



Arieli, I. and Levy, Y. J. (2015) Determinacy of games with Stochastic Eventual Perfect Monitoring. *Games and Economic Behavior*, 91, pp. 166-185. (doi:[10.1016/j.geb.2015.04.003](https://doi.org/10.1016/j.geb.2015.04.003))

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/129889/>

Deposited on: 08 November 2019

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

Determinacy of Games with Stochastic Eventual Perfect Monitoring*

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

April 19, 2015

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.

JEL Codes: D83, C73.

Key words: Zero-sum stochastic games, Determinacy, Stochastic Eventual Perfect Monitoring.

Highlights:

1. We consider zero-sum infinite-horizon stochastic games with imperfect monitoring structures.
2. We study two types of monitoring structures.
3. We show that such games with Borel winning sets are determined under the stronger structure.
4. We provide an example in which determinacy fails under the weaker structure.

*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. The authors thank an anonymous referee for his remarks.

[†]Faculty of Industrial Engineering and Management, Technion- Israel Institute of Technology, iarielei@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

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 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 the other 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 in [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 stochastic 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 - that WSEPM is not sufficient to guarantee determinacy. (In particular, this shows that WSEPM does not imply SEPM.) Properties and equivalent reformulations of the SEPM condition, as well as other applications, can be found in [Arieli and Levy (In Preparation)].

Our result and the distinction with the result of [Shmaya (2011)] can be nicely illustrated using terms from *interactive epistemology*. Shmaya's EPM condition implies that every action played, eventually becomes common knowledge among the players. More precisely, every finite history h of size n that is played with positive probability becomes commonly known among the players after a fixed deterministic time t . In contrast, [Cripps et al. (2008)] consider a two player game with incomplete information and show that in order for the two players to be able to play the efficient equilibrium (coordinate an attack) the state and the fact that the other player is going to play the efficient equilibrium strategy must become approximate common knowledge among the players.

In epistemic terms the SEPM condition is fundamentally different than that of [Shmaya (2011)] and [Cripps et al. (2008)]. 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 Preparation)]. Our determinacy result shows that the eventual common learning - or more precisely, common *p*-belief - is not what is required for determinacy - but rather only mutual learning in the appropriate sense.

2 Model, Examples, and Results

2.1 Preliminary Notation

For a Borel space S ,¹ let $\Delta(S)$ denote the space of regular Borel probability measure on S , endowed with the topology of narrow convergence. Given $\mu, \nu \in \Delta(S)$, the total variation distance² is

$$\|\mu - \nu\| = 2 \sup_{A \subseteq S} |\mu(A) - \nu(A)|$$

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 $n \in \mathbb{N}$, let $[n] = j \in \{1, 2\}$ be such that $j = [n] \bmod 2$.

For Borel spaces X, Y , a mapping $f : X \rightarrow Y$, and an $\nu \in \Delta(X)$, we let $f_*(\nu)$ denote the induced measure on Y given by $f_*(\nu)(A) = \nu(f^{-1}(A))$. By a transition kernel $\eta(\cdot | \cdot)$ from X to Y , we mean a measurable mapping from X to $\Delta(Y)$ - i.e., for each $x \in X$, $\eta(\cdot | x)$ (or, for brevity, $\eta(x)$) is a probability measure on Y such that for each Borel $A \subseteq Y$, $\eta(A | \cdot)$ is Borel measurable. If $\nu \in \Delta(X)$, then $\eta(\nu)$ denotes the induced measure on Y ,

$$\eta(\nu)(A) = \int_X \eta(A | x) d\nu(x)$$

2.2 Definition

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

- A_n is the finite action space used at stage n , respectively.
- W is a subset of $H_\infty := \prod_{n \in \mathbb{N}} A_n$.
- Θ is a standard Borel space of signals.³
- For each $n \in \mathbb{N}$, $\eta_n : \prod_{k < n} A_k \rightarrow \Delta(\Theta^2)$ is the transition kernel of signals.

Denote $H_n = \prod_{k < n} A_k$, $H_* = \cup_n H_n$.

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 player⁴ - denote

¹That is, a space homeomorphic to a Borel subset of a complete separable metric space.

²The supremum is taken over Borel sets, of course. In some of the appendices, the total variation distance is applied to pairs of measures which are not necessarily probability measures, and this definition does not apply; we refer the reader to standard references, e.g., [Rudin (1986)].

³One could also allow for time-dependent signals - as we will later - by taking Θ to be a disjoint union of all signalling spaces.

⁴Since Player 1 plays only at odd stages and Player 2 only at even stages, it's actually not necessary that each player receive a signal at each stage; it would have sufficed if they receive before they play; i.e., every other stage. It is simpler for our notation, however, to assume they each receive a signal at each stage.

the signal to Player j before stage n by θ_n^j ; given $h \in H_n$ the pair (θ_n^1, θ_n^2) is chosen by Nature according to the distribution $\eta_n(h)$; we will denote the marginal on each coordinate by η_n^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 inductively on $(\Theta^2)^n$ - by

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

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

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 let⁵ $\pi_H^j : H_n \rightarrow \prod_{[s]=j, s < n} A_s$ be the projection of j 's actions. Then, for any two $\rho, \lambda \in \Delta(H_n)$ which satisfy $(\pi_H^j)_*(\rho) \perp (\pi_H^j)_*(\lambda)$, we have⁶ $\eta^j(\rho) \perp \eta^j(\lambda)$.

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_∞^j , $j \in \{1, 2\}$, from H_∞ to Θ^∞ ; each infinite history $h \in H_\infty$ induces probability distributions $\eta_\infty^j(h)$ on Θ^∞ for $j = 1, 2$ - that is, probability distributions on the sequence of each players' signals defined for cylindrical sets by

$$\eta_\infty^j(a_1, a_2, \dots)(\{\bar{\theta} \in \Theta^\infty \mid \bar{\theta}_n = p\}) = \eta^j(a_1, \dots, a_{n-1})(\bar{\theta}_n = p), \quad \forall p \in \Theta^n$$

Examples of signalling structures are given in Section 2.8.

2.4 Strategies

A behavioral strategy for Player 1 is a sequence of measurable 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

⁵This is different than π^j used later, which will be the projection to Player j 's signal space.

⁶Recall the notation of Section 2.1.

assumption we made above, players 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 in [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 functions⁷ 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.*

Informally, Shmaya's EPM condition holds if for every two distinct infinite histories of play, there exists a finite time n such that at time n each player can distinguish between the two histories.

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

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

where Σ (resp. Υ) is the space of behavioral strategies for Player 1 (resp. 2). The main result of [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

When dealing with stochastic signals rather than deterministic, there are two natural extensions to Shmaya's EPM condition.

Definition 2.5. *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 $j \in \{1, 2\}$.*

This condition, while simple to state, lacks sufficient strength for our needs, and hence we define:

⁷For 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.

Definition 2.6. Let $\pi^H : \tilde{H}_\infty \rightarrow H_\infty$, (resp. $\pi^1, \pi^2 : \tilde{H}_\infty \rightarrow \Theta^\infty$) be the 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 $j = 1, 2$.

Remark 2.7. Observe that $\pi_*^j(P_{\sigma, \tau}) = \eta^j(\pi_*^H(P_{\sigma, \tau}))$.

The *SEPM* condition states that if two pairs of strategy profiles induce mutually singular distributions over the set of plays then they also induce mutually singular distributions over the set of signals for each player j . In contrast the *WSEPM* require only that every two distinct infinite histories induce singular distributions over the set of signals for each player j .

Remark 2.8. Clearly, these two conditions coincide in the case of deterministic signalling.

We further highlight the distinction between these two conditions in Section 2.8. The main result for sequential games is:

Theorem 2.9. *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 *WSEPM* does not guarantee determinacy. (In particular, *WSEPM* does not imply *SEPM*.)

Section 3 is dedicated to the proof of Theorem 2.13. We rely on a technique introduced by [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 2.10. A two-player zero-sum (sequential⁸) stochastic game with signals is given by a sextuple $\Gamma(W) = ((S_n)_{n \in \mathbb{N}}, (A_n)_{n \in \mathbb{N}}, q, \Theta, (\eta_n)_{n \in \mathbb{N}}, W)$ where:

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

⁸To differentiate from the standard stochastic game models in which players play simultaneously.

We will also denote $\tilde{H}_n = \prod_{k < n} (S_k \times \Theta^2 \times A_k)$, $\tilde{H}_\infty = \prod_{k < \infty} (S_k \times \Theta^2 \times A_k)$, $H_*^\diamond = \cup_n H_n^\diamond$, and $H_* = \cup_n H_n$. 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 , the history of the game up to that point being $h = (z_1, a_1, \dots, z_{n-1}, a_{n-1}, z_n) \in H_n^\diamond$. 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 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_{n+1}^\diamond replacing H_n in Assumption 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 2.11. $\hat{q} = (\hat{q}_1, \hat{q}_2, \dots)$ 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 2.12. Let $\pi^H : \tilde{H}_\infty \rightarrow H_\infty = \prod_{n \in \mathbb{N}} S_t \times A_t$, (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 2.13. *Theorem 2.9 holds for stochastic games with SEPM as well.*

Remark 2.14. [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 2.13 below.

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 2.13 to hold.

Remark 2.16. If one wants to define WSEPM for stochastic games in such a way that SEPM implies WSEPM, one should phrase Definition 2.5 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 SEPM implies WSEPM.

Remark 2.17. 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.

2.8 Examples of Monitoring Structures

Example #1: Suppose $A_1 = \{L, R\}$, while A_n is a singleton for $n \geq 2$; i.e., Player 1 plays once and no one plays thereafter. There are no states. There is a sequence $(\alpha_n)_{n=2}^\infty$ with $0 < \alpha_n < \frac{1}{2}$ and $\alpha_n \rightarrow 0$. At stage $n \geq 2$, a public coin gives H with probability $\frac{1}{2} + \alpha_n$ if L was played (and T otherwise), and with probability $\frac{1}{2} - \alpha_n$ if R was played.⁹

Formally, Player 2's strategy space Υ is trivial; Player 1's space of behavioural strategies Σ are the mixtures of σ_L and σ_R which play L, R , respectively; clearly, the only pair of behavioural strategies which induce singular distributions on $H_\infty = A_1$ are the pair of pure strategies σ_L, σ_R . Hence, to verify SEPM, we need to check (using the notation of Definition 2.6) that $\pi_*^2(P_{\sigma_L}) \perp \pi_*^2(P_{\sigma_R})$. It's also immediate that in this case, SEPM and WSEPM coincide, since there are only two possible histories (they indeed coincide whenever there are finitely many histories - i.e., if after some point, the players' action spaces turn trivial). Clearly,

$$\pi_*^2(P_{\sigma_L}) = \otimes_{n=2}^\infty \left(\frac{1}{2} + \alpha_n, \frac{1}{2} - \alpha_n \right), \quad \pi_*^2(P_{\sigma_R}) = \otimes_{n=2}^\infty \left(\frac{1}{2} - \alpha_n, \frac{1}{2} + \alpha_n \right)$$

By assumption, $\limsup_n \alpha_n < \frac{1}{2}$. The Kakutani criterion, [Kakutani (1948)], establishes that if $\sum_{n=2}^\infty \alpha_n^2 = \infty$, then $\pi_*^2(P_{\sigma_L}) \perp \pi_*^2(P_{\sigma_R})$, and we have SEPM, but if $\sum_{n=2}^\infty \alpha_n^2 < \infty$, then they are equivalent¹⁰ and hence SEPM does not

⁹Player 1 is deterministically informed of his action, and we do not directly specify his signalling structure.

¹⁰That is, absolutely continuous w.r.t. each other.

hold. In particular, if $\alpha_n \equiv \alpha$ for some fixed $0 < \alpha < \frac{1}{2}$, then SEPM holds. For such a case, this can be seen without resorting to the Kakutani criterion, rather as a simple application of the law of large numbers: Define

$$V_L = \{(\theta_1, \theta_2, \dots) \mid \lim_{n \rightarrow \infty} \frac{|\{k \leq n \mid \theta_k = H\}|}{n} = \frac{1}{2} + \alpha\}$$

and V_R similarly with T replacing H . $V_L \cap V_R = \emptyset$, but $\pi_*^2(P_{\sigma_L})(V_L) = 1$ while $\pi_*^2(P_{\sigma_R})(V_R) = 1$.

Example #2: Delayed monitoring is another classical example. The signalling is deterministic (for simplicity, assume no states): There exist two functions $\psi^1, \psi^2 : \mathbb{N} \rightarrow \mathbb{N}$ which are increasing and satisfy $\psi^j(n) \geq n + 1$ for all n . The interpretation is that Player j is informed of his opponent's n -th action at the beginning of stage $\psi^j(n)$. Such deterministic kernels are then injective functions, and hence SEPM is satisfied; see Section 2.5.

Example #3: This is the model studied in [Cripps et al. (2008)]. There are no actions, $S_1 = S$ is a finite stage space in the first stage, and S_n is trivial for $n \geq 2$. Nature chooses an element of S with distribution q (w.l.o.g., $q(s) > 0$ for all $s \in S$). The signal spaces for the players are finite, Θ^1 and Θ^2 . For each $s \in S$, there is a joint distribution $\eta_s \in \Delta(\Theta^1 \times \Theta^2)$, such that if $s \neq s'$, then the marginals η_s^j and $\eta_{s'}^j$ of Player j differ for each $j \in \{1, 2\}$. After the initial choice s by Nature, at each stage $n = 1, 2, 3, \dots$, private signals are chosen i.i.d. according to the distribution η_s at each stage.

Using the notations of Definition 2.12, let \tilde{q}, \tilde{q}' be two beliefs of Nature such that $\pi_*^H(P_{\tilde{q}}) \perp \pi_*^H(P_{\tilde{q}'})$. In this case, this clearly means that $P_{\tilde{q}} = P_q(\cdot \mid A)$ and $P_{\tilde{q}'} = P_q(\cdot \mid A')$ for disjoint $A, A' \subseteq S$. The law of large numbers implies that if for $j \in \{1, 2\}$, we define $(V_s^j)_{s \in S}$ by

$$V_s^j = \{(\theta_1, \theta_2, \dots) \mid \forall \theta^* \in \Theta^j, \lim_{n \rightarrow \infty} \frac{|\{k \leq n \mid \theta_k = \theta^*\}|}{n} = \eta_s^j[\theta^*]\}$$

then for each $j \in \{1, 2\}$, $(V_s^j)_{s \in S}$ are disjoint,

$$\pi_*^j(P_{\tilde{q}})(\cup_{s \in A} V_s^j) = 1 \text{ and } \pi_*^j(P_{\tilde{q}'}) (\cup_{s \in A'} V_s^j) = 1$$

hence SEPM hold (and hence WSEPM, since the set of possible histories is again finite).

Cripps et al. have in fact showed that a stronger type of learning holds in this example, that of *common p-learning*. The differences between the types of learning are discussed in more detail in [Arieli and Levy (In Preparation)].

Example #4: Here is an example where WSEPM holds but SEPM does not. Player 2 has no actions, so for simplicity of notation, let Player 1 play

at each stage (not just the odd stages). Fix some $\alpha \in (0, 1)$. At each stage, Player 1 chooses an action in $\{L, R\}$. (There are no states.) After his n -th choice, there are 2^n possible histories in $H_{n+1} = \{L, R\}^n$. For each of these histories, Nature independently performs a lottery which chooses values 1 or 0: Formally,¹¹ $\theta_{n+1} \in \{0, 1\}^{H_{n+1}}$; for the true history $h \in H_{n+1}$, $\theta_{n+1}(h) = 1$, while for other $h' \in H_{n+1}$, $\theta_{n+1}(h') = 1$ (resp. = 0) with probability α (resp. $1 - \alpha$), and for $h' \neq h''$, $\theta_{n+1}(h')$ and $\theta_{n+1}(h'')$ are independent.

Clearly, then, a false history will be found out because it will almost certainly receive a 0 at some point. Hence, for every $h = (a_1, a_2, \dots) \in H_\infty = \{L, R\}^\infty$, define $V_h = \{(\theta_2, \dots) \mid \forall n \in \mathbb{N}, \theta_{n+1}(a_1, \dots, a_n) = 1\}$, then (in the notation of Definition 2.6), $\eta_\infty^2(h)(V_h) = 1$ while for all other $h' \in H_\infty$, $\eta_\infty^2(h')(V_h) = 0$. Hence WSEPM is satisfied.

However, for $\alpha > \frac{1}{2}$, we contend that SEPM is *not* satisfied; we argue heuristically since a similar but more involved example will be treated formally in Section 4. The signalling process is similar to a branching process, which begins with a single organism, and in which in each generation, each live organism gives birth to two offsprings (after it gives birth it dies); each offspring lives to the next generation with probability α . It is known that if $\alpha > \frac{1}{2}$, then almost surely there are many infinite generation lines. Hence, even though in our case the signal process will lead to an 'infinite generation line' along the true history, it leads in any case to many such infinite lines. Hence, if we take two singular distributions on H_∞ , each of which is very dispersed - for example, one chooses L at the first stage while the other chooses R , and thereafter they both randomise uniformly - if α is close to 1, the resulting distributions will be very very similar. (The dispersion of the measures is required, otherwise one can pin-point the exact infinite generation lines we expect to find, as was done above to show WSEPM.)

On the other hand, if $\alpha < \frac{1}{2}$, it follows readily from the discussion in [Karlin and Taylor, Sec. 3, p. 396] that the probability of having at least one generation line surviving at n generations goes to 0 as $n \rightarrow \infty$, and hence if we impose an infinite generation line along the true history, it will stand out noticeably and the true history could be recovered as the others die out.¹²

¹¹Here the signal space is time-dependent. As was remarked in a footnote in the presentation of the model how to incorporate time-dependent signals: by taking the signal space to be the disjoint union of signal spaces over all periods.

¹²More concretely, given a specific finite history $h \in H_*$, players will learn whether h occurred or not with arbitrarily high reliability, which is what is needed to SEPM (as adverse to WSEPM); see Proposition 3.1 and the remarks after it.

3 Proof of Theorem 2.13

3.1 Preliminaries

We henceforth assume that $\Gamma(W)$ satisfies the SEPM assumption. We shall first make a reduction to the case in which the signal spaces are finite but stage-dependent. Deriving the general form of the theorem from the case in which signals spaces are finite is relegated to Appendix B, and henceforth we will continue under the assumption that signals spaces are finite but state-dependent, with signal space Θ_n (for each player) used at stage n .

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

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

In other words, given that some $h \in H_*^\diamond$ has occurred, as time goes by, there is high probability that a player performing Bayesian updating will associate high probability to h having occurred. When there are no states, this condition was shown to be equivalent to SEPM,¹³ [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.

Corollary 3.2. *Let $h \in H_*^\diamond$, and let (σ, τ) be a strategy profile with $P_{\sigma, \tau}(h) > 0$. If SEPM holds, then for each $j = 1, 2$, $P_{\sigma, \tau}(h|\theta_1^j, \dots, \theta_n^j) \rightarrow 1_h$ in $P_{\sigma, \tau}$ -probability, i.e., for each $\varepsilon > 0$, there is $N \in \mathbb{N}$ such that for $j = 1, 2$ and $n \geq N$,*

$$P_{\sigma, \tau}(|P_{\sigma, \tau}(h|\theta_1^j, \dots, \theta_n^j) - 1_h| > \varepsilon) < \varepsilon \quad (3.2)$$

Proof. Assume $h \in H_\ell^\diamond$; SEPM implies (3.1) holds for h as well as for any of the finitely many other $h' \in H_\ell^\diamond$ for large enough n and each $j = 1, 2$, and hence one deduces that for large enough n and $j = 1, 2$, for each $h' \in H_\ell^\diamond$ with $h' \neq h$,

$$P_{\sigma, \tau}(P_{\sigma, \tau}(h|\theta_1^j, \dots, \theta_n^j) < \varepsilon \mid h') > 1 - \varepsilon \quad (3.3)$$

and multiplying by $P(h')$ and summing over all $h' \in H_\ell^\diamond$ with $h' \neq h$ gives

$$P_{\sigma, \tau}(P_{\sigma, \tau}(h|\theta_1^j, \dots, \theta_n^j) < \varepsilon \mid H_\ell^\diamond \setminus \{h\}) > 1 - \varepsilon \quad (3.4)$$

This together with (3.1) gives (3.2) for large enough n and $j = 1, 2$. \square

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

¹³It's not clear if this equivalence still holds when states are added.

Lemma 3.3. *Let $\varepsilon > 0$, recall that the stage dependent signal spaces $\Theta_1, \Theta_2, \dots$ are finite.¹⁴ There exists a sequence of finite sets $(\Delta_n^\varepsilon)_{n \in \mathbb{N}}$, with $\Delta_n^\varepsilon \subseteq \Delta(A_n)$, such that for any pair of behavioral strategies σ, τ , there is a pair σ', τ' which choose, at each stage n , mixed actions in Δ_n^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.*

For each $\varepsilon > 0$, fix such $(\Delta_k^\varepsilon)_{k \in \mathbb{N}}$, recall that 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_k^ε) . Let $\bar{\Theta}_n = \prod_{k \leq n} \Theta_k$.

Lemma 3.4. *Assume SEPM. For each $h \in H_*^\diamond$, $\varepsilon > 0$, there is $N = N(h, \varepsilon)$, such that for all strategy profiles (σ, τ) in $\Sigma^\varepsilon \times \Upsilon^\varepsilon$, each $j = 1, 2$, and all $n \geq N$, (3.2) holds.*

Proof. Let $h \in H_\ell^\diamond$ and $\varepsilon > 0$. For each n , let Σ_n^ε be the strategies of Player 1 in the first n stages which make choices in $(\Delta_k^\varepsilon)_{k \leq n}$, i.e., $\Sigma_n^\varepsilon = \prod_{[k]=1, k \leq n} (\Delta_k^\varepsilon)^{\bar{\Theta}_k}$, and similarly define $\Upsilon_n^\varepsilon = \prod_{[k]=2, k \leq n} (\Delta_k^\varepsilon)^{\bar{\Theta}_k}$. Denote $V = \cup_{n=0}^\infty \Sigma_n^\varepsilon \times \Upsilon_n^\varepsilon$, and on this set of vertices, define a tree structure, in which v', v'' are connected if for some n , $v' = (\sigma', \tau') \in \Sigma_n^\varepsilon \times \Upsilon_n^\varepsilon$, $v'' = (\sigma'', \tau'') \in \Sigma_{n+1}^\varepsilon \times \Upsilon_{n+1}^\varepsilon$, and the restriction of (σ'', τ'') to the domain of (σ', τ') agrees with (σ', τ') (i.e., if they agree through n stages).

Clearly, every infinite branch of the resulting tree defines an element (σ, τ) of $\Sigma^\varepsilon \times \Upsilon^\varepsilon$. For $n \geq \ell$, let V_n denote all the vertices $(\sigma_n, \tau_n) \in \Sigma_n^\varepsilon \times \Upsilon_n^\varepsilon$ such that for $j = 1, 2$,

$$P_{\sigma_n, \tau_n}(|P_{\sigma_n, \tau_n}(h|\theta_1^j, \dots, \theta_n^j) - 1_h| > \frac{\varepsilon^3}{2}) < \frac{\varepsilon^3}{8} \quad (3.5)$$

and let $V^0 = \cup_{n=\ell}^\infty V_n$. Then, let W be the tree resulting from V when all elements of V^0 and any of their descendants are removed. By Corollary 3.2, W has no infinite branches - every infinite branch includes an element of V^0 . Hence, since each node in W has finite degree, König's lemma implies that W has finite depth $N = N(h, \varepsilon)$; i.e., for $n > N$, $(\Sigma_n^\varepsilon \times \Upsilon_n^\varepsilon) \cap W = \emptyset$. Hence, for $j = 1, 2$, and any $(\sigma, \tau) \in \Sigma^\varepsilon \times \Upsilon^\varepsilon$, there is $K \leq N$ (specifically, some K for which (σ, τ) restricted to the first K stages is in V_K) such that

$$P_{\sigma, \tau}(|P_{\sigma, \tau}(h|\theta_1^j, \dots, \theta_K^j) - 1_h| > \frac{\varepsilon^3}{2}) < \frac{\varepsilon^3}{2} \quad (3.6)$$

and hence by Lemma 5.5, for $n > K$ (and, in particular, for $n > N$) and $j = 1, 2$, (3.2) holds. \square

¹⁴This lemma remains correct even if the state spaces are general; but the proof is straightforward for the case of finite signal spaces.

Remark 3.5. We introduce the following notation, based on Lemma 3.4. Fix $\varepsilon > 0$. For each $k \in \mathbb{N}$, let

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

Proposition 3.6. *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)], and we sketch it for convenience:

Proof. The mapping from pure behavioral strategies - which are compact spaces in the product topology - to expected payoff¹⁵ is upper semi-continuous in this game. Hence, by Fan's minimax theorem, a value exists in mixed strategies.¹⁶ Since the game has perfect recall, behavioral strategies are equivalent to mixed strategies, by Kuhn's theorem. \square

3.2 The Auxiliary Game Λ

Fix $0 < \varepsilon < 1$, let (Δ_n^ε) be the finite action spaces that correspond to ε as in Lemma 3.3, and let $N(\cdot, \cdot)$ be the function in Remark 3.5. (Recall that we are in the case where the signal spaces are finite but state-dependent, $(\Theta_n)_{n=1}^\infty$.) 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 (recall $\bar{\Theta}_n = \prod_{k \leq n} \Theta_k$). Denote $\bar{B}_n = \prod_{k \leq n} B_k$. Define $\mathcal{N} : \mathbb{N} \rightarrow \mathbb{N}$ by $\mathcal{N}(k) = N(k, \frac{\varepsilon}{2^k |H_k^\circ|})$, and let¹⁷

$$K_n = \{k \in \mathbb{N} \mid \mathcal{N}(k) = n\}, \quad T_n = \prod_{k \in K_n} A_{k-1} \times S_k$$

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 $\mathcal{K}(n) = \max\{k \mid \mathcal{N}(k) \leq n\} = \max[\cup_{m \leq n} K_m]$ - that is, the maximal stage approximately learned by stage n - and define for $j = 1, 2$, $\tilde{f}_n^j : \bar{B}_{n-1} \times (\bar{\Theta}_n)^2 \rightarrow H_{\mathcal{K}(n)}^\circ$, where $\tilde{f}_n^j(\beta_{n-1}, \bar{\theta}_n)$, for $\beta_{n-1} \in \bar{B}_{n-1}$

¹⁵This is the only difference from the corresponding lemma in [Shmaya (2011)] - pure strategy profiles here do not yield deterministic payoffs.

¹⁶That is, mixtures of pure behavioral strategies.

¹⁷Note that K_n may be empty - indeed, for growing information lags, it will be empty for most n - and in this case, $T_n = \{\emptyset\}$.

and $\bar{\theta}_n \in \bar{\Theta}_n$, is the $h \in H_{\mathcal{K}(n)}^\circ$ such that¹⁸¹⁹

$$P_{\bar{\beta}_{\mathcal{N}(\mathcal{K}(n))-1}}(h|\bar{\theta}_{\mathcal{N}(\mathcal{K}(n))}^j) > 1 - \frac{\varepsilon}{2^{\mathcal{K}(n)} \cdot |H_{\mathcal{K}(n)}^\circ|}, \quad j = 1, 2 \quad (3.7)$$

(If such h exists, it is unique.²⁰) Otherwise, let $\tilde{f}_n^j(\bar{\beta}_{n-1}, \bar{\theta}_n^j)$ be an arbitrary element of $H_{\mathcal{K}(n)}^\circ$. That is, \tilde{f}_n^j returns, for Player j , what he thinks with high probability has occurred, given the signals $\bar{\theta}_{\mathcal{N}(\mathcal{K}(n))}^j$ he's seen and given the strategies $\bar{\beta}_{\mathcal{N}(\mathcal{K}(n))-1}$ that have been employed, *at the last time* $\mathcal{N}(\mathcal{K}(n))$ *at which he had enough information to make a new deduction* (and *not* at each stage anew).

For brevity, we will write $\tilde{f}_n(\bar{\beta}_{n-1}, \bar{\theta}_n)$ instead of $\tilde{f}_n^1(\bar{\beta}_{n-1}, \bar{\theta}_n)$ (as the following lemma shows, \tilde{f}^1 and \tilde{f}^2 coincide with high probability anyway).

Lemma 3.7. *For any pair of behavioral strategies σ, τ in $\Gamma^\varepsilon, \Upsilon^\varepsilon$,*

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

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

Proof. Let σ, τ in $\Sigma^\varepsilon, \Upsilon^\varepsilon$, $j \in \{1, 2\}$. Observe that for each $k \in \text{range}(\mathcal{K})$, and each $h \in H_k^\circ$, by the definition of f_n^j ,

$$\begin{aligned} \{ \tilde{f}_{\mathcal{N}(k)}^j(\bar{\beta}_{\mathcal{N}(k)-1}, \bar{\theta}_{\mathcal{N}(k)}^j) \neq h \wedge h = (z_1, a_1, \dots, z_k) \} \\ \subseteq \{ P_{\sigma, \tau}(h|\bar{\theta}_{\mathcal{N}(k)}^j) \leq 1 - \frac{\varepsilon}{2^k \cdot |H_k^\circ|} \wedge 1_h(z_1, a_1, \dots, z_k) = 1 \} \\ \subseteq \{ |P_{\sigma, \tau}(h|\bar{\theta}_{\mathcal{N}(k)}^j) - 1_h| \geq \frac{\varepsilon}{2^k \cdot |H_k^\circ|} \} \end{aligned}$$

and the latter event has $P_{\sigma, \tau}$ probability less than $\frac{\varepsilon}{2^k \cdot |H_k^\circ|}$, by definition of $\mathcal{N}(\cdot)$.

¹⁸The distribution induced on $H_{\mathcal{K}(n)}^\circ \times \bar{\Theta}_n$ is dependent only on the strategies in the first $n - 1$ stages, hence the term $P_{\bar{\beta}_{n-1}}(h|\bar{\theta}_n^j)$ is well-defined.

¹⁹Note that $n \geq \mathcal{N}(\mathcal{K}(n))$, and $\mathcal{N}(k) = n$ when $k = \mathcal{K}(n)$.

²⁰Since $1 - \frac{\varepsilon}{2^{\mathcal{K}(n)} |H_{\mathcal{K}(n)}^\circ|} > \frac{1}{2}$.

Hence, noticing that by definition, $\tilde{f}_{\mathcal{N}(\mathcal{K}(n))}^j(\bar{\beta}_{\mathcal{N}(\mathcal{K}(n))-1}, \bar{\theta}_{\mathcal{N}(\mathcal{K}(n))}^j) = \tilde{f}_n^j(\bar{\beta}_{n-1}, \bar{\theta}_n^j)$,

$$\begin{aligned}
& P_{\sigma, \tau}(\exists n \in \mathbb{N}, \tilde{f}_n^j(\bar{\beta}_{n-1}, \bar{\theta}_n^j) \neq (z_1, a_1, \dots, z_{\mathcal{K}(n)})) \\
&= P_{\sigma, \tau}(\exists n \in \mathbb{N}, \tilde{f}_{\mathcal{N}(\mathcal{K}(n))}^j(\bar{\beta}_{\mathcal{N}(\mathcal{K}(n))-1}, \bar{\theta}_{\mathcal{N}(\mathcal{K}(n))}^j) \neq (z_1, a_1, \dots, z_{\mathcal{K}(n)})) \\
&\leq \sum_{n \in \mathbb{N}} P_{\sigma, \tau}(\tilde{f}_{\mathcal{N}(\mathcal{K}(n))}^j(\bar{\beta}_{\mathcal{N}(\mathcal{K}(n))-1}, \bar{\theta}_{\mathcal{N}(\mathcal{K}(n))}^j) \neq (z_1, a_1, \dots, z_{\mathcal{K}(n)})) \\
&= \sum_{k \in \text{range}(\mathcal{K})} P_{\sigma, \tau}(\tilde{f}_{\mathcal{N}(k)}^j(\bar{\beta}_{\mathcal{N}(k)-1}, \bar{\theta}_{\mathcal{N}(k)}^j) \neq (z_1, a_1, \dots, z_k)) \\
&= \sum_{k \in \text{range}(\mathcal{K})} \sum_{h \in H_k^\circ} P_{\sigma, \tau}(\tilde{f}_{\mathcal{N}(k)}^j(\bar{\beta}_{\mathcal{N}(k)-1}, \bar{\theta}_{\mathcal{N}(k)}^j) \neq h \wedge h = (z_1, a_1, \dots, z_k)) \\
&\leq \sum_{k \in \text{range}(\mathcal{K})} \sum_{h \in H_k^\circ} P_{\sigma, \tau}(|P_{\sigma, \tau}(h | \bar{\theta}_{\mathcal{N}(k)}^j) - \mathbf{1}_h| > \frac{\varepsilon}{2^k \cdot |H_k^\circ|}) \leq \sum_{k \in \text{range}(\mathcal{K})} \sum_{h \in H_k^\circ} \frac{\varepsilon}{2^k \cdot |H_k^\circ|} \leq \varepsilon
\end{aligned}$$

and this was for $j = 1, 2$. \square

Let $\pi_n : H_n^\circ \rightarrow T_n$ be defined by projection, and let $F : T_1 \times T_2 \times \dots \rightarrow H_\infty$ be defined by

$$F((\pi_n(u|_n))_{n \in \mathbb{N}}) = u$$

By mildly abusive notation, $F : T_1 \times \dots \times T_n \rightarrow H_{\mathcal{K}(n)}^\circ$, which is the projection of F defined above onto $H_{\mathcal{K}(n)}^\circ$; this is well-defined, as these first $\mathcal{K}(n)$ coordinates in the output of F depends only on the first n coordinates of $\prod_m T_m$. Let $f_k = \pi_k \circ \tilde{f}_k$.

In Λ , Player 1 plays at odd stages, Player 2 plays at even stages, with perfect monitoring. Given $p = (t_1, b_1, \dots, t_{n-1}, b_{n-1}) \in T_1 \times B_1 \times \dots \times T_{n-1} \times B_{n-1}$, we need to define the distribution induced on the next state - i.e., the distribution on T_n - that Nature employs. To do that we define a sequence of auxiliary random variables $(\hat{\theta}_n)_{n=1}^\infty$, $(\hat{z}_n)_{n=1}^\infty$, and $(\hat{a}_n)_{n=1}^\infty$. These variables satisfy for every n that $\hat{\theta}_n \in \Theta_n^2$, $\hat{z}_n \in S_n$, and $\hat{a}_n \in A_n$, and are distributed as follows:

$$P(\hat{z}_n | \hat{z}_1, \hat{a}_1, \dots, \hat{z}_{n-1}, \hat{a}_{n-1}) = q(\hat{z}_1, \hat{a}_1, \dots, \hat{z}_{n-1}, \hat{a}_{n-1})[\hat{z}_n] \quad (3.8)$$

$$P(\hat{a}_n | \hat{\theta}_1, \dots, \hat{\theta}_n) = b_n(\hat{\theta}_1^{[n]}, \dots, \hat{\theta}_n^{[n]})[\hat{a}_n] \quad (3.9)$$

and

$$\begin{aligned}
& P(\hat{\theta}_n | \hat{\theta}_1, \dots, \hat{\theta}_{n-1}, \hat{z}_1, \hat{a}_1, \dots, \hat{z}_{n-1}, \hat{a}_{n-1}, \hat{z}_n) \\
&= \eta_n(\hat{z}_1, \hat{a}_1, \dots, \hat{z}_{n-1}, \hat{a}_{n-1}, \hat{z}_n)[\hat{\theta}_n]
\end{aligned} \quad (3.10)$$

In other words, $(\hat{\theta}_k)_{k \leq n}$, $(\hat{z}_k)_{k \leq n}$, $(\hat{a}_k)_{k < n}$ distribute like a play in Γ if the strategies b_1, b_2, \dots, b_{n-1} were used.

We define the transition function in Λ by

$$\begin{aligned} q_\Lambda(t_1, b_1, \dots, t_{n-1}, b_{n-1})[t] = \\ P(f_n(b_1, \dots, b_{n-1}, \hat{\theta}_1, \dots, \hat{\theta}_n) = t | f_k(b_1, \dots, b_{k-1}, \hat{\theta}_1, \dots, \hat{\theta}_k) = t_k, \forall k < n) \end{aligned} \quad (3.11)$$

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 $F(t_1, t_2, \dots) \in W$ (and receives a payoff of 1 unit from Player 2), and loses otherwise (and receives a payoff of -1). The following is a variant of Lemma 4.4 in [Shmaya (2011)], and the reader is encouraged to read the background material presented succinctly and briefly in Appendix A there.

Lemma 3.8. *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 $(t_1, b_1, t_2, b_2, \dots) \in V$.²¹ Define $G : (\prod_{j \in \mathbb{N}} T_j \times B_j) \rightarrow H_\infty$ by

$$G(t_1, b_1, t_2, b_2, \dots) = F(t_1, t_2, \dots)$$

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

$$F(t_1, t_2, \dots) \in W \iff G(t_1, b_1, t_2, \dots) \in W \iff (t_1, b_1, t_2, \dots) \in G^{-1}(W)$$

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

Lemma 3.9. *For every Borel set $W \subseteq H_\infty$,*

$$\text{val}(\Lambda(W)) - 5\varepsilon \leq \overline{\text{val}}(\Gamma(W))$$

Lemma 3.9 is the heart of the proof. If we take this Lemma as a given, then it is easy to complete the proof of Theorem 2.13: Let $\Gamma(W)$ with W Borel be a game which satisfies the SEPM assumption. Let $\varepsilon > 0$ and let W_0 be a compact subset of W as in Lemma 3.8. Then

$$\text{val}\Gamma(W) \geq \text{val}\Gamma(W_0) = \overline{\text{val}}\Gamma(W_0) \geq \text{val}\Lambda(W_0) - 5\varepsilon > \text{val}\Lambda(W) - 6\varepsilon$$

²¹This is fundamentally different than $\Lambda(W)$, which is defined via the function F .

where the first inequality follows from $W_0 \subseteq W$, the equality from Proposition 3.6, the second inequality from Lemma 3.9, and the third inequality from the choice of W_0 via Lemma 3.8.

In a symmetric fashion, $\overline{val}\Gamma(W) < val\Lambda(W) + 6\varepsilon$. Therefore, $\overline{val}\Gamma(W) < \underline{val}\Gamma(W) + 12\varepsilon$ for any $\varepsilon > 0$, so $\overline{val}\Gamma(W) \leq \underline{val}\Gamma(W)$, as required.

3.3 Proof of Lemma 3.9

Fix Borel $W \subseteq H_\infty$. For $m \leq n$, denote the mapping $g_{n,m} : \overline{\Theta}_n \times \overline{B}_{n-1} \rightarrow H_{\mathcal{K}(m)}^\diamond$ given by

$$g_{n,m}(b_1, \dots, b_{n-1}, \theta_1, \dots, \theta_n) = \tilde{f}_m(b_1, \dots, b_{m-1}, \theta_1, \dots, \theta_m)$$

Let $y = (y_n)_{[n]=2}$ be an ε -optimal strategy for Player 2 in $\Gamma(W)$ for which $y_n(\bar{\theta}_n) \in \Delta_n^\varepsilon$ for every $\bar{\theta}_n \in \overline{\Theta}_n^2$ at each even stage n ; such exists by Lemma 3.3. Define a pure strategy y^* of Λ defined by

$$y_n^*(t_1, b_1, \dots, t_n)(\bar{\theta}_n^2) = y_n(\bar{\theta}_n^2) \quad (3.12)$$

Let x^* be any pure strategy of Player 1 in Λ ; define a behavioral strategy $x = (x_n)_{[n]=1}$ for Player 1 in Γ given by

$$x_n(\bar{\theta}_n^1) = x_n^*(t_1, b_1, \dots, t_n)(\bar{\theta}_n^1) \quad (3.13)$$

where t_1, b_1, \dots, t_n is the finite history of Λ defined inductively by

$$b_k = \begin{cases} x_k^*(t_1, b_1, \dots, t_k) & \text{if } k \text{ is odd} \\ y_k & \text{if } k \text{ is even} \end{cases} \quad (3.14)$$

and

$$t_k = \pi_k(g_{n,k}(\theta_1, \dots, \theta_n, b_1, \dots, b_{n-1})) = f_k(\theta_1, \dots, \theta_k, b_1, \dots, b_{k-1}) \quad (3.15)$$

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 a sequence of random variables with distribution Q such that for all n , $\Pi_n = (\Pi_n^1, \Pi_n^2) \in \Theta_n^2$, $\zeta_n \in T_n$, $\xi_n \in S_n$, $\beta_n \in B_n$, $\alpha_n \in A_n$, and (denote $\overline{\Pi}_n = (\Pi_1, \dots, \Pi_n)$, similarly $\overline{\Pi}_n^j$ for $j = 1, 2$, and for $\bar{\xi}_n, \bar{\alpha}_n, \bar{\beta}_n, \bar{\zeta}_n$):

$$Q(\Pi_n = (\theta_n^1, \theta_n^2) \mid \bar{\xi}_n, \bar{\alpha}_{n-1}, \overline{\Pi}_{n-1}) = \eta_n(\bar{\xi}_n, \bar{\alpha}_{n-1})[(\theta_n^1, \theta_n^2)] \quad (3.16)$$

$$\zeta_n = f_n(\bar{\beta}_{n-1}, \overline{\Pi}_n) \quad (3.17)$$

$$\beta_n = x_n^*(\bar{\zeta}_n, \bar{\beta}_{n-1}) \text{ for odd } n, \quad (3.18)$$

$$\beta_n = y_n^*(\bar{\zeta}_n, \bar{\beta}_{n-1}) \text{ for even } n, \quad (3.19)$$

$$Q(\alpha_n | \bar{\Pi}_n^{[n]}) = \beta_n(\bar{\Pi}_n^{[n]})[\alpha_n] \quad (3.20)$$

$$Q(\xi_n = s | \bar{\xi}_{n-1}, \bar{\alpha}_{n-1}) = q(\bar{\xi}_{n-1}, \bar{\alpha}_{n-1})[s] \quad (3.21)$$

where recall that f_n is defined via projection to T_n of \tilde{f}_n . From (3.17) follows then that, for all n and all $k \leq n$,

$$(\zeta_1, \dots, \zeta_k) = g_{n,k}(\bar{\Pi}_n, \bar{\beta}_{n-1}) \quad (3.22)$$

$\beta_n = y_n$ for even n ; in particular,

$$y_n(\bar{\Pi}_n^2) = \beta_n(\bar{\Pi}_n^2) \quad (3.23)$$

Comparing (3.18) and (3.19) with (3.14) and (3.12), and (3.17) with (3.15), we see inductive that

$$x_n(\bar{\Pi}_n^1) = \beta_n(\bar{\Pi}_n^1) \quad (3.24)$$

for all odd n . Putting these last two into (3.20) gives:

$$Q(\alpha_n | \bar{\Pi}_n^{[n]}) = \begin{cases} x_n(\bar{\Pi}_n^1)[\alpha_n] & \text{if } n \text{ is odd} \\ y_n(\bar{\Pi}_n^2)[\alpha_n] & \text{if } n \text{ is even} \end{cases} \quad (3.25)$$

As such, from (3.16), (3.25), and (3.21) it is deduced that (ξ_n, Π_n, α_n) is an (x, y) -random play of Γ ; that is, this sequence of random variables distributes under Q 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 $P_{x,y}$.

Now, define $\bar{b}_0 = \emptyset$ 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, by (3.23) and (3.24), $\beta_j = b_j$ for $j \leq n$. Then, by comparing (3.8),(3.9),(3.10) with (3.16),(3.20), (3.21), the distribution of $(\hat{\theta}_n, \hat{z}_n, \hat{a}_n, b_n)_{n=1}^\infty$ under P is the same as the distribution of $(\Pi_n, \xi_n, \alpha_n, \beta_n)_{n=1}^\infty$ under Q . Therefore (recall that both x^* and y^* are pure in Λ , and hence by (3.18) and (3.19) the sequence (ζ_n) determines the sequence (β_n)), we have by (3.11),

$$\begin{aligned} Q(\zeta_n = t \mid \zeta_1, \beta_1, \dots, \zeta_{n-1}, \beta_{n-1}) &= Q(\zeta_n = t \mid \zeta_1, \dots, \zeta_{n-1}) \\ &= Q(f_n(\bar{\beta}_{n-1}, \bar{\Pi}_n) = t \mid \forall k < n, f_k(\bar{\beta}_{k-1}, \bar{\Pi}_k) = \zeta_k) \\ &= P(f_n(\bar{b}_{n-1}, \hat{\theta}_1, \dots, \hat{\theta}_n) = t \mid \forall k < n, f_k(b_{k-1}, \hat{\theta}_1, \dots, \hat{\theta}_{n-1}) = \zeta_k) \\ &= q_\Lambda(t \mid \zeta_1, b_1, \dots, \zeta_{n-1}, b_{n-1}) = q_\Lambda(t \mid \zeta_1, \beta_1, \dots, \zeta_{n-1}, \beta_{n-1}) \end{aligned}$$

Hence $\zeta_1, \beta_1, \zeta_2, \beta_2, \dots$ is an (x^*, y^*) -random play of Λ (i.e., distributes by P_{x^*, y^*}).

Hence, by (3.22) and Lemma 3.7, and noting that Q -a.s. $F(\bar{\zeta}_n) = \tilde{f}_n(\bar{\Pi}_n, \bar{\beta}_{n-1})$,

$$\begin{aligned}
& |P_{x,y}((z_n, a_n)_{n=1}^\infty \in W) - P_{x^*,y^*}(F((s_n)_{n=1}^\infty) \in W)| \\
&= |Q((\xi_n, \alpha_n)_{n=1}^\infty \in W) - Q(F((\zeta_n)_{n=1}^\infty) \in W)| \\
&\leq Q(((\xi_n, \alpha_n)_{n=1}^\infty \in W) \Delta (F((\zeta_n)_{n=1}^\infty) \in W)) \\
&\leq Q(\exists n \in \mathbb{N}, F(\bar{\zeta}_n) \neq (\xi_1, \alpha_1, \dots, \xi_n)) \\
&= P_{x,y}(\exists n \in \mathbb{N}, \tilde{f}_n(\bar{b}_{n-1}, \bar{\theta}_n) \neq (z_1, a_1, \dots, z_{\mathcal{K}(n)})) \leq 2\varepsilon
\end{aligned}$$

where $\beta_k = b_k = x_k$ for odd k and $\beta_k = b_k = y_k$ for even k . Hence, since y was ε -optimal in $\Gamma(W)$ and recalling that payoffs to the players are ± 1 , gives

$$2 \cdot P_{x^*,y^*}(F((s_n)_{n=1}^\infty) \in W) - 1 \leq 2 \cdot (P_{x,y}((z_n, a_n)_{n=1}^\infty \in W) + 2\varepsilon) - 1 \leq \overline{val}(\Gamma(W)) + 5\varepsilon$$

Taking the supremum over all x^* gives and recalling that $\Lambda(W)$ is determined (Lemma 3.8) gives

$$val(\Lambda(W)) = \overline{val}(\Lambda(W)) \leq \overline{val}\Gamma(W) + 5\varepsilon$$

which completes the proof of Lemma 3.9.

4 Insufficiency of WSEPM

In this section, we will show that even if the game satisfies WSEPM, 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 an action from²² $C = \{S, L\} \times \{\uparrow, \downarrow\}$, resulting in a signal in I_{n+1} defined below whose conditional probability depends on all actions chosen until now, as we describe below. This defines a transition kernel η from infinite sequences of actions to infinite sequences of signals. Denote $D = \{S, L\}$ and $B = \{\uparrow, \downarrow\}$; so that $C = D \times B$. Fix a parameter $0 < \alpha < 1$.

Let $I_{n+1} = \{0, 1\}^{C^n}$ (note that $I_1 = \{\emptyset\}$). We define a transition kernel η from $C^\mathbb{N}$ to $\bar{I}_\infty := \prod_{n \in \mathbb{N}} I_n$. The simplest way to describe it is by specifying the distribution $\eta_{n+1}(\cdot | c_1, \dots, c_n)$ 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): for any $i_{n+1} \in I_{n+1}$,

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

²² $\{S, L\}$ is interpreted as 'Stay' or 'Leave'; this interpretation will be used in our example in the next section.

where $K_{n+1}(i_{n+1}) = \#\{(c'_1, \dots, c'_n) \in C^n \mid i_{n+1}(c'_1, \dots, c'_n) = 1\}$. In other words, for each of the $|C^n| (= 4^n)$ possible histories up through n actions, Nature performs independent lotteries: The true history is assigned 1, while all other histories are assigned 1 or 0 with probabilities $\alpha, 1 - \alpha$ respectively.

Now, call a behavioural strategy σ *normal* if it does not depend on the previous outcomes in B or on previous signals - it can depend on the previous outcomes in D - and which plays, at each stage n , a product distribution on $D \times B$ whose marginal on B is the uniform $(\frac{1}{2}, \frac{1}{2})$. Formally,

$$\sigma(d_1, b_1, i_1, \dots, d_n, b_n, i_n) = \tilde{\sigma}(d_1, \dots, d_n) \otimes \left(\frac{1}{2}, \frac{1}{2}\right)$$

for some $\tilde{\sigma} : \cup_{n=0}^{\infty} D^n \rightarrow \Delta(D)$. The proof of the following Proposition is somewhat tedious and relies on results from the theory of branching processes, and is given in Appendix D.

Proposition 4.1. *For each $\varepsilon > 0$, there is $0 < \alpha_\varepsilon < 1$, such that if $1 > \alpha > \alpha_\varepsilon$, then $\|P_\sigma - P_{\sigma'}\| < \varepsilon$ for any normal strategies σ, σ' , where P_σ and $P_{\sigma'}$ denote the distributions induced on \bar{I}_∞ as a result of using σ or σ' , 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 it (and it can be generalized much further). Begin with the endurance game Γ given 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 $\inf \emptyset = \infty$, 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 2 has no monitoring, then $\underline{val}(\Gamma) = 0$ and $\overline{val}(\Gamma) = 1$;²³ this is regardless of Player 1's monitoring structure.

Now, let $0 < \varepsilon < \frac{1}{4}$, let $B = \{\uparrow, \downarrow\}$ and $D = \{S, L\}$. Define another game Γ' in which:

- The action set of Player 1 each time he plays is $A = D \times B$.
- The action set of Player 2 each time he plays is $E = D = \{S, L\}$.

²³In [Shmaya (2011)], a loss for Player 1 has payoff 0; while in the current paper, it has payoff -1 .

- Player 1 has an arbitrary monitoring structure (with perfect recall).
- 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 for some fixed $1 > \alpha > \alpha_\varepsilon$. Specifically, if η is the transition kernel from that section,

$$\eta^2(a_1, e_1, a_2, e_2, a_3, \dots, a_n) = \eta(a_1, \dots, a_n) \otimes_{j=1}^{n-1} \delta_{e_j}$$

where the term $\otimes_{j=1}^{n-1} \delta_{e_j}$ refers to the perfect recall.

- W' is the inverse image of W via the projection from $(A \times E)^\mathbb{N}$ to $(D \times E)^\mathbb{N}$.

Note that in our notation, a_n (resp. e_n) is played by Player 1 (resp. 2) at stage $2n - 1$ (resp. $2n$).

Proposition 4.2. *The game satisfies WSEPM.*

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 $\eta^2(h)(V) = 1$, while $\eta^2(h)(V') = 0$, since if $a_K \neq a'_K$, then²⁴ $\eta^2(h)(V') \leq \lim_{n \rightarrow \infty} \alpha^{n-K} = 0$, and symmetrically, $\eta^2(h')(V') = 1$, $\eta^2(h')(V) = 0$ \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 $L \times (\frac{1}{2}, \frac{1}{2})$ only at his n -th turn and $S \times (\frac{1}{2}, \frac{1}{2})$ otherwise (σ_∞ always plays $S \times (\frac{1}{2}, \frac{1}{2})$). Fix a strategy τ of Player 2. Let $\delta > 0$, and let $N \in \mathbb{N}$ be such that

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

²⁴Since false histories only get signal 1 with probability α at each stage.

Such N clearly exists. We contend that

$$P_{\sigma_{N+1},\tau}(W') > 1 - \delta - \varepsilon \quad (4.3)$$

Indeed, by the definition of W' ,

$$P_{\sigma_{N+1},\tau}(W') = P_{\sigma_{N+1},\tau}(\{h \mid N \geq n^2(h)\}) + P_{\sigma_{N+1},\tau}(\{h \mid n^2(h) = \infty\}) \quad (4.4)$$

Observe that

$$P_{\sigma_{N+1},\tau}(\{h \mid N \geq n^2(h)\}) = P_{\sigma_\infty,\tau}(\{h \mid N \geq n^2(h)\}) \quad (4.5)$$

We will show that

$$P_{\sigma_{N+1},\tau}(\{h \mid n^2(h) = \infty\}) \geq P_{\sigma_\infty,\tau}(\{h \mid n^2(h) = \infty\}) - \varepsilon > (1 - P_{\sigma_\infty,\tau}(\{h \mid N \geq n^2(h)\}) - \delta) - \varepsilon \quad (4.6)$$

The right inequality follows by (4.2), and the left inequality follows from the following argument (which makes repeated use of Lemma 5.2): The distribution sequence of signals of Player 2 that he receives about Player 1's actions - denote the space of all such sequences of signals²⁵ 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 ε . 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 ε , and $n^2(\cdot)$ depends only on these actions.

(4.4), (4.5), and (4.6) complete the proof of (4.3), which completes the proof of the lemma since $\delta > 0$ is arbitrary. \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 2 has no monitoring of his opponent's actions such 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 $0 < \alpha < 1$. The description of the game Γ' - its derivation from Γ - now follows precisely as above, with D_n replacing $D = \{S, L\}$ (resp. E_n replacing $E = \{S, L\}$) at stage n . We state without proof:²⁶

Proposition 4.4. *The monitoring structure of Γ' satisfies WSEPM but, if α is close enough to 1, it does not possess a value.*

²⁵This space excludes his monitoring of his own past actions.

²⁶In our example, in the proof of Lemma 4.3, we relied on the existence of near-optimal 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.

5 Appendix A

This section contains several short but technical results, which may be of independent interest; recall that $\|\cdot\|$ denotes the total-variation norm.²⁷

Lemma 5.1. *Let (X, \mathcal{F}) be a measurable space, let μ be a complex measure on \mathcal{F} , and let $(\mathcal{F}_n)_{n=1}^\infty$ be a filtration generating \mathcal{F} . Let μ_n be the restriction of μ to \mathcal{F}_n . Then $\|\mu_n\| \rightarrow \|\mu\|$.*

Proof. That $\|\mu_n\| \leq \|\mu\|$ for all n follows by definition. If $\|\mu\| = 0$, we're done; otherwise, normalise μ so that $\|\mu\| = 1$, hence²⁸ $|\mu| \in \Delta(X)$. We have $|\frac{d\mu}{d|\mu|}| \equiv 1$, and for any \mathcal{F}_n -measurable function g , $\int_X g d\mu_n = \int_X g d\mu$. Hence,²⁹

$$\int_X E_{|\mu|}[\frac{\overline{d\mu}}{d|\mu|} | \mathcal{F}_n] d\mu_n = \int_X E_{|\mu|}[\frac{\overline{d\mu}}{d|\mu|} | \mathcal{F}_n] d\mu = \int_X E_{|\mu|}[\frac{\overline{d\mu}}{d|\mu|} | \mathcal{F}_n] \frac{d\mu}{d|\mu|} d|\mu|$$

Hence, by the martingale convergence theorem and applying the bounded convergence theorem to both sides,

$$\int_X E_{|\mu|}[\frac{\overline{d\mu}}{d|\mu|} | \mathcal{F}_n] d\mu_n \rightarrow \int_X \frac{\overline{d\mu}}{d|\mu|} \cdot \frac{d\mu}{d|\mu|} d|\mu| = \int_X d\mu = \|\mu\|$$

Since $|E_{|\mu|}[\frac{\overline{d\mu}}{d|\mu|} | \mathcal{F}_n]| \leq 1$, this implies $\liminf_{n \rightarrow \infty} \|\mu_n\| \geq \|\mu\|$. \square

Lemma 5.2. *Let X, Y be Borel space, η be a transition kernel from X to Y . Let $\mu, \nu \in \Delta(X)$. Then³⁰*

$$\|\eta(\mu) - \eta(\nu)\| \leq \|\mu - \nu\|$$

Proof.

$$\begin{aligned} \|\eta(\mu) - \eta(\nu)\| &= \left\| \int_X \eta(\cdot | x) d\mu(x) - \int_X \eta(\cdot | x) d\nu(x) \right\| \\ &\leq \int_X \|\eta(\cdot | x)\| |d\mu - \nu|(x) = \int_X |d\mu - \nu|(x) = \|\mu - \nu\| \end{aligned}$$

\square

Lemma 5.3. *Let (X, \mathcal{F}) be a measurable space, let (\mathcal{F}_n) be a filtration generating \mathcal{F} , let $\mu, \nu \in \Delta(X)$ with $\mu \ll \nu$, and define $\mu_n \in \Delta(X)$ on (X, \mathcal{F}) by³¹*

$$\mu_n(B) = \int_B E_\nu[\frac{d\mu}{d\nu} | \mathcal{F}_n] d\nu$$

Then $\mu_n \rightarrow \mu$ in norm and $\frac{d\mu_n}{d\nu}$ is \mathcal{F}_n -measurable.

²⁷The reader is referred to standard references, e.g., [Rudin (1986)], for background measure-theoretical tools such as the total variation norm, total variation measure, and the Radon-Nikodym derivative.

²⁸ $|\mu|$ is the total variation measure.

²⁹ \bar{z} denotes the complex conjugate of z .

³⁰Recall the notation of Section 2.1.

³¹Note that $\mu_n(X) = 1$ and $\mu_n \geq 0$; hence, $\mu_n \in \Delta(X)$.

Proof. Clearly, $\frac{d\mu_n}{d\nu} = E_\nu[\frac{d\mu}{d\nu} \mid \mathcal{F}_n]$ and hence is \mathcal{F}_n -measurable. Since $\mu_n \ll \nu$ and $\mu \ll \nu$, it suffices to show that $E_\nu[\frac{d\mu}{d\nu} \mid \mathcal{F}_n] \rightarrow \frac{d\mu}{d\nu}$ in $L^1(\nu)$;³² this follows from the martingale convergence theorem. \square

Corollary 5.4. *Let (X, \mathcal{F}) be a measurable space, let (\mathcal{F}_n) be a filtration generating \mathcal{F} , let $\nu_1, \dots, \nu_M \in \Delta(X)$, and let $\varepsilon > 0$. Then there exists $\mu_1, \dots, \mu_M \in \Delta(X)$ and $N \in \mathbb{N}$ such that $\|\mu_k - \nu_k\| < \varepsilon$ and $\nu_k \ll \mu_k$ for all k , and for any j, k , μ_j, μ_k are equivalent³³ and $\frac{d\mu_k}{d\mu_j}$ is \mathcal{F}_N -measurable.*

Proof. First define for each $k = 1, \dots, M$,

$$\eta_k = (1 - \frac{\varepsilon}{2})\nu_k + \frac{\varepsilon}{2} \frac{1}{M} \sum_{m=1}^M \nu_m$$

The (η_k) are equivalent. Then, for each $k = 1, \dots, M - 1$ and each $n \in \mathbb{N}$, denote

$$\mu_k^n(B) = \int_B E_{\eta_{k+1}}[\frac{d\eta_k}{d\eta_{k+1}} \mid \mathcal{F}_n] d\eta_{k+1}$$

Also let $\mu_M^n = \nu_M$ for all n . Then for each n and each $k = 1, \dots, M - 1$, $\frac{d\mu_k^n}{d\mu_{k+1}^n}$ is \mathcal{F}_n -measurable (by Lemma 5.3), and hence if $j < k$, $\frac{d\mu_j^n}{d\mu_k^n} = \prod_{m=j}^{k-1} \frac{d\mu_m^n}{d\mu_{m+1}^n}$ and $\frac{d\mu_k^n}{d\mu_j^n} = (\frac{d\mu_j^n}{d\mu_k^n})^{-1}$ are \mathcal{F}_n -measurable too. Hence, letting $\mu_k = \mu_k^N$ for each $k = 1, \dots, M$ and N large enough gives the corollary by Lemma 5.3. \square

Lemma 5.5. *Let (Ω, \mathcal{B}, P) be a probability space, let $\mathcal{F} \subseteq \mathcal{G} \subseteq \mathcal{B}$ be σ -algebras, let $0 < \varepsilon < 1$ and let $f : \Omega \rightarrow [0, 1]$ be \mathcal{B} -measurable with $|f| \leq 1$, and suppose that*

$$P(|E(f \mid \mathcal{F}) - f| > \frac{\varepsilon^3}{2}) < \frac{\varepsilon^3}{8} \quad (5.1)$$

Then

$$P(|E(f \mid \mathcal{G}) - f| > \varepsilon) < \varepsilon \quad (5.2)$$

Proof. By (5.1), since $|E(f \mid \mathcal{F}) - f| \leq 2$,

$$\int_{\Omega} |E(f \mid \mathcal{F}) - f|^2 dP < \frac{\varepsilon^3}{8} \cdot 4 + (1 - \frac{\varepsilon^3}{8})(\frac{\varepsilon^3}{2})^2 < \varepsilon^3$$

But since the conditional expectation operator $E(\cdot \mid \mathcal{F})$ is a projection in $L^2(P)$ to the space of \mathcal{F} -measurable functions, and similarly for \mathcal{G} ,

$$\int_{\Omega} |E(f \mid \mathcal{G}) - f|^2 dP \leq \int_{\Omega} |E(f \mid \mathcal{F}) - f|^2 dP$$

Furthermore, by Markov's inequality,

$$\varepsilon^2 P(|E(f \mid \mathcal{G}) - f| > \varepsilon) \leq \int_{\Omega} |E(f \mid \mathcal{G}) - f|^2 dP$$

³²If μ, η, ν are measures with $\nu \geq 0$ and $\mu, \eta \ll \nu$, then $\|\mu - \eta\| = \|\frac{d\mu}{d\nu} - \frac{d\eta}{d\nu}\|_{L^1(\nu)}$

³³I.e., for each j, k , ν_j and ν_k have the same null sets.

Hence,

$$\varepsilon^2 P(|E(f | \mathcal{G}) - f| > \varepsilon) < \varepsilon^3$$

as required. \square

6 Appendix B: Proof of the General Theorem 2.13

In this section, we prove Theorem 2.13 when the signal space Θ is a general Borel space, assuming that it has already been proven when the signal space is finite (but stage-dependent). We begin with an auxiliary result:

Lemma 6.1. *Let Λ be any Borel space. Let $\{\mathcal{F}_k\}_k$ be a filtration of measurable sets of Λ that generates the Borel σ -algebra. Let $V \subseteq \Delta(\Lambda)$ be finite, and let A be a finite set. For every $\varepsilon > 0$ it holds for every N large enough that for every measurable strategy $\sigma : \Lambda \rightarrow \Delta(A)$ there exists an \mathcal{F}_N -measurable³⁴ strategy $\tilde{\sigma} : \Lambda \rightarrow \Delta(A)$ such that*

$$\forall \nu \in V, \|\sigma_*(\nu) - \tilde{\sigma}_*(\nu)\| < \varepsilon$$

Proof. Enumerate $V = \{\nu_1, \dots, \nu_M\}$ and let $\varepsilon > 0$; by Corollary 5.4, there are $V' = \{\mu_1, \dots, \mu_M\}$ and $N \in \mathbb{N}$ such that $\|\mu_k - \nu_k\| < \frac{\varepsilon}{2}$ and $\nu_k \ll \mu_k$ for each k , any two measures in V' are equivalent, and the Radon-Nikodym derivatives of any two measures of V' w.r.t. each other is \mathcal{F}_N -measurable. We contend that for any $\mu_k, \mu_m \in V'$ and $n \geq N$,

$$\int E_{\mu_k}[\sigma | \mathcal{F}_n] d\mu_m = \int \sigma d\mu_m = \sigma_*(\mu_m) \quad (6.1)$$

Indeed, for $n \geq N$, $\frac{d\mu_m}{d\mu_k}$ is \mathcal{F}_n -measurable, hence

$$\begin{aligned} \int E_{\mu_k}[\sigma | \mathcal{F}_n] d\mu_m &= \int E_{\mu_k}[\sigma | \mathcal{F}_n] \frac{d\mu_m}{d\mu_k} d\mu_k \\ &= \int E_{\mu_k}[\sigma \frac{d\mu_m}{d\mu_k} | \mathcal{F}_n] d\mu_k = \int \sigma \frac{d\mu_m}{d\mu_k} d\mu_k = \int \sigma d\mu_m \end{aligned}$$

Given σ , define³⁵

$$\tilde{\sigma} = \frac{1}{|V|} \sum_{k=1}^M E_{\mu_k}[\sigma | \mathcal{F}_N]$$

Then for each $k = 1, \dots, M$, by Lemma 5.2,

$$\|\sigma_*(\mu_k) - \tilde{\sigma}_*(\nu_k)\| < \frac{\varepsilon}{2}$$

³⁴Hence, also \mathcal{F}_n -measurable for all $n \geq N$.

³⁵These are only defined μ_k -a.e.; they can be completed in an arbitrary measurable way on null sets.

$$\|\tilde{\sigma}_*(\mu_k) - \tilde{\sigma}_*(\nu_k)\| < \frac{\varepsilon}{2}$$

Hence, for $j = 1, \dots, M$,

$$\begin{aligned} \|\tilde{\sigma}_*(\nu_k) - \sigma_*(\nu_k)\| &< \varepsilon + \|\tilde{\sigma}_*(\mu_k) - \sigma_*(\mu_k)\| \\ &\leq \varepsilon + \frac{1}{|V|} \sum_{k=1}^M \left\| \int_{\Lambda} E_{\mu_k}[\sigma \mid \mathcal{F}_N] d\mu_j - \sigma_*(\mu_j) \right\| = \varepsilon \end{aligned}$$

where we have used (6.1). □

Proof of Theorem 2.9 (The General Case). Let $\Gamma(W)$ be any general game that satisfies SEPM. In order to prove the determinacy of $\Gamma(W)$ we shall use a reduction of the game $\Gamma(W)$ to a new game $\Gamma^\varepsilon(W)$ which has a finite, stage dependent set of signals. We begin with a simplifying assumption that will make the proof much easier:

- (PR) 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 that already prescribed to him in the game) which is deterministically and uniquely determined by his previous actions. (I.e., each stage before he plays, he is told his previous action.) 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 possible past sequences of actions of the currently active player.

Fix $\varepsilon > 0$. Recall that $\eta^j(h)$ for $j \in \{1, 2\}$, $n \in \mathbb{N}$, and $h \in H_n$ 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. We shall define inductively two sequences of finite measurable partitions $\{\mathcal{F}_n^j\}_n$ of $\{\Theta^n\}_n$, one for each player $j = 1, 2$ such that for every n , \mathcal{F}_n^j is a finite Borel partition of Θ^n . For such partitions and a given n let $\bar{\mathcal{F}}_n^j = \otimes_{k \leq n} \mathcal{F}_k^j$ be a partition over Θ^n that is generated by $\{\mathcal{F}_k^j\}_{k \leq n}$. We note that $\bar{\mathcal{F}}_n^j$ is clearly finer than \mathcal{F}_n^j . This would represent the information of Player j at stage n in the auxiliary game $\Gamma^\varepsilon(W)$. Let $\Sigma_n \subseteq \Sigma$ be the set of strategies σ of Player 1 such that at each odd stage $k \leq n$, σ_k is measurable with respect to $\bar{\mathcal{F}}_k^1$. (In stages $k > n$, σ_k is not restricted; that is, it can be any measurable mapping from Θ^k to $\Delta(A)$.) Similarly let Υ_n be the set of such restricted behavioral strategies of Player 2. We shall show first that we can find a sequence of partitions that satisfy the following condition for each $n \in \mathbb{N}$:

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

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

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

We construct the partition inductively as follows. Assume that $\{\mathcal{F}_k^j\}_{k < n}$ has been defined such that condition (C_k) holds for every $k < n$. W.l.o.g., assume that n is odd so that player 1 is the active player. Let $V = \{\eta^1(h) \mid h \in H_n^\diamond\} \subseteq \Delta(\Theta^n)$ be the finite set of measures induced on Player 1's first n signals by the histories through the n -th state.

We use Lemma 6.1 to construct a finite partition \mathcal{F}_n^1 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,

$$\forall \nu \in V \quad \|(\sigma_n)_*(\nu) - (\sigma'_n)_*(\nu)\| < \frac{\epsilon}{2^n}. \quad (6.2)$$

For player 2 let \mathcal{F}_n^2 be the null partition over Θ^n .

We claim that the strategy spaces³⁶ Σ_n, Υ_n satisfy condition (C_n) . To see this let $\sigma \in \Sigma$ be 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

$$\|P_{\sigma, \tau} - P_{\tilde{\sigma}, \tau}\| < (1 - \frac{1}{2^{n-1}})\epsilon, \quad (6.3)$$

Define then $\sigma' = (\tilde{\sigma}_1, \dots, \tilde{\sigma}_{n-1}, \sigma'_n, \sigma_{n+1}, \sigma_{n+2}, \dots)$. For each $h_n \in H_n^\diamond$, the measure $\eta^1(h_n)$ is in V , and since for all $h_{n+1} \in H_{n+1}^\diamond$, $P_{\sigma, \tau}(\cdot \mid h_{n+1}) = P_{\sigma', \tau}(\cdot \mid h_{n+1})$,

$$\|P_{\sigma', \tau}(\cdot \mid h_n) - P_{\sigma, \tau}(\cdot \mid h_n)\| = \|(\sigma_n)_*(\eta^1(h_n)) - (\sigma'_n)_*(\eta^1(h_n))\| < \frac{\epsilon}{2^n} \quad (6.4)$$

Therefore, since $P_{\sigma', \tau}$ and $P_{\tilde{\sigma}, \tau}$ induce the same distributions on H_n^\diamond , (6.3) and (6.4) yield that,³⁷

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

The above inductive construction yields two sequences of finite measurable partitions $\{\mathcal{F}_n^j\}_n$ of $\{\Theta^n\}_n$ one for each player $j \in \{1, 2\}$ 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 $\bar{\mathcal{F}}_n$ measurable for every n , and for every strategy τ of player 2 in $\Gamma(W)$,

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

and similarly for Player 2.

We shall make the following two additional refinements on the partitions $\{\mathcal{F}_n^j\}_n$ of every player $j = 1, 2$:

1. For every every k the generated partition $\bar{\mathcal{F}}_k^j$ of Θ^k preserves perfect recall for Player j . That is, the generated partition $\bar{\mathcal{F}}_k^j$ is measurable with respect to the sequence of perfectly informative signals Player j receives about the actions taken by himself up to stage k .

³⁶Note that in this case, where n is odd, $\Upsilon_n = \Upsilon_{n-1}$.

³⁷This can be seen via a generalisation of Lemma 5.2, which we do not state explicitly.

2. For every two distinct infinite vectors of signals $\bar{\theta}, \bar{\theta}' \in \Theta^\infty$ there exists a k such that their projections $\bar{\theta}_k, \bar{\theta}'_k$ to $\bar{\Theta}_k = \Theta^k$ lie in two distinct partition element of $\bar{\mathcal{F}}_k^j$.

The two conditions can be achieved by simply refining the partition elements in the construction stage; the first, in particular, follows easily if the structure discussed early in (PR) is assumed.

Let $\Gamma^\varepsilon(W)$ be the game in which, for each $n \in \mathbb{N}$, each player $j = [n]$ uses only $\bar{\mathcal{F}}_n^j$ -measurable strategies. The game $\Gamma^\varepsilon(W)$ is of course equivalent to a game where at stage n the set of signals of Player j is Θ_n^j - which is the collection of atoms of $\bar{\mathcal{F}}_n^j$ - is finite. To see this note that the game $\Gamma^\varepsilon(W)$ is equivalent to the game where at every stage n each player j is informed of his partition element in $\bar{\mathcal{F}}_n^j$. By condition 1 the game $\Gamma^\varepsilon(W)$ clearly satisfies the perfect recall assumption.

In addition we claim that $\Gamma^\varepsilon(W)$ satisfies SEPM: To see this let \hat{q}, \hat{q}' be a pair of beliefs on Nature (see Definition 2.11), and let $(\sigma, \tau), (\sigma', \tau')$ be a pair of strategy profiles of $\Gamma^\varepsilon(W)$ such that $\pi_*^H(P_{\hat{q}, \sigma, \tau}) \perp \pi_*^H(P_{\hat{q}', \sigma', \tau'})$. Since both (σ, τ) and (σ', τ') are also strategies profiles of $\Gamma(W)$ it holds that $\pi_*^j(P_{\hat{q}, \sigma, \tau}) \perp \pi_*^j(P_{\hat{q}', \sigma', \tau'})$ for $j = 1, 2$ (see Definition 2.12) where $\pi^j : \tilde{H}_\infty \rightarrow \Theta^\infty$ are the projections in the original game $\Gamma(W)$. By the above Condition 2, for every $j \in \{1, 2\}$ there exists a Borel bijection ψ between Θ^∞ and $\prod_n \mathcal{F}_n^j$. Hence if we let $\bar{\pi}^j : \tilde{H}_\infty = \prod_n (S_n \times A_n \times \Theta^2) \rightarrow \prod_n \mathcal{F}_n^j$ be the projection in the auxiliary game $\Gamma^\varepsilon(W)$, $\bar{\pi}_*^j = \psi_* \circ \pi_*^j$ and then it must hold for every player j that $\bar{\pi}_*^j(P_{\hat{q}, \sigma, \tau}) \perp \bar{\pi}_*^j(P_{\hat{q}', \sigma', \tau'})$.³⁸ Therefore the game $\Gamma^\varepsilon(W)$ is determined, by the version of Theorem 2.13 for finite signal spaces.

We claim that

$$\underline{val}\Gamma(W) \geq \underline{val}(\Gamma^\varepsilon(W)) - 2\varepsilon.$$

To see this note that by construction every strategy σ' in $\Gamma^\varepsilon(W)$ guarantees the same payoff up to 2ε in $\Gamma(W)$.³⁹ Similarly,

$$\overline{val}\Gamma(W) \leq \overline{val}(\Gamma^\varepsilon(W)) + 2\varepsilon.$$

Hence,

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

Since ε is arbitrary it follows that $\Gamma(W)$ 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 $j \in \{1, 2\}$. There exists for each $n \in \mathbb{N}$ and for each $h \in H_n^\circ$ satisfying $P_{\sigma, \tau}(h) >$*

³⁸Since ψ is injective, ψ_* preserves mutual singularity.

³⁹Recall that the payoffs are ± 1 .

0 a strategy profile σ_h, τ_h and a belief of Nature \tilde{q}_h such that $P_{\sigma_h, \tau_h, \tilde{q}_h}(\cdot) = P_{\sigma, \tau}(\cdot | h)$, and hence for $h \neq h' \in H_n^\circ$, $\pi_*^j(P_{\sigma, \tau}(\cdot | h)) \perp \pi_*^j(P_{\sigma, \tau}(\cdot | h'))$, where π_*^j denotes the projection to Player j 's signal space as in Definition 2.6.

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,

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

where π_*^H is the projection to H_∞ , as in Definition 2.6. In other words, $\pi_*^j(P_{\sigma, \tau}(\cdot | h))$ 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 $(A_\alpha) = (h)_{h \in H_n^\circ}$):

Lemma 7.2. *Let X, Y be standard Borel spaces, let $\mu \in \Delta(X)$ and let $A_1, \dots, A_n \subseteq X$ 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 $\eta(\mu_i) \perp \eta(\mu_j)$ for $i \neq j$. Then for any filtration $(\mathcal{F}_t)_t$ of Y generating the Borel σ -algebra⁴⁰ \mathcal{F} and each j , $P_\mu(A_j | \mathcal{F}_t) \xrightarrow[t \rightarrow \infty]{} 1$ $\eta(\mu_j)$ -a.s..*

Proof. Since $(P_\mu(A_j | \mathcal{F}_t))_t$ is a martingale, by the martingale convergence theorem, it suffices to show that $P_\mu(A_j | \mathcal{F}) = 1$ $\eta(\mu_j)$ -a.s.. 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 $P_\mu(A_j | \mathcal{F}) = 1$ $\eta(\mu)$ -a.s. in B_j ; by (7.1), it suffices to show that $P_\mu(B_j | \mathcal{F}) = 1$ $\eta(\mu)$ -a.s. in B_j , which is immediate. \square

8 Appendix D: Proof of Proposition 4.1

Fix $\varepsilon > 0$, and $0 < \alpha < 1$, and on \bar{I}_∞ , define the following measure P : We can write $\bar{I}_\infty = \prod_{n \in \mathbb{N}} \{1, 0\}^{C^n}$ and hence define $P = \otimes_{n \in \mathbb{N}} \otimes_{(c_1, \dots, c_n) \in C^n} (\alpha, 1 - \alpha)$.

Proposition 8.1. *For any normal strategy σ , $\|P_\sigma - P\| \rightarrow 1$ as $\alpha \rightarrow 1$, where P_σ denotes the marginal on \bar{I}_∞ .⁴¹*

This proposition clearly implies Proposition 4.1.

⁴⁰A σ -algebra \mathcal{F} on Y implicitly introduces the σ -algebra $\{\emptyset, X\} \times \mathcal{F}$ on $X \times Y$.

⁴¹In one calculation below, P_σ will refer to the product measure on $D^{\mathbb{N}} \times \bar{I}_\infty$.

Proof. By the bounded convergence theorem, it suffices to consider normal σ which makes pure choices in D ; let $\bar{d}_\infty = (d_1, d_2, d_3, \dots)$ be the sequence in $D^\mathbb{N}$ that σ chooses. For $k \in \mathbb{N}$, let $B_n^d = \{(d_1, b_1, \dots, d_n, b_n) \in \prod_{k \leq n} (d_k \times B)\}$ - i.e., histories of length n which agree with \bar{d}_∞ in their choices in D . Let $\bar{I}_{n+1} = \prod_{k \leq n} \{0, 1\}^{C^k}$. By Lemma 5.1, in order to prove the proposition it is sufficient to show that for every δ there exists α_0 , such that if $\alpha > \alpha_0$, then for all $n \in \mathbb{N}$, $\|P_\sigma^n - P^n\| < \delta$, where P_σ^n, P^n are the marginals of P_σ, P on \bar{I}_{n+1} .

Let $\bar{i}_{n+1} \in \bar{I}_{n+1}$. Let

$$K(\bar{i}_{n+1}) = |\cup_{k \leq n} \{(c_1, \dots, c_k) \in C^k \mid i_{k+1}(c_1, \dots, c_k) = 1\}|$$

and denote $J_n = \sum_{k=1}^n |C^k|$. Now,

$$P(\bar{i}_{n+1}) = \alpha^{K(\bar{i}_{n+1})} (1 - \alpha)^{J_n - K(\bar{i}_{n+1})}$$

while (since conditional on $h \in B_n^d$, the coordinates of the signal are chosen independently, with the n coordinates along the true path being given 1 for certain, and the others randomised independently with probabilities $(\alpha, 1 - \alpha)$),

$$\begin{aligned} P_\sigma^n(\bar{i}_{n+1}) &= \sum_{h \in B_n^d} P_\sigma(\bar{i}_{n+1} \mid h) P_\sigma(h) \\ &= \sum_{h=(d_1, b_1, \dots, d_n, b_n) \in B_n^d} \frac{1}{|B_n^d|} \alpha^{K(\bar{i}_{n+1}) - n} (1 - \alpha)^{J_n - K(\bar{i}_{n+1})} \cdot \prod_{j=1}^n i_{j+1}(d_1, b_1, \dots, d_j, b_j) \\ &= \alpha^{K(\bar{i}_{n+1})} (1 - \alpha)^{J_n - K(\bar{i}_{n+1})} [\alpha^{-n} \cdot \frac{1}{|B_n^d|} \sum_{h=(d_1, b_1, \dots, d_n, b_n) \in B_n^d} \prod_{j=1}^n i_{j+1}(d_1, b_1, \dots, d_j, b_j)] \end{aligned}$$

Hence, since $|B_n^d| = 2^n$,

$$\begin{aligned} \frac{P_\sigma^n(\bar{i}_{n+1})}{P^n(\bar{i}_{n+1})} &= \alpha^{-n} \cdot \frac{1}{2^n} \sum_{h=(d_1, b_1, \dots, d_n, b_n) \in B_n^d} \prod_{j=1}^n i_{j+1}(d_1, b_1, \dots, d_j, b_j) \\ &= \frac{\#\{(d_1, b_1, \dots, d_n, b_n) \in B_n^d \mid \forall j \leq n, i_{j+1}(d_1, b_1, \dots, d_j, b_j) = 1\}}{(2\alpha)^n} \end{aligned} \tag{8.1}$$

Since d_1, d_2, \dots are fixed, the proof of the proposition will be complete if we show that:

Lemma 8.2. *Let $Q = \otimes_{n \in \mathbb{N}} (\alpha, 1 - \alpha)^{2^n}$ on $\Theta = \Theta_1 \times \Theta_2 \times \dots = \prod_{n \in \mathbb{N}} \{1, 0\}^{2^n}$, and let*

$$Y_n(\theta_1, \theta_2, \dots) = \frac{\#\{(x_1, \dots, x_n) \in 2^n \mid \forall k \leq n, \theta_k(x_1, \dots, x_k) = 1\}}{(2\alpha)^n}$$

Then for any $\varepsilon > 0$, there is $\alpha_0 < 1$ such that if $1 > \alpha > \alpha_0$, for each $n \in \mathbb{N}$, there is a set $A_n \subseteq \bar{I}_{n+1}$ such that $Q(A_n) > 1 - \varepsilon$ and $|Y_n - 1| < \varepsilon$ in A_n .

The lemma can also be stated: In a branching process in which each node splits into two branches, each of which 'lives' with probability α , if α is close enough to 1, then for any $n \in \mathbb{N}$, it holds on a set of arbitrarily large (uniformly in n) measure the set of surviving branches at stage n is close to an $(2\alpha)^n$ -fraction of the total number of 2^n possible branches that could have existed.

Proof. Let $\gamma = 2\alpha$; it follows from [Karlin and Taylor, Sec 8.2] that $E(Y_n) = 1$, and if $\gamma \neq 1$,

$$\text{Var}(Y_n) = \text{Var}(Y_1)\gamma^{n-1} \frac{\gamma^n - 1}{\gamma^{2n}(\gamma - 1)} \leq \text{Var}(Y_1) \frac{1}{\gamma(\gamma - 1)} = \frac{1 - \alpha}{2\alpha - 1}.$$

Hence $\sup_n \text{Var}(Y_n)$ approaches zero as α approaches one. The lemma follows directly from Chebyshev's inequality. \square

\square

References

- [Arieli and Levy (2011)] Arieli, D. and Levy, I. (2011) Infinite Sequential Games with Perfect but Incomplete Information, *Int. J. Game Theory*, 40, 207–213.
- [Arieli and Levy (In Preparation)] Arieli, D. and Levy, I. (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, *Appliciones Mathematicae*, Vol. X., pp. 99–101, Hugo Steinhaus Jubilee Volume.
- [Cripps et al. (2008)] CRIPPS, M.W., ELY, J.C., MAILATH, G.J., AND SAMUELSON, L. (2008), Common Learning, *Econometrica*, 76, 909–933.
- [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
- [Kakutani (1948)] KAKUTANI, S. (1948), On Equivalence of Infinite Product Measures, *Annals of Mathematics*, 49, 214–224.
- [Karlin and Taylor] KARLIN, S. AND TAYLOR, H.M. (1975), *A First Course in Stochastic Processes*, 2nd Ed., Academic Press, Inc..

- [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.
- [Rudin (1986)] RUDIN, W. (1986), *Real and Complex Analysis*, 3rd Ed., International Series in Pure and Applied Mathematics, McGraw-Hill Science.
- [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.