > 0$ and $0 < \varrho < 1$. For each $k \in \mathbb{N}$, let

$$N(k, \varepsilon, \varrho) = \max_{h \in H_k^\diamond} N(h, \varepsilon, \varrho)$$

We will also assume that for each strategy pair σ, τ , each player j , each $h \in H_*^\diamond$, $N(h, \sigma, \tau, j, \varepsilon, \varrho)$ is the smallest possible selection such that (3.1) holds for all $n \geq N(h, \sigma, \tau, j, \varepsilon, \varrho)$, and denote like above,

$$N(k, \sigma, \tau, j, \varepsilon, \varrho) = \max_{h \in H_k^\diamond} N(h, \sigma, \tau, j, \varepsilon, \varrho)$$

In particular, if $j \in \{1, 2\}$, $k \in \mathbb{N}$, $(\sigma, \tau), (\sigma', \tau')$ are strategy profiles in $\Sigma^\varepsilon, \Upsilon^\varepsilon$ which agree through the first $N = N(k, \sigma, \tau, j, \varepsilon, \varrho)$ stages, then $N(k, \sigma', \tau', j, \varepsilon, \varrho) = N$. As such, if β_1, \dots, β_n are such that $\beta_k : \Theta^k \rightarrow \Delta(A_j)$ for each $1 \leq k \leq N$ - i.e., β_1, \dots, β_N dictate the first N stages of a strategy profile - and $j \in \{1, 2\}$, we can define $N(k, \beta_1, \dots, \beta_n, j, \varepsilon, \varrho) = N(k, \sigma, \tau, j, \varepsilon, \varrho)$ for any (equivalently, some) strategy profile (σ, τ) which dictates the same rules as β_1, \dots, β_n through N stages, as long as $N(k, \sigma, \tau, j, \varepsilon, \varrho) \leq N$; otherwise, $N(k, \beta_1, \dots, \beta_n, j, \varepsilon, \varrho)$ is undefined.

Proposition 3.5. *If W is compact, then $\Gamma(W)$ is determined.*

Note that we do not make any requirements of the signalling structure for this result. The proof is essentially the same as Lemma 3.1 of Shmaya (2011), [Shmaya (2011)], and we sketch it for convenience:

⁷A set is finitary if it is determined by finitely many coordinates.

Proof. The mapping from pure behavioral strategies - which are compact spaces in the product topology - to expected payoff⁸ is upper semi-continuous in this game. Hence, by Fan's minimax theorem, a value exists in mixed strategies.⁹ Since the game has perfect recall, behavioral strategies are equivalent to mixed strategies, by Kuhn's theorem. \square

3.2 The Auxiliary Game Λ

We henceforth assume that $\Gamma(W)$ satisfies the SEPM assumption and has countable state space.

Fix $\varepsilon > 0$ with $\varepsilon < 1$, let (Δ_j) be the finite signal spaces, and let $N(\cdot, \cdot, \cdot)$ be the function in Remark 3.4. Denote $\bar{\Theta}_n = \prod_{j \leq n} \Theta_n$. We define an auxiliary stochastic game Λ of perfect information:
Let $B_n = \{b : \bar{\Theta}_n \rightarrow \Delta_n^\varepsilon\}$ be the set of actions at stage $n = 1, 2, \dots$ in Λ . Denote $\bar{B}_n = \prod_{j \leq n} B_n$. Define $\Xi : \mathbb{N} \rightarrow \mathbb{N}$ by $\Xi(k) = N(k, \frac{\varepsilon}{2^{k+1}}, \frac{\varepsilon}{2^{k+1} \cdot |H_k^\circ|})$, and let

$$K_n = \{k \in \mathbb{N} \mid \Xi(k) = n\}, \quad T'_n = \prod_{k \in K_n} A_{k-1} \times S_k, \quad T_n = T'_n \cup \{*\}$$

where we take A_0 to be trivial. T_n is the set of states of stage n in Λ . K_n can be described as those stages 'approximately' learned by stage n and not necessarily earlier; T'_n is what is actually learned.

For each $n \in \mathbb{N}$, define $M_n = \max\{k \mid \Xi(k) \leq n\}$ - that is, the stages necessarily learned by stage n - and define for $j = 1, 2$, $\tilde{f}_n^j : \bar{B}_{n-1} \times (\bar{\Theta}_n)^2 \rightarrow H_{M_n}^\circ$, where $\tilde{f}_n^j(\bar{\beta}_{n-1}, \bar{\theta}_n)$, for $\bar{\beta}_{n-1} \in \bar{B}_{n-1}$ and $\bar{\theta}_n \in \bar{\Theta}_n$, is an $h \in H_{M_n}^\circ$ such that:

$$P_{\bar{\beta}_{n-1}}(h \mid \bar{\theta}_{N^*}^j) > 1 - \frac{\varepsilon}{2^n}, \quad j = 1, 2 \quad (3.2)$$

where $N^* = N(M_n, \bar{\beta}_{n-1}, j, \frac{\varepsilon}{2^{k+1}}, \frac{\varepsilon}{2^{k+1} \cdot |H_k^\circ|})$, if such h exists, in which case it is unique.¹⁰ Otherwise, denote $\tilde{f}_n^j(\bar{\beta}_{n-1}, \bar{\theta}_n^j) = *$. Then, define

$$\tilde{f}_n(\bar{\beta}, \bar{\theta}_n) = \begin{cases} \tilde{f}_n^1(\bar{\beta}, \bar{\theta}_n^1) & \text{if } \tilde{f}_n^1(\bar{\beta}, \bar{\theta}_n^1) = \tilde{f}_n^2(\bar{\beta}, \bar{\theta}_n^2) \neq * \\ * & \text{otherwise} \end{cases}$$

Lemma 3.6. *For any pair of behavioral strategies σ, τ in $\Gamma^\varepsilon, \Upsilon^\varepsilon$, and each $j \in \{1, 2\}$,*

$$P_{\sigma, \tau}(\exists n \in \mathbb{N}, j \in \{1, 2\}, \tilde{f}_n(\bar{\beta}_n, \bar{\theta}_n^j) \neq (z_1, a_1, \dots, z_{M_n})) < 2\varepsilon$$

where $\beta_j = \sigma_j$ for odd j , $\beta_j = \tau_j$ for even j .

⁸This is the only difference from the corresponding lemma in [Shmaya (2011)] - pure strategy profiles here do not yield deterministic payoffs.

⁹That is, mixtures of pure behavioral strategies.

¹⁰Since $1 - \frac{\varepsilon}{2^n} > \frac{1}{2}$

Proof. Denote $Z_n = \{h \in H_{M_n}^\diamond \mid P_{\sigma,\tau}(h) \geq \frac{\varepsilon}{2^{n+1} \cdot |H_{M_n}^\diamond|}\}$. For each $j \in \{1, 2\}$,

$$\begin{aligned}
& P_{\sigma,\tau}(\exists n \in \mathbb{N}, \tilde{f}_n^j(\bar{\beta}_n, \bar{\theta}_n^j) \neq (z_1, a_1, \dots, z_{M_n})) \\
& \leq \sum_{n \in \mathbb{N}} P_{\sigma,\tau}(\tilde{f}_n^j(\bar{\beta}_n, \bar{\theta}_n^j) \neq (z_1, a_1, \dots, z_{M_n})) \\
& = \sum_{n \in \mathbb{N}} \sum_{h \in H_{M_n}^\diamond} P_{\sigma,\tau}(\tilde{f}_n^j(\bar{\beta}_n, \bar{\theta}_n^j) \neq h \wedge h = (z_1, a_1, \dots, z_{M_n})) \\
& = \sum_{n \in \mathbb{N}} \sum_{h \in Z_n} P_{\sigma,\tau}(h) \cdot P_{\sigma,\tau}(\tilde{f}_n^j(\bar{\beta}_n, \bar{\theta}_n^j) \neq h|h) \\
& + \sum_{n \in \mathbb{N}} \sum_{h \in H_{M_n}^\diamond \setminus Z_n} P_{\sigma,\tau}(\tilde{f}_n^j(\bar{\beta}_n, \bar{\theta}_n^j) \neq h \wedge h = (z_1, a_1, \dots, z_{M_n})) \\
& \leq \sum_{n \in \mathbb{N}} \sum_{h \in Z_n} P_{\sigma,\tau}(h) \frac{\varepsilon}{2^{n+1}} + \sum_{n \in \mathbb{N}} \sum_{h \in H_{M_n}^\diamond \setminus Z_n} P_{\sigma,\tau}(h = (z_1, a_1, \dots, z_{M_n})) \\
& \leq \frac{\varepsilon}{2} + \sum_{n \in \mathbb{N}} |H_{M_n}^\diamond| \cdot \frac{\varepsilon}{2^{n+1} \cdot |H_{M_n}^\diamond|} \leq \varepsilon
\end{aligned}$$

where we have used the definition of $\Xi(\cdot)$. \square

Let $\pi_n : H_n^\diamond \rightarrow T_n'$ be defined by projection, and let $F : T_1 \times T_2 \times \dots \rightarrow H_\infty \cup \{*\}$ be defined by

$$\begin{aligned}
F((\pi_n(u|_n))_{n \in \mathbb{N}}) &= u, \text{ if } u \in H_\infty \\
F((t_n)_{n \in \mathbb{N}}) &= *, \text{ if } \exists n \in \mathbb{N}, t_n = *
\end{aligned}$$

By Lemma (3.3), this is well-defined. By mildly abusive notation, $F : T_1 \times \dots \times T_n \rightarrow H_{M_n}^\diamond \cup \{*\}$, which is the projection of F defined above onto $H_{M_n}^\diamond$, and $*$ projects to $*$; this is well-defined, as these first M_n coordinates in the output of F dependent only on the first n coordinates of $\prod_j T_j$. Let $f_k = \pi_k \circ \tilde{f}_k$, where $\pi_j(*) = *$.

In Λ , Player 1 plays at odd stages, Player 2 plays at even stages, with perfect monitoring. Given $p = (t_1, b_1, \dots, t_n, b_n) \in T_1 \times B_1 \times \dots \times T_n \times B_n$, we need to define the distribution induced on the next state - i.e., the distribution on T_n - that Nature employs. We define the transition function in Λ by

$$\tilde{q}(t_1, b_1, \dots, t_{n-1}, b_{n-1})[*] = 1, \text{ if } \exists j < n, t_j = *$$

and otherwise,

$$\begin{aligned}
& \tilde{q}(t_1, b_1, \dots, t_{n-1}, b_{n-1})[t] = \\
& P(f_n(b_1, \dots, b_n, \hat{\theta}_1, \dots, \hat{\theta}_n) = t \mid f_k(b_1, \dots, b_k, \hat{\theta}_1, \dots, \hat{\theta}_k) = t_k, \forall k < n) \quad (3.3)
\end{aligned}$$

where $(\hat{\theta}_j)_{j=1}^n$ is a sequence of random variables, with $\hat{\theta}_j \in \Theta_j^2$, such that there are sequences $(\hat{z}_j)_{j=1}^n$, $(\hat{a}_j)_{j=1}^{n-1}$ of random variables, $\hat{z}_j \in S_j$ and $\hat{a}_j \in A_j$ satisfying

$$P(\hat{z}_k \mid \hat{z}_1, \hat{a}_1, \dots, \hat{z}_{k-1}, \hat{a}_{k-1}) = q(\hat{z}_1, \hat{a}_1, \dots, \hat{z}_{k-1}, \hat{a}_{k-1})[\hat{z}_k] \quad (3.4)$$

$$P(\hat{a}_k \mid \hat{\theta}_1, \dots, \hat{\theta}_k) = b_k(\hat{\theta}_1^{[k]}, \dots, \hat{\theta}_k^{[k]})[\hat{a}_k] \quad (3.5)$$

and

$$\begin{aligned} P(\hat{\theta}_k \mid \hat{\theta}_1, \dots, \hat{\theta}_{k-1}, \hat{z}_1, \hat{a}_1, \dots, \hat{z}_{k-1}, \hat{a}_{k-1}, \hat{z}_k) \\ = \eta(\hat{\theta}_1, \dots, \hat{\theta}_{k-1}, \hat{z}_1, \hat{a}_1, \dots, \hat{z}_{k-1}, \hat{a}_{k-1}, \hat{z}_k)[\hat{\theta}_k] \end{aligned} \quad (3.6)$$

What we have described above is the dynamics of the game Λ . For a $W \subseteq H_\infty$, Player 1 wins in $\Lambda(W)$ if $F(s_1, s_2, \dots) \in W$ (and receives a payoff of 1 unit from Player 2), and loses otherwise (and receives a payoff of -1). The following is essentially Lemma 4.4 in Shmaya (2011), [Shmaya (2011)], and for convenience we again recall the proof.

Lemma 3.7. *The game $\Lambda(W)$ is determined when W is Borel, and $\text{val}(\Lambda(W_0)) \geq \text{val}(\Lambda(W)) - \varepsilon$ for some compact set $W_0 \subseteq W$.*

Proof. For $V \subseteq \prod_{n \in \mathbb{N}} T_n \times B_n$, let $\Lambda_0(V)$ denote the game with dynamics as in Λ such that Player 1 wins if $(s_1, b_1, s_2, b_2, \dots) \in V$.¹¹ Define $G : (\prod_{j \in \mathbb{N}} (T_j \setminus \{*\}) \times B_j) \rightarrow H_\infty$ by

$$G(s_1, b_1, s_2, b_2, \dots) = F(s_1, s_2, \dots)$$

G is continuous, and $\Lambda(W) = \Lambda_0(G^{-1}(W))$ is a stochastic game with winning set $G^{-1}(W)$, since

$$F(s_1, s_2, \dots) \in W \iff G(s_1, b_1, s_2, \dots) \in W \iff (s_1, b_1, s_2, \dots) \in G^{-1}(W)$$

Therefore, $\Lambda(W)$ has a value, see [Martin (1998)]. Furthermore, there is a compact subset $C \subseteq G^{-1}(W)$ such that $\text{val}(\Lambda_0(C)) > \text{val}(\Lambda_0(G^{-1}(W))) - \varepsilon$ ([Maitra et al (1992)]; see also [Maitra and Sudderth (1996), Ch. 6]). But $\text{val}(\Lambda_0(C)) \leq \text{val}(\Lambda_0(G^{-1}(G(C)))) = \text{val}(\Lambda(G(C)))$. We can take $W_0 = G(C)$. W_0 is compact, satisfies the require inequality as

$$\text{val}(\Lambda(W_0)) = \text{val}(\Lambda(G(C))) \geq \text{val}(\Lambda_0(C)) \geq \text{val}(\Lambda_0(G^{-1}(W))) - \varepsilon = \text{val}(\Lambda(W)) - \varepsilon$$

and $W_0 = G(C) \subseteq G(G^{-1}(W)) = W$. \square

Lemma 3.8. *For every Borel set W of H_∞ ,*

$$\text{val}(\Lambda(W)) - 4\varepsilon \leq \overline{\text{val}}(\Gamma(W))$$

¹¹This is fundamentally different than $\Lambda(W)$, which is defined via the function F .

Lemma 3.8 is the heart of the proof. If we take this Lemma as a given, then it is easy to complete the proof of Theorem 2.12: Let $\Gamma(W)$ with W Borel be a game which satisfies the SEPM assumption. Let $\varepsilon > 0$; we have already discussed that we may assume that only finitely many signals are possible at each stage. Let W_0 be a compact subset of W as in Lemma 3.7. Then

$$\underline{val}\Gamma(W) \geq \underline{val}\Gamma(W_0) = \overline{val}\Gamma(W_0) \geq \underline{val}\Lambda(W_0) - 4\varepsilon > \underline{val}\Lambda(W) - 5\varepsilon$$

where the first inequality follows from $W_0 \subseteq W$, the equality from Proposition 3.5, the second inequality from Lemma 3.8, and the third inequality from the choice of W_0 .

In a symmetric fashion, $\overline{val}\Gamma(W) < \underline{val}\Lambda(W) + 5\varepsilon$. Therefore, $\overline{val}\Gamma(W) < \underline{val}\Gamma(W) + 10\varepsilon$ for any $\varepsilon > 0$, so $\overline{val}\Gamma(W) \leq \underline{val}\Gamma(W)$, as required.

3.3 Proof of Lemma 3.8

For $k \leq n$, denote the mapping $g_{n,k} : \bar{\Theta}_n \times \bar{B}_{n-1} \rightarrow H_{M_k}^\circ \cup \{*\}$ given by

$$g_{n,k}(\theta_1, \dots, \theta_n, b_1, \dots, b_{n-1}) = \tilde{f}_k(\theta_1, \dots, \theta_k, b_1, \dots, b_{k-1})$$

Let y be an ε -optimal strategy for Player 2 in $\Gamma(W)$ which chooses mixed actions in Δ_n^ε at each even stage n ; such exists by Lemma 3.2. Define a pure strategy y^* of Λ defined by

$$y^*(s_1, b_1, \dots, s_n) = y_n$$

Let x^* be any pure strategy of Player 1 in Λ ; define a behavioral strategy in Γ given, for each $n \in \mathbb{N}$ and $p \in \Theta^n$, by

$$x_n(p) = x_n^*(s_1, b_1, \dots, s_n)(p)$$

where s_1, b_1, \dots, s_n is the finite history of Λ defined inductively by $b_k = x_k^*(s_1, b_1, \dots, s_k)$ for odd k and $b_k = y_k$ for even k , and $s_k = \pi_k(g_{n,k}(p_1, \dots, p_n, \beta_1, \dots, \beta_{n-1}))$.

We will join an (x, y) -random play of Γ and an (x^*, y^*) -random play of Λ with 'almost equal' payoffs. Let $(\Pi_k, \zeta_k, \xi_k, \beta_k, \alpha_k)_{k \in \mathbb{N}}$ be sequence of random variables such that for all n , $\Pi_n = (\Pi_n^1, \Pi_n^2) \in \Theta_n^2$, $\xi_n \in S_n$, $\zeta_n \in T_n$, $\alpha_n \in A_n$, $\beta_n \in B_n$, and (denote $\bar{\Pi}_n = (\Pi_1, \dots, \Pi_n)$, and similarly $\bar{\Pi}_n^j$, $j = 1, 2$, and for $\bar{\xi}_n, \bar{\alpha}_n, \bar{\beta}_n, \bar{\zeta}_n$):

$$P(\Pi_n = p \mid \bar{\xi}_n, \bar{\alpha}_{n-1}, \bar{\Pi}_{n-1}) = \eta(\bar{\xi}_n, \bar{\alpha}_{n-1}, \bar{\Pi}_{n-1})[p] \quad (3.7)$$

$$\zeta_n = f_n(\bar{\Pi}_n, \bar{\beta}_{n-1}) \quad (3.8)$$

$$\beta_n = x_n^*(\bar{\zeta}_n, \bar{\beta}_{n-1}) \text{ for odd } n, \quad (3.9)$$

$$\beta_n = y_n^*(\bar{\zeta}_n, \bar{\beta}_{n-1}) \text{ for even } n, \quad (3.10)$$

$$P(\alpha_n | \bar{\xi}_n, \bar{\alpha}_{n-1}) = \beta_n(\bar{\Pi}_n^{[n]})[\alpha_n] \quad (3.11)$$

$$P(\xi_n = s | \bar{\xi}_{n-1}, \bar{\alpha}_{n-1}) = q(\bar{\xi}_{n-1}, \bar{\alpha}_{n-1})[s] \quad (3.12)$$

where recall that f_n is defined via projection to T'_n of \tilde{f}_n . From (3.8) follows then that, for all n and all $k \leq n$,

$$(\zeta_1, \dots, \zeta_k) = g_{n,k}(\bar{\Pi}_n, \bar{\beta}_{n-1}) \quad (3.13)$$

$\beta_n = y_n$ for even n ; that is,

$$y_n(\bar{\Pi}_n^2) = \beta_n(\bar{\Pi}_n^2) \quad (3.14)$$

Then, using (3.13), (3.9), and (3.14), it follows inductively that

$$x_n(\bar{\Pi}_n^1) = \beta_n(\bar{\Pi}_n^1) \quad (3.15)$$

for all odd n . Putting these last two into (3.11) gives:

$$P(\alpha_n | \bar{\xi}_n, \bar{\alpha}_{n-1}) = \begin{cases} x_n(\bar{\Pi}_n^1) & \text{if } n \text{ is odd} \\ y_n(\bar{\Pi}_n^2) & \text{if } n \text{ is even} \end{cases} \quad (3.16)$$

As such, from (3.7), (3.16), and (3.12) it is deduced that $\xi_1, \alpha_1, \xi_2, \alpha_2, \dots$ is an (x, y) -random play of Γ ; that is, this sequence of random variables distributes as the sequence of states and actions distribute under $P_{x,y}$. Indeed, these three equalities are precisely the dynamics of the stochastic process (ξ_n, Π_n, α_n) under $P_{x,y}$.

Furthermore, by comparing (3.4),(3.5),(3.6) with (3.7),(3.11), (3.12), the distribution of $(\hat{\theta}_n, \hat{z}_n, \hat{a}_n)_{n=1}^\infty$ is the same as the distribution of $(\Pi_n, \xi_n, \alpha_n)_{n=1}^\infty$, given that the choices β_1, \dots, β_n are the same as b_1, \dots, b_n in (3.3). Therefore, for a sequence $\zeta_1, \beta_1, \dots, \zeta_{n-1}, \beta_{n-1}$ which is consistent with x^*, y^* (i.e., for each $k \neq n-1$ odd, $\beta_k = x^*(\zeta_1, \beta_1, \dots, \zeta_k)$, and similarly for k even with y^* instead - note that both x^* and y^* are pure in Λ , and hence the sequence (ζ_n) determines the sequence (β_n)), we have by (3.3)

$$\begin{aligned} P(\zeta_n = z | \zeta_1, \beta_1, \dots, \zeta_{n-1}, \beta_{n-1}) &= P(\zeta_n = z | \zeta_1, \dots, \zeta_{n-1}) \\ &= P(f_n(\bar{\beta}_{n-1}, \bar{\Pi}_n) = z | \forall k < n, f_k(\bar{\beta}_{k-1}, \bar{\Pi}_k) = \zeta_k) \\ &= P_{x^*, y^*}(f_n(\bar{b}_{n-1}, \bar{\theta}_n) = z | \forall k < n, f_k(\bar{b}_{k-1}, \bar{\theta}_k) = \zeta_k) \end{aligned} \quad (3.17)$$

where $\bar{b}_0 = \emptyset$ and, inductively, $b_j = x^*(\zeta_1, b_1, \dots, \zeta_j)$ for odd j , and similarly for even j with y^* (again, recall that x^*, y^* are pure); equivalently, $\beta_j = b_j$ for $j \leq n$. Hence $\zeta_1, \beta_1, \zeta_2, \beta_2, \dots$ is an (x^*, y^*) -random play of Λ .

By (3.13) and Lemma 3.6, we have

$$\begin{aligned} &|P_{x,y}((z_n, a_n)_{n=1}^\infty \in W) - P_{x^*, y^*}(F((s_n)_{n=1}^\infty) \in W)| \\ &= |P((\xi_n, \alpha_n)_{n=1}^\infty \in W) - P(F((\zeta_n)_{n=1}^\infty) \in W)| \\ &\leq P(((\xi_n, \alpha_n)_{n=1}^\infty \in W) \Delta (F((\zeta_n)_{n=1}^\infty) \in W)) \\ &\leq P_{x,y}(\exists n \in \mathbb{N}, \tilde{f}_n(\bar{\beta}_{n-1}, \bar{\theta}_n) \neq (z_1, a_1, \dots, z_{M_n})) < 2\varepsilon \end{aligned}$$

where $\beta_j = x_j$ for odd j and $\beta_j = y_j$ for even j , which completes the proof of Lemma 3.8.

4 Insufficiency of WSEPM

In this section, we will show that even if the game satisfies WSEMP, the value need not exist. In fact, in the example we construct, Player 1 will be fully informed - i.e., will possess perfect monitoring - and it's only Player 2 whose signals are 'blurred'.

4.1 Blurring Signals

Here, we will begin by defining what can be thought of as a decision maker (that is, a single player) choosing at each stage n and action of B_n , resulting in a signal in I_n whose conditional probability depends on all actions chosen until now (but not on previous signals). This defines a transition kernel η from infinite sequences of actions to infinite sequences of signals.

Let D_1, D_2, D_3, \dots be a fixed sequence of action sets (or choice sets) for a decision maker. For every sequence of integers $m = (m_1, m_2, \dots) \in \mathbb{N}^{\mathbb{N}}$, make the following definitions: For each $n \in \mathbb{N}$, let B_n be a set of size m_n , let $C_n = D_n \times B_n$, let $\bar{C}_n = \prod_{k \leq n} C_k$, $M_n = |\bar{C}_n| = \prod_{k \leq n} |D_k| \cdot m_k$, $\bar{C}_\infty = \prod_{k \in \mathbb{N}} C_k$, and let $I_{n+1} = \{0, 1\}^{\bar{C}_n}$ (note that I_1 is a singleton). We define a transition kernel η from \bar{C}_∞ to $\bar{I}_\infty := \prod_{n \in \mathbb{N}} I_n$. The simplest way to describe it is by specifying the distribution $\eta_n(\cdot | b_1, \dots, b_n)$ on I_{n+1} given that c_1, \dots, c_n have been chosen, and this will determine the kernel η as in Section 2.3: η_n is independent of any previous signals, and for any $i_{n+1} \in I_{n+1}$,

$$\eta_{n+1}(i_{n+1} | c_1, \dots, c_n) = \begin{cases} 0 & \text{if } i_{n+1}(c_1, \dots, c_n) = 0 \\ (\frac{1}{2})^{M_n-1} & \text{if } i_{n+1}(c_1, \dots, c_n) = 1 \end{cases} \quad (4.1)$$

In other words, for each of the M_n possible histories up through n actions, Nature performs independent lotteries (also independent of previous lotteries): The true history is assigned 1, while all other histories are assigned 1 or 0 with equal probability.

Now, let ν_n (resp. μ_n) denote the uniform measure on B_n (resp. D_n), and define a strategy for a decision maker who needs to choose a decision from C_n at stage n : The strategy σ_n , at stage n , plays mixed the action $\mu_n \otimes \nu_n$. Call a behavioural strategy σ *normal* if it does not depend on the previous outcomes in B_1, B_2, \dots or on previous signals - it can depend on the previous outcomes in D_1, D_2, \dots - and which plays, at each stage n , a product distribution on $D_n \times B_n$ whose marginal on B_n is the uniform ν_n . The proof of the following Proposition is given in Appendix D.

Proposition 4.1. *For each $\varepsilon > 0$, there are functions $(g_k)_{k=1}^\infty$, $g_k : \mathbb{N} \rightarrow \mathbb{N}$, such that if for all $k \in \mathbb{N}$ it holds that¹² $m_k \geq g_k(\prod_{i < k} m_i)$, then $\|P_\sigma - P_\% \| < \varepsilon$ for any normal strategy σ , where P_σ and $P_\%$ denote the distributions induced on \bar{I}_∞ as a result of using σ or $\sigma_\%$, respectively.*

4.2 The Example

This construction could be done much more generally, but for simplicity, we build a single example and remark below how to generalize (and it can be generalized much further). Begin with the endurance game given in in Section 2.3 of [Shmaya (2011)]: Starting with Player 1, at each stage, players alternatively choose to stay (S) or leave (L), resulting a sequence of choices $h = (d_1, e_1, d_2, e_2, \dots)$. (There are no states.) Let $n^1(h) = \inf\{n \in \mathbb{N} \mid d_n = L\}$, $n^2(h) = \inf\{n \in \mathbb{N} \mid e_n = L\}$, where the infimum of the empty set is ∞ , and define the winning set of Player 1 of this game Γ to be

$$W = \{h \mid n^1(h) > n^2(h) \text{ or } n^1(h) < n^2(h) = \infty\}$$

That is, Player 1 wishes to leave after Player 2, but even if Player 2 is never going to leave, Player 1 wants to leave at some point. It is shown there that if Player 1 has perfect monitoring and Player 2 has no monitoring, then $\underline{val}(\Gamma) = 0$ and $\overline{val}(\Gamma) = 1$.¹³

Now, let $0 < \varepsilon < \frac{1}{4}$, let B_1, B_2, \dots be sets of sizes m_1, m_2, \dots corresponding to this ε and to $D_n = \{S, L\}$ for all n as in Proposition 4.1. Define another game Γ' in which:

- The action sets of Player 1 are A_1, A_2, \dots where $A_k = D_n \times B_k$.
- The action sets of Player 2 are also $E_n = \{S, L\}$.
- Player 1 has perfect monitoring.
- Player 2 has perfect recall (after he plays a move, he observes it perfectly), and his monitoring structure of his opponent's actions is given via the kernel as in Section 4.1. Specifically, if η is the transition kernel from that section,

$$\eta^2(a_1, \theta_1, e_1, a_2, \theta_2, e_2, a_3, \dots, a_n) = \eta(a_1, \dots, a_n)$$

where we have suppressed reference to Player 1's signals and to the perfectly informative signals Player 2 receives about his own actions.

- W' is the inverse image of W via the projection from $A_1 \times E_1 \times A_2 \times E_2 \times \dots$ to $D_1 \times E_1 \times D_2 \times E_2 \times \dots$.

Proposition 4.2. *The game satisfies WSEMP.*

¹²The empty product is 1.

¹³In [?], a loss for Player 1 has payoff 0; while in the current paper, it has payoff -1 .

Proof. Player 1 has perfect monitoring. As for Player 2: Let $h = (a_1, e_1, a_2, e_2, \dots) \neq h' = (a'_1, e'_1, a'_2, e'_2, \dots) \in H_\infty$. If they are different in some action of Player 2, then $\eta^2(h) \perp \eta^2(h')$ since Player 2 has perfect recall. So assume they are different in some action of Player 1. Let V be the subset of Player 2's signal space given by

$$V = \{(i_1, i_2, \dots) \mid \forall n \in \mathbb{N}, i_{n+1}(a_1, e_1, \dots, a_n) = 1\}$$

i.e., which always give a signal 1 all along the history h , and define V' similarly w.r.t. $h' = (a'_1, e'_1, a'_2, e'_2, \dots)$. Then clearly $\eta^2(h)(V) = 1$, $\eta^2(h)(V') = 0$, and the converse equalities as well, since any false history will with probability 1 eventually get a 0 at some point - see the explanation following (4.1). \square

Clearly, since Player 2's strategy set is richer in Γ' than in Γ ,

$$\underline{val}(\Gamma') \leq \underline{val}(\Gamma) = 0$$

So it suffices to show that

Lemma 4.3. $\overline{val}(\Gamma') \geq 1 - 4\varepsilon$

In fact, as our proof will show, this remains true even if we remove all of Player 1's monitoring of Player 2's actions.

Proof. Let σ_n for $n = 1, 2, 3, \dots, \infty$ be the strategy for Player 1 that (ignoring Player 2 actions) plays $L \times \nu_n$ only at time n and $S \times \nu_n$ otherwise (σ_∞ always plays $S \times \nu_n$), where recall that ν_n is uniform on B_n . Fix a strategy τ of Player 2. Let $\delta > 0$, and let $N \in \mathbb{N}$ be such that

$$P_{\sigma_\infty, \tau}(\{h \mid N < n^2(h) < \infty\}) < \delta \quad (4.2)$$

Such N clearly exists. We contend that

$$P_{\sigma_{N+1}, \tau}(W') > 1 - \delta$$

Indeed, by the definition of W' ,

$$P_{\sigma_{N+1}, \tau}(W') = P_{\sigma_{N+1}, \tau}(\{h \mid N \geq n^2(h)\}) + P_{\sigma_{N+1}, \tau}(\{h \mid n^2(h) = \infty\}) \quad (4.3)$$

Observe that

$$P_{\sigma_{N+1}, \tau}(\{h \mid N \geq n^2(h)\}) = P_{\sigma_\infty, \tau}(\{h \mid N \geq n^2(h)\})$$

We will show that

$$P_{\sigma_{N+1}, \tau}(\{h \mid n^2(h) = \infty\}) \geq P_{\sigma_\infty, \tau}(\{h \mid n^2(h) = \infty\}) - 2\varepsilon > (1 - P_{\sigma_\infty, \tau}(\{h \mid N \geq n^2(h)\}) - \delta) - 2\varepsilon$$

These last three equations complete the proof, since $\delta > 0$ is arbitrary. The right inequality follows by (4.2), and the left inequality follows from the following argument: The sequence of signals of Player 2 that he receives about Player 1's actions - denote the space of all such sequences of signals by Θ_∞^2 - is determined

by the strategy of Player 1 only, assuming Player 1 uses strategies like $\sigma_{N+1}, \sigma_\infty$ which disregard Player 2's actions. By our assumptions and Proposition 4.1, the total variation distance between the induced measures on Θ_∞^2 from playing σ_{N+1} or σ_∞ is less than 2ε . τ then determines a transition kernel from Θ_∞^2 to sequences of actions in $\{S, L\}^\infty$ of Player 2. Hence the resulting marginals of $P_{\sigma_{N+1}, \tau}$ and $P_{\sigma_\infty, \tau}$ on Player 2's sequences of actions $\{S, L\}^\infty$ differ in total variation by at most 2ε , and $n^2(\cdot)$ depends only on these actions. \square

These construction could be generalized without too much difficulty in the following way: Begin with any game Γ with winning set W in which Player 1 has perfect monitoring of his opponent's actions, Player 2 has no monitoring of his opponent's actions, and assume that $\Gamma(W)$ does not possess a value. (There are no states.) For convenience, denote the action spaces of Player 1 as D_1, D_2, D_3, \dots , and the action sets of Player 2 as E_1, E_2, E_3, \dots (that is, D_k is used by Player 1 at stage $2k - 1$, and E_k is used by Player 2 at stage $2k$.) Let $0 < \varepsilon$, let B_1, B_2, B_3, \dots be of sizes m_1, m_2, \dots which correspond to D_1, D_2, D_3, \dots and to ε as in Proposition 4.1. The description of the game Γ' - its derivation from Γ - now follows precisely as above. We state without proof:¹⁴

Proposition 4.4. *The monitoring structure of Γ' satisfies WSEMP but, if $\varepsilon > 0$ is small enough, it does not possess a value.*

5 Appendix A

Lemma 5.1. *Let μ be a probability measure on a Borel space X , $\varepsilon > 0$, let $B \in \mathcal{B}(X)$ satisfy $\mu(B) > 1 - \varepsilon$. Define a measure μ_B on X by $\mu_B(A) = \mu(A | B)$. Then $\|\mu - \mu_B\| < 2\varepsilon$.*

Proof. Define $\mu'_B = \mu(B) \cdot \mu_B$; that is, $\mu'_B(A) = \mu(A \cap B)$. Then

$$\|\mu - \mu_B\| \leq \|\mu - \mu'_B\| + \|\mu'_B - \mu_B\| < \mu(B) + \|\mu_B\| \cdot (1 - \mu(B)) < 2\varepsilon$$

\square

Lemma 5.2. *Let X_1, X_2, \dots be finite sets, $X = \prod_{n=1}^\infty X_n$. Let B_n be the partition induced by the first n coordinates, and let μ be a complex measure on X . For each n , let μ_n be the induced measure on¹⁵ B_n defined by $\mu_n(\{B\}) = \mu(B)$ for all $B \in B_n$. Then $\|\mu_n\| \xrightarrow{n \rightarrow \infty} \|\mu\|$.*

Proof. Clearly, $(\|\mu_n\|)_n$ is a non-decreasing sequence which is bounded by above by $\|\mu\|$. Let $\varepsilon > 0$. The space of measures is the dual of $C(X)$, the space of

¹⁴In our example, in the proof of Lemma 4.3, we relied on the existence of near-optional strategies of strategies of Player 1 (in the original game, without signalling) which ignore Player 2's actions - but a more cumbersome argument could have gone through without this assumption.

¹⁵View B_n as a finite set of elements, where each element is a partition class.

continuous functions on X with supremum norm, so there is $f \in C(X)$, $\|f\| \leq 1$, such that

$$\int_X f d\mu > \|\mu\| - \frac{\varepsilon}{2}$$

The Stone-Weierstrauss theorem implies that the set of functions g for which there exists n such that g which is constant on each atom of B_n are dense in $C(X)$. Let g and n be such that g is constant on each atom in B_n and

$$\|g - f\| < \frac{\varepsilon}{2\|\mu\|}, \quad \|g\| \leq 1$$

and we can view g as a function on B_n as well. Therefore,

$$\|\mu_n\| \geq \int_{B_n} g d\mu_n = \int_X g d\mu > \int_X f d\mu - \frac{\varepsilon}{2} > \|\mu\| - \varepsilon$$

□

Lemma 5.3. *Let $X = \prod_{n \in \mathbb{N}} X_n$ with each X_n finite, let $\Omega \subseteq X$ be finitary, and let (Ξ_n) be finitary subsets of Ω such that for each $x \in \Omega$, x belongs to all but finitely many of the (Ξ_n) . Then $\Omega \subseteq \Xi_n$ for all but finitely many n .*

Proof. On the set of vertices, $G = \cup_{n=0}^{\infty} \prod_{j=1}^n X_j$, let $T = \{v \in G \mid (v, w) \in \Omega \text{ for some word } w\}$ - that is, all the vertices which are initial segments of some element in Ω . Edges in T are naturally between an element (x_1, \dots, x_n) and an element $(x_1, \dots, x_n, x_{n+1})$. We can view each vertex $v \in G$ as a subset of X consisting of all elements which have v as an initial segment, and each Ξ_n as a subset of T consisting of all the $v \in T$ such that $v \subseteq \Xi_n$. We observe that if $v \in \Xi_n$ for some n then so are all of its descendants, and hence $T' = T \setminus \Xi$ is also a tree, where $\Xi = \cup \Xi_n$. By assumption it has no infinite branch. By König's lemma, this implies that T' has finite depth - i.e., there is N such that all T' branches are of length N - and this completes the proof. □

6 Appendix B: Proof of the General Theorem 2.12

In this section, we show how to prove Theorem 2.12 when the state space is a general Borel space, assuming that it has already been proven when the state space is finite (but stage-dependent). We begin with some auxiliary results:

Let Λ be any compact space. Let ν be any measure over Λ and let $\sigma : \Lambda \rightarrow \Delta(A)$ be Borel measurable where A is finite. Let $\mu(\nu, \sigma)$ be the measure over $\Lambda \times A$ induced by ν and σ (σ can be viewed as a transition kernel from Λ to A), and let d_w be the Prokhorov metric¹⁶ over $\Delta(\Lambda \times A)$.

Let $\{\mathcal{F}_k\}_k$ be a filtration of measurable sets of Λ that induces the Borel σ -algebra, such that \mathcal{F}_k is finite for every k .

¹⁶Or any other metric that induces the weak* topology.

Lemma 6.1. *Let $\nu \in \Delta(\Lambda)$. For every ϵ there exists a k' such that for every measurable strategy $\sigma : \Lambda \rightarrow \Delta(A)$ there exists a $\mathcal{F}_{k'}$ measurable strategy $\tilde{\sigma} : \Lambda \rightarrow \Delta(A)$ such that,*

$$d_w(\mu(\nu, \sigma), \mu(\nu, \tilde{\sigma})) < \epsilon. \quad (6.1)$$

Proof. It is easy to show that for every fixed strategy σ and ϵ there exists a k' and a $\mathcal{F}_{k'}$ measurable strategy $\tilde{\sigma}$ such that,

$$d_w(\mu(\nu, \sigma), \mu(\nu, \tilde{\sigma})) < \epsilon.$$

We shall show that this property holds uniformly for every σ .

Assume by contradiction that for some $\epsilon > 0$ it holds that for every k there exists a strategy σ_k such that for every \mathcal{F}_k measurable strategy $\tilde{\sigma}$,

$$d_w(\mu(\nu, \sigma_k), \mu(\nu, \tilde{\sigma})) \geq \epsilon. \quad (6.2)$$

By compactness of $\Delta(\Lambda \times A)$, by taking taking subsequence, one can assume that the sequence $\{\mu(\nu, \sigma_k)\}_k$ converges weakly to some μ . Clearly $\mu = \mu(\nu, \sigma)$ for some strategy σ .¹⁷ By the above we can find a strategy $\bar{\sigma}$ that is $\mathcal{F}_{k'}$ measurable for some k' such that,

$$d_w(\mu(\nu, \sigma), \mu(\nu, \bar{\sigma})) < \frac{\epsilon}{2}.$$

On the other hand since $\{\mu(\nu, \sigma_k)\}_k$ converges to $\mu(\nu, \sigma)$ there exists a k_0 such that for every $k > k_0$,

$$d_w(\mu(\nu, \sigma), \mu(\nu, \sigma_k)) < \frac{\epsilon}{2}.$$

Hence by the triangle inequality we get that for every $k > k'$,

$$d_w(\mu(\nu, \bar{\sigma}), \mu(\nu, \sigma_k)) < \epsilon.$$

This stands in contradiction to (6.2). □

Corollary 6.2. *Let $V \subset \Delta(\Lambda)$ be a finite set. For every $\epsilon > 0$ there exists a k' such that for every measurable strategy $\sigma : \Lambda \rightarrow \Delta(A)$ there exists a $\mathcal{F}_{k'}$ -measurable strategy $\tilde{\sigma} : \Lambda \rightarrow \Delta(A)$ such that, if $\mu_A(\nu, \sigma)$ denotes the marginal of $\mu(\nu, \sigma)$ on A , then*

$$\forall \nu \in V, \|\mu_A(\nu, \sigma) - \mu_A(\nu, \tilde{\sigma})\| < \epsilon$$

Proof. For each $\delta > 0$ small enough, one can choose k' and $\mathcal{F}_{k'}$ so that

$$\forall \nu \in V, d_w(\mu(\nu, \sigma), \mu(\nu, \tilde{\sigma})) < \delta. \quad (6.3)$$

and since A is finite, it holds that for $\delta > 0$ small enough and any two measures $\mu, \mu' \in \Delta(\Lambda \times A)$,

$$d_w(\mu, \mu') < \delta \implies \|\mu_A - \mu'_A\| < \epsilon$$

□

¹⁷Since the marginal of each $\mu(\nu, \sigma_k)$ on Λ is ν , so is the marginal of μ .

Proof of Theorem 2.6 (The General Case). Let $\Gamma(W)$ be any general game. By the isomorphism theorem (see Theorem 15.6 in [Kechris(1995)]) the signal space Θ is isomorphic, in a measure theoretic sense, to some compact metric space. Hence, we can assume that the signal space is compact, since we don't use any topological properties of the space, just the measurable structure.

We begin with two simplifications that will make the proof much easier:

- First, we assume that players only receive signals at stages at which they are active - i.e., Player 1 (resp. 2) receives a signal only before he comes to play at the odd (resp. even) stages. This is no loss of generality, as one can have a player receive two signals at once - i.e., the signal space can be modelled as a product space - while allowing a deterministic history-independent 'dummy' signal (i.e., an information-less pre-determined signal) for the players at stages at which they are not active.
- We want the reduction to preserve the perfect recall requirement. To avoid any operation that may ruin the perfect recall, enlarge the signal spaces in the following way: At each stage, the active player receives an additional signal (that is, in addition to the other already prescribed to him in the game) which is deterministically and uniquely determined by his previous actions. (I.e., Each time, before he plays, he is given a list of his previous actions.) Each of these signals is from a discrete, finite space. Formally, the new signal space at each stage is a product of the original signal space Θ and the finite space of possibly past sequences of actions of the currently active player. We will then, implicitly, throughout our construction, only allow partitions of the state space which *preserve perfect recall* - i.e., no partition element contains two signals which give different past sequences of actions of the currently active player.

Recall that $\eta_n^j(h)$ for $j \in \{1, 2\}$, $n \in \mathbb{N}$, and $h \in H_{n-1}$ denotes the distribution on Θ^n , the space of Player j 's first n signals, and that Σ, Υ denote the spaces of Player 1, 2's strategies, respectively. For odd n , let $\overline{\mathcal{F}}_n = \otimes_{k|2k+1 \leq n} \mathcal{F}_{2k+1}$ be the σ -algebra of Player 1's signals up to time n , and similarly for even n for Player 2.

Let $\Gamma(W)$ be any game with a compact set of signals Θ . Fix $\epsilon > 0$. We shall define a sequence of finite measurable partitions $\{\mathcal{F}_n\}_n$ of $\{\Theta^n\}_n$ and a corresponding sequence of classes of strategies Σ_n, Υ_n such that the following conditions hold:

- For $\sigma, \tau \in \Sigma_n, \Upsilon_n$, at each stage $k \leq n$, the active player $i = [k]$ has to choose an $\overline{\mathcal{F}}_k$ -measurable strategy. In stages $k > n$ the players are not restricted in their play.
- The partitions preserve perfect recall, in the sense described above.

- For every strategy $\sigma \in \Sigma$, there is $\sigma' \in \Sigma_n$ such that for every strategy $\tau \in \Upsilon$,

$$\|P_{\sigma,\tau} - P_{\sigma',\tau}\| < (1 - \frac{1}{2^n})\epsilon.$$

and similarly for every $\tau \in \Upsilon$ there is $\tau' \in \Upsilon_n$ satisfying the similar inequality.

We construct the partition inductively as follows.

Assume that $\{\mathcal{F}_k\}_{k < n}$ has been defined such that the two properties hold for every $k < n$. W.l.o.g., assume that n is odd so that player 1 is the active player. Let $V = \{\eta^1(h) \mid h \in H_n^\circ\} \subseteq \Delta(\Theta^n)$ be the finite set of measures induced on Player 1's first n signals by the histories of length $n - 1$, and let

$$U = \{\lambda(\cdot \mid G) \in V \mid \lambda \in V, G \in \prod_{k \mid 2k+1 < n} \mathcal{F}_{2k+1} \text{ s.t. } \lambda(G) > 0\}$$

be the conditional distributions on any element of V w.r.t. atoms of Player 1's knowledge in the coarsened knowledge partitions. Note that U is finite.

For every pair of strategies σ, τ and $n \in \mathbb{N}$ we let $P_{\sigma,\tau}^n$ be the marginal of $P_{\sigma,\tau}$ over H_n .

For every measure $\nu \in U$ and $\sigma_n : \Theta^n \rightarrow A_n$ let $\mu(\nu, \sigma_n)$ be the induced measure over $\Theta^n \times A_n$. We use Corollary 6.2 to construct a partition \mathcal{F}_n over Θ^n such that for every $\sigma_n : \Theta^n \rightarrow A_n$ there exists a \mathcal{F}_n -measurable strategy σ'_n such that,

$$\forall \nu \in U \quad \|\mu_{A_n}(\nu, \sigma_n) - \mu_{A_n}(\nu, \sigma'_n)\| < \frac{\epsilon}{2^n}. \quad (6.4)$$

Since we can also refine this partition and still have the equation hold, we may assume the choice of partition preserves perfect recall.

We claim that the strategies spaces¹⁸ Σ_n, Υ_n have the three properties listed above. The first is just a matter of definition, and the partition was chosen such that the second holds. To see the third let $\sigma \in \Sigma$ any strategy of Player 1. By the inductive construction one can find a strategy $\tilde{\sigma} \in \Sigma_{n-1}$ such that for every strategy $\tau \in \Upsilon$ of Player 2

$$\|P_{\sigma,\tau}^{n-1} - P_{\tilde{\sigma},\tau}^{n-1}\| < (1 - \frac{1}{2^{n-1}})\epsilon, \quad (6.5)$$

and for every $h_{n-1} \in H_{n-1}$,

$$P_{\sigma,\tau}(\cdot \mid h_{n-1}) = P_{\tilde{\sigma},\tau}(\cdot \mid h_{n-1}). \quad (6.6)$$

Define then $\sigma' = (\tilde{\sigma}_1, \dots, \tilde{\sigma}_{n-1}, \sigma'_n, \sigma_{n+1}, \sigma_{n+2}, \dots)$. Furthermore, for each $h_{n-1} \in H_{n-1}$, the measure $\eta^1(h_{n-1})$ is in V , and

$$\|P_{\sigma',\tau}^n(\cdot \mid h_{n-1}) - P_{\sigma,\tau}^n(\cdot \mid h_{n-1})\| = \|\mu_{A_n}(\eta^1(h_{n-1}), \sigma_n) - \mu_{A_n}(\eta^1(h_{n-1}), \sigma'_n)\| < \frac{\epsilon}{2^n} \quad (6.7)$$

¹⁸Note that in this case, where n is odd, $\Upsilon_n = \Upsilon_{n-1}$.

Hence, the marginals of $P_{\sigma,\tau}(\cdot|h_{n-1})$ and $P_{\sigma',\tau}(\cdot|h_{n-1})$ over A_n have a distance of less than $\frac{\epsilon}{2^n}$ between them. Therefore, (6.5) and (6.6) yield that,

$$\|P_{\sigma,\tau} - P_{\sigma',\tau}\| < (1 - \frac{1}{2^n})\epsilon.$$

The above inductive construction yields a sequence of finite measurable partitions $\{\mathcal{F}_n\}_n$ of $\{\Theta^n\}_n$ respectively such that for every strategy $\sigma \in \Sigma$ of Player 1 there exists a strategy $\sigma' = (\sigma'_1, \sigma'_2, \dots)$ such that for every $n \in \mathbb{N}$, σ'_n is $\overline{\mathcal{F}}_n$ measurable for every n , and for every strategy τ of player 2 in $\Gamma(W)$,

$$\|P_{\sigma,\tau} - P_{\sigma',\tau}\| < \epsilon.$$

Let $\Gamma_\infty(W)$ be the game in which players use only $\{\mathcal{F}_n\}_n$ -measurable strategies. The game $\Gamma_\infty(W)$ is of course equivalent to a game where at every stage n the set of signals Θ_n for every player is finite. Therefore the game $\Gamma_\infty(W)$ is determined, by the version of Theorem 2.12 for finite signal spaces. We claim that

$$\underline{val}\Gamma(W) \geq \underline{val}(\Gamma_\infty(W)) - 2\epsilon.$$

To see this note that by construction every strategy σ' in $\Gamma_\infty(W)$ guarantees the same payoff up to 2ϵ in $\Gamma(W)$.¹⁹ Similarly,

$$\overline{val}\Gamma(W) \leq \overline{val}(\Gamma_\infty(W)) + 2\epsilon.$$

Hence,

$$\overline{val}\Gamma(W) - \underline{val}\Gamma(W) \leq 4\epsilon.$$

Since ϵ is arbitrary the game $\Gamma(W)$ is determined. \square

7 Appendix C: Proof of Proposition 3.1

Lemma 7.1. *Let σ, τ be any strategy profile of a game satisfying SEPM and let $k \in \{1, 2\}$. There exists for each $n \in \mathbb{N}$ and for each $h \in H_n^\diamond$ satisfying $P_{\sigma,\tau}(h) > 0$ a strategy profile σ_h, τ_h and a belief of Nature \tilde{q}_h such that $P_{\sigma,\tau,\tilde{q}_h} = P_{\sigma,\tau}(\cdot | h)$, and hence for $h \neq h' \in H_n^\diamond$, $\pi_*^k(P_{\sigma,\tau}(\cdot | h)) \perp \pi_*^k(P_{\sigma,\tau}(\cdot | h'))$, where π^k denotes the projection to Player k 's signal space as in Definition 2.5.*

Proof. Simply define $\sigma_h, \tau_h, \tilde{q}_h$ to make pure choices up through the choice of the n 's state which agree with h , and to agree with σ, τ, q thereafter; since $P_{\sigma,\tau}(h) > 0$, this indeed defines a belief of Nature. The second part then follows from the definition of SEPM. \square

Recall that,

$$\pi_*^k(P_{\sigma,\tau}(\cdot | h)) = \int_{H_\infty} \eta(\omega)(\cdot) d\pi_*^H(P(\omega | h))$$

¹⁹Recall that the payoffs are ± 1 .

where π^H is the projection to H_∞ , as in Definition 2.5. In other words, $\pi_*^k(P_{\sigma,\tau}(\cdot | h))$ is the measure induced by η and $\pi_*^H(P_{\sigma,\tau}(\cdot | h))$. Hence, Proposition 3.1 follows from the previous lemma and the following lemma (by taking $(A_j)_j = (h)_{h \in H_\infty^0}$):

Lemma 7.2. *Let X, Y be standard Borel spaces, let $\mu \in \Delta(X)$ and let A_1, \dots, A_n be disjoint Borel sets which satisfy $\mu(\cup A_j) = 1$, $\mu(A_j) > 0$ for all j ; denote $\mu_j = \mu(\cdot | A_j)$. Let η be a transition kernel from X to Y such that $\eta(\mu_i) := \int_X \eta(\omega) d\mu_i(\omega) \perp \eta(\mu_j)$ for $i \neq j$. Then for any filtration $(\mathcal{F}_n)_n$ of Y generating the Borel σ -algebra²⁰ and each j , $P_\mu(A_j | \mathcal{F}_t) \xrightarrow[t \rightarrow \infty]{} 1 \eta(\mu_j)$ -a.s., where $P_\mu = \int_X \delta_\omega \otimes \eta(\omega) d\mu(\omega)$ is the measure induced on $X \times Y$ by μ and η .*

Proof. Since $(P_\mu(A_j | \mathcal{F}_t))_t$ is a martingale, by the martingale convergence theorem, it suffices to show that $P_\mu(A_j | \mathcal{B}_Y) = 1 \eta(\mu_j)$ -a.s., where \mathcal{B}_Y is the Borel σ -algebra on Y . By assumption there are disjoint B_1, \dots, B_n such that $\eta(\mu_i)(B_j) = 1$ if $i = j$ and 0 if $i \neq j$, and in particular that:

$$P_\mu((A_j \times Y) \Delta (X \times B_j)) = 0, \quad j = 1, \dots, n \quad (7.1)$$

It suffices to show that $P_\mu(A_j | \mathcal{B}) = 1 \eta(\mu)$ -a.s. in B_j ; by the last equation, it suffices to show that $P_\mu(B_j | \mathcal{B}) = 1 \eta(\mu)$ -a.s. in B_j , which is immediate. \square

8 Appendix D: Proof of Lemma 4.1

Fix $\varepsilon > 0$; let $m \in \mathbb{N}^{\mathbb{N}}$, and we will later specify how large its coordinates need to be. It suffices to consider normal σ which makes pure choices in D_1, D_2, D_3, \dots ; let $\bar{d}_\infty = (d_1, d_2, d_3, \dots)$ be the sequence in D_1, D_2, D_3, \dots that σ chooses. For each $k \in \mathbb{N}$, let $P_\sigma^k, P_{\%}^k$ be the induced measures on \bar{I}_{k+1} . It suffices to show by Lemma 5.2 that $\|P_\sigma^k - P_{\%}^k\| < \varepsilon$ for all k in our construction of m . For each $\bar{b}_k = (b_1, \dots, b_k) \in \bar{B}_k$, define an element in \bar{C}_k by:

$$\bar{b}_k^d = (d_1, b_1, d_2, b_2, \dots, d_k, b_k) \in \bar{C}_k$$

Denote $T = \sum_{j=1}^k |\bar{C}_j| = \sum_{j=1}^k \prod_{i \leq j} |C_i|$. Under $P_{\%}^k$, there are 2^T possible sequences of signal in \bar{I}_{k+1} that could arise - that is, $|\bar{I}_{k+1}| = 2^T$ - each with equal probability; hence, for any $\bar{i}_{k+1} \in \bar{I}_{k+1}$,

$$P_{\%}^k(\bar{i}_{k+1}) = \frac{1}{2^T}$$

On the other hand, for each $\bar{b}_k \in \bar{B}_k$, there are 2^{T-k} possible $\bar{i}_{k+1} = (i_1, \dots, i_{k+1}) \in \bar{I}_{k+1}$ that can result if σ is played - those which have $i_j(\bar{b}_k^d|_j) = 1$ for all $j \leq k$, where $|_j$ denotes the projection to $\prod_{i \leq j} C_i$ - all with equal probability. Hence, for any $\bar{i}_{k+1} \in \bar{I}_{k+1}$,

$$P_\sigma^k(\bar{i}_{k+1}) = \frac{1}{|\bar{B}_k|} \sum_{\bar{b}_k \in \bar{B}_k} \frac{1}{2^{T-k}} \prod_{j=1}^k 1_{i_j(\bar{b}_k^d|_j)=1}$$

²⁰A σ -algebra \mathcal{F} on Y implicitly introduces the σ -algebra $\{\emptyset, X\} \times \mathcal{F}$ on $X \times Y$.

Therefore, letting

$$J_k^d(\bar{i}_{k+1}) = \{\bar{b}_k \in \bar{B}_k \mid \forall j \leq k, i_j(\bar{b}_k|_j) = 1\}$$

we have, letting μ denote the uniform measure on \bar{I}_{k+1} , and denoting $B_0^d = \{\emptyset\}$ and $J_0^d = 1$, we see that

$$\begin{aligned} \|P_{\%}^k - P_{\sigma}^k\| &= \sum_{\bar{i}_{k+1} \in \bar{I}_{k+1}} |P_{\%}^k(\bar{i}_{k+1}) - P_{\sigma}^k(\bar{i}_{k+1})| = \frac{1}{2^T} \sum_{\bar{i}_{k+1} \in \bar{I}_{k+1}} \left| 1 - \frac{2^k}{|\bar{B}_k|} J^d(\bar{i}_{k+1}) \right| \\ &= E_{\mu} \left[\left| 1 - \frac{2^k}{|\bar{B}_k|} J_k^d(\bar{i}_{k+1}) \right| \right] \leq \sum_{n=1}^k E_{\mu} \left[\left| \frac{2^{n-1}}{|\bar{B}_{n-1}|} J_{n-1}^d(\bar{i}_{k+1}) - \frac{2^n}{|\bar{B}_n|} J_n^d(\bar{i}_{k+1}) \right| \right] \\ &\leq \sum_{n=1}^k \frac{2^{n-1}}{|\bar{B}_{n-1}|} J_{n-1}^d(\bar{i}_{k+1}) E_{\mu} \left[\left| 1 - \frac{2}{|B_n|} \frac{J_n^d(\bar{i}_{k+1})}{J_{n-1}^d(\bar{i}_{k+1})} \right| \right] \leq \sum_{n=1}^k 2^{n-1} E_{\mu} \left[\left| 1 - \frac{2}{|B_n|} \frac{J_n^d(\bar{i}_{k+1})}{J_{n-1}^d(\bar{i}_{k+1})} \right| \right] \end{aligned}$$

since $J_{n-1}^d(\bar{i}_{k+1}) \leq |\bar{B}_{n-1}|$. Hence, it suffices to prove:

Lemma 8.1. *For all $\varepsilon > 0$, there are functions $(g_k)_{k=1}^{\infty}$, $g_k : \mathbb{N} \rightarrow \mathbb{N}$, such that if for all $k \in \mathbb{N}$, $m_k \geq g_k(\prod_{i < k} m_i)$, then for all $k \in \mathbb{N}$,*

$$E_{\mu} \left[\left| 1 - \frac{2}{|B_k|} \frac{J_k^d(\bar{i}_{\infty})}{J_{k-1}^d(\bar{i}_{\infty})} \right| \right] < \frac{\varepsilon}{2^{2k}} \quad (8.1)$$

(Since we are looking at arbitrary duration in the lemma, we observe the entire sequence of signals \bar{i}_{∞} .) Intuitively, we expect $J^d(\bar{i}_{\infty})$ to be on the order of $\frac{|\bar{B}_k|}{2^k}$, since about half of the possible histories - out of those histories of the form \bar{b}_k^d - are disqualified by the signals at each round.

Proof. Fix $\varepsilon > 0$. Define g_k in the following manner: $g_k(M)$ is such that if $N \geq g_k(M)$, and t_1, \dots, t_N are i.i.d. with values in $\{0, 1\}$ with equal probability, then,

$$P\left(\left| \frac{\sum_{j=1}^n t_j}{N} - \frac{1}{2} \right| > \frac{\varepsilon}{2^{2k+2}}\right) < \frac{\varepsilon}{2^{2k+1}M}$$

Such $g_k(M)$ exists by the weak law of large numbers. Fix $k \in \mathbb{N}$. By assumption, for each $\bar{b}_k \in \bar{B}_k$,

$$P_{\mu} \left(\left| \frac{\sum_{b_k \in B_k} \bar{i}_{\infty}(\bar{b}_{k-1}^d, d_k, b_k)}{|B_k|} - \frac{1}{2} \right| > \frac{\varepsilon}{2^{2k+2}} \right) < \frac{\varepsilon}{2^{2k+1} \prod_{i < k} m_i}$$

and hence since $\prod_{i < k} m_i = |\bar{B}_{k-1}|$,

$$P_{\mu}(\exists \bar{b}_{k-1} \in \bar{B}_{k-1}, \left| \frac{2 \cdot \sum_{b_k \in B_k} \bar{i}_{\infty}(\bar{b}_{k-1}^d, d_k, b_k)}{|B_k|} - 1 \right| > \frac{\varepsilon}{2^{2k+1}}) < \frac{\varepsilon}{2^{2k+1}} \quad (8.2)$$

Observe that

$$J_{k-1}^d(\bar{i}_\infty) = \sum_{\bar{b}_{k-1} \in \bar{B}_{k-1}} \left(\prod_{j < k} \bar{i}_\infty(\bar{b}_j) \right) \quad (8.3)$$

We calculate,

$$\begin{aligned} \frac{J_k^d(\bar{i}_\infty)}{J_{k-1}^d(\bar{i}_\infty)} &= \frac{\sum_{\bar{b}_k \in \bar{B}_k} \left(\left(\prod_{j < k} \bar{i}_\infty(\bar{b}_j) \right) \cdot \bar{i}_\infty(\bar{b}_k) \right)}{J_{k-1}^d(\bar{i}_\infty)} \\ &= \frac{\sum_{\bar{b}_{k-1} \in \bar{B}_{k-1}} \left(\left(\prod_{j < k} \bar{i}_\infty(\bar{b}_j) \right) \cdot \sum_{b_k \in B_k} \bar{i}_\infty(\bar{b}_{k-1}, d_k, b_k) \right)}{J_{k-1}^d(\bar{i}_\infty)} \\ &= \frac{\sum_{\bar{b}_{k-1} \in \bar{B}_{k-1}} \left(\left(\prod_{j < k} \bar{i}_\infty(\bar{b}_j) \right) \cdot \frac{|B_k|}{2} \right)}{J_{k-1}^d(\bar{i}_\infty)} \\ &\quad + \frac{\sum_{\bar{b}_{k-1} \in \bar{B}_{k-1}} \left(\left(\prod_{j < k} \bar{i}_\infty(\bar{b}_j) \right) \cdot \left(\sum_{b_k \in B_k} \bar{i}_\infty(\bar{b}_{k-1}, d_k, b_k) - \frac{|B_k|}{2} \right) \right)}{J_{k-1}^d(\bar{i}_\infty)} \\ &= \frac{|B_k|}{2} + \frac{\sum_{\bar{b}_{k-1} \in \bar{B}_{k-1}} \left(\left(\prod_{j < k} \bar{i}_\infty(\bar{b}_j) \right) \cdot \left(\sum_{b_k \in B_k} \bar{i}_\infty(\bar{b}_{k-1}, d_k, b_k) - \frac{|B_k|}{2} \right) \right)}{J_{k-1}^d(\bar{i}_\infty)} \end{aligned}$$

where \bar{b}_j in the sum denotes the restriction of \bar{b}_k to j coordinates, and we have used the representation (8.3) in the last equality. Hence, by (8.2), repeated use of (8.3), and since

$$\left| \frac{2}{|B_k|} \sum_{b_k \in B_k} \bar{i}_\infty(\bar{b}_{k-1}, d_k, b_k) - 1 \right| \leq 1$$

we have

$$\begin{aligned} E_\mu \left[\left| 1 - \frac{2}{|B_k|} \frac{J_n^d(\bar{i}_\infty)}{J_{k-1}^d(\bar{i}_\infty)} \right| \right] &\leq E_\mu \left[\frac{\sum_{\bar{b}_{k-1} \in \bar{B}_{k-1}} \left(\left(\prod_{j < k} \bar{i}_\infty(\bar{b}_j) \right) \cdot \left| \frac{2}{|B_k|} \sum_{b_k \in B_k} \bar{i}_\infty(\bar{b}_{k-1}, d_k, b_k) - 1 \right| \right)}{J_{k-1}^d(\bar{i}_\infty)} \right] \\ &\leq \frac{\varepsilon}{2^{2k+1}} + \frac{\varepsilon}{2^{2k+1}} = \frac{\varepsilon}{2^{2k}} \end{aligned}$$

□

References

- [Arieli and Levy (2011)] Infinite Sequential Games with Perfect but Incomplete Information, *Int. J. Game Theory*, 40, 207–213.
- [Arieli and Levy (In Preparation.)] Stochastic Eventual Perfect Monitoring, In Preparation.

- [Bertsekas and Shreve (1996)] BERTSEKAS, D. AND SHREVE, S. (1996) *Stochastic Optimal Control: The Discrete Time Case*, Athena Scientific, Belmont, Mass.
- [Blackwell (1969)] BLACKWELL, D. (1969) Infinite G_δ Games of Imperfect Information, *Applicationes Mathematicae*, Vol. X., pp. 99–101, Hugo Steinhaus Jubilee Volume.
- [Fan (1953)] FAN, K. (1953) Minimax theorems, *Proceedings of the National Academy of Sciences of the U.S.A.* 39, 42-47
- [Gale and Stewart (1953)] GALE, D. AND STEWART, F.M. (1953) Infinite Games with Perfect Information, in *Contributions to Theory Theory of Games*, Vol. 2, *Annals of Mathematical Studies* 28, 245–266
- [Kechris(1995)] KECHRIS, A. S. (1995), *Classical Descriptive Set Theory*, Graduate Texts in Mathematics 156. Springer.
- [Maitra et al (1992)] MAITRA, A., PURVES, R., SUDDERTH, W. (1992), A Capacability Theorem in Measurable Gambling Theory, *Trans. of the American Math. Soc.*, 333, 221–249
- [Maitra and Sudderth (1996)] MAITRA, A. AND SUDDERTH, W. (1996), *Discrete Gambling and Stochastic Games* Springer-Verlag, New York.
- [Martin (1975)] MARTIN, D. (1975), Borel Determinacy, *Annals of Mathematics* 102, 363–371
- [Martin (1998)] MARTIN, D. (1998), The Determinacy of Blackwell Games, *The Journal of Symbolic Logic*, 63, 1565–1581
- [Mauldin (1983)] MAULDIN, D., PREISS, D., AND HEINRICH, W. (1983), Orthogonal Transition Kernels, *Annals of Probability*, 11, 970–988
- [Mauldin (1991)] MAULDIN, D. AND WEIZSACKER, H. (1991), Some Orthogonality Preserving Kernels which are not Completely Orthogonal, *Annals of Probability*, 19, 396–400
- [Monderer and Samet (1989)] MONDERER, D., AND SAMET, D. (1989), Approximating Common Knowledge with Common Beliefs, *Games Econ. Behav.*, 1, 170–190.
- [Shmaya (2008)] SHMAYA, E. (2008), Many Inspections are Manipulable, *Theoretical Economics*, 3, 376-382.
- [Shmaya (2011)] SHMAYA, E. (2011) The Determinacy of Infinite Games with Eventual Perfect Monitoring, *Proc. Amer. Math. Soc.*, 139, 3665–3678.
- [Wolfe (1956)] WOLFE, P. (1956), The Strict Determinateness of Certain Infinite Games, *Pacific J. Math.*, 5, 841–847.