

Approximations of expectations under infinite product measures*

Galit Ashkenazi-Golan[†] János Flesch[‡] Arkadi Predtetchinski[§]
Eilon Solan[¶]

January 7, 2025

Abstract

We are given a bounded Borel-measurable real-valued function on a product of countably many Polish spaces, and a product probability measure. We are interested in points in the product space that can be used to approximate the expected value of this function. We define two notions. A point is called a *weak ϵ -approximation*, where $\epsilon \geq 0$, if the Dirac measure on this point, except in finitely many coordinates where another measure can be taken, gives an expected value that is ϵ -close to the original expected value. A point is called a *strong ϵ -approximation* if the same holds under the restriction that in those finitely many coordinates the measure is equal to the original one. We prove that both the set of weak 0-approximation points and the set of strong ϵ -approximation points, for any $\epsilon > 0$, have measure 1 under the original measure. Finally, we provide two applications: (i) in Game Theory on the minmax guarantee levels of the players in games with infinitely many players, and (ii) in Decision Theory on the set of feasible expected payoffs in infinite duration problems.

Keywords: Infinite product space, product measure, universally measurable set, strategic-form game, minmax value.

MSC 2020 Classification: 28A35, 28C20, 91A07.

*Solan acknowledges the support of the Israel Science Foundation, grant #211/22. We are grateful to John (Yehuda) Levy for a stimulating discussion.

[†]London School of Economics and Political Science, Houghton Street, London WC2A 2AE, UK, E-mail: galit.ashkenazi@gmail.com.

[‡]Department of Quantitative Economics, Maastricht University, P.O.Box 616, 6200 MD, The Netherlands. E-mail: j.flesch@maastrichtuniversity.nl.

[§]Department of Economics, Maastricht University, P.O.Box 616, 6200 MD, The Netherlands. E-mail: a.predtetchinski@maastrichtuniversity.nl.

[¶]School of Mathematical Sciences, Tel-Aviv University, Tel-Aviv, Israel, 6997800, E-mail: eilons@tauex.tau.ac.il.

1 Introduction

Approximations of expected payoffs. Suppose that $\sigma = \otimes_{i \in \mathbb{N}} \sigma_i$ is a product probability measure on a product space $X = \times_{i \in \mathbb{N}} X_i$, where each X_i is a non-empty Polish space, and $f : X \rightarrow \mathbb{R}$ is a bounded and Borel-measurable function. In this paper, we ask whether $\mathbb{E}_\sigma[f]$ can be approximated by an expectation $\mathbb{E}_\tau[f]$, for some “simple” product measure τ .

If f is a tail function, then there necessarily exists $x \in X$ such that $f(x) = \mathbb{E}_\sigma[f]$; in fact, almost every (under σ) point of X has the stated property. In general, no such point x need to exist.

We are interested in two types of approximating measures. We say that $x = (x_i)_{i \in \mathbb{N}} \in X$ is a *weak ϵ -approximation* of the expectation of f under σ if there is some product probability measure $\tau = \otimes_{i \in \mathbb{N}} \tau_i$, such that (i) $|\mathbb{E}_\tau[f] - \mathbb{E}_\sigma[f]| \leq \epsilon$, and (ii) τ_i is the Dirac measure on x_i for all but finitely many i ’s. We will say that $x \in X$ is a *strong ϵ -approximation* of the expectation of f under σ if the same holds, where in addition (iii) τ_i is either σ_i or the Dirac measure on x_i , for every $i \in \mathbb{N}$.

Our main results are that (i) the set of weak 0-approximations have σ -measure 1 (cf. Theorem 4), and (ii) the set of strong ϵ -approximations have σ -measure 1 for every $\epsilon > 0$ (cf. Theorem 6). As an example demonstrates (cf. Example 7), a strong 0-approximation does not always exist.

Application in Game Theory. The first motivation for our study comes from Game Theory, and more specifically, from the study of strategic-form games with countably many players. Games with countably many players have been studied, e.g., by Peleg [8], Voorneveld [14], Rachmilevitch [9], and Ashkenazi-Golan, Flesch, Predtetchinski and Solan [2]. These studies have mainly revolved around the question of (non)existence of Nash (ϵ -)equilibrium. In contrast, here we focus on the notion of the minmax value.

The *minmax value* of a player is one of the central concepts in game theory. It is defined as the maximal payoff that the player can guarantee to obtain, when all other players try to decrease her payoff. The underlying assumption on the behavior of the punishing players is that they cannot correlate their lotteries, and thus their joint strategy is a product measure.

Ashkenazi-Golan, Flesch, Predtetchinski and Solan [2] introduced a notion of the *finitistic minmax value*. This is a version of the minmax value where only finitely many of the punishing players are allowed to randomize; the others are required to choose a Dirac measure. It is clear that the finitistic minmax value is not lower than the minmax value. Whether the two values are in fact equal is an open problem. Here we answer this question in the affirmative (cf. Theorem 8). This is achieved using the result (ii) on strong ϵ -approximations.

Application in Decision Theory. The second motivation comes from decision theory. In dynamic decision problems with infinite horizon, *Markov strategies* play an important role (e.g., Hill [5], Kiefer, Mayr, Shirmohammadi, Totzke [6], Sudderth [13]). These are strategies in which, at each stage, the choice of action may depend on the current stage

and the current state, but not on past events. Such strategies are desirable because of their simplicity, as the decision maker's behavior depends only on the current payoff relevant parameters. A strategy is called *eventually pure* if after some stage it does not use randomization (cf. Flesch, Herings, Maes, Predtetchinski [4]).

Using our result (i) on weak 0-approximations, we show that in a class of decision problems with infinite horizon, the set of expected payoffs that the decision maker can obtain by Markov strategies is equal to the set of expected payoffs under eventually pure Markov strategies. In other words, using Markov strategies that involve a randomization at infinitely many stages does not generate additional expected payoffs.

Structure of the paper. Definitions appear in Section 2, the main results and their proofs are provided in Section 3, and the applications to Game Theory and Decision Theory are discussed in Section 4. A few concluding remarks are given, including an open problem, in Section 5.

2 Approximations of expectations

Let $\mathbb{N} = \{1, 2, \dots\}$. For each $i \in \mathbb{N}$, let X_i be a non-empty Polish space. Let $X = \times_{i \in \mathbb{N}} X_i$ denote the product space, which we endow with the product topology.

For each $i \in \mathbb{N}$, let Σ_i be the set of Borel probability measures on X_i , and let Σ be the set of all product Borel probability measures on X , i.e., the measures $\sigma = \otimes_{i \in \mathbb{N}} \sigma_i$ where $\sigma_i \in \Sigma_i$ for each $i \in \mathbb{N}$. With a slight abuse of notation, we write x_i to denote both a point of X_i and the Dirac measure concentrated on x_i , and we use the same symbol σ to denote both a measure in Σ and its unique extension to the sigma-algebra of all universally measurable sets of X .

We define two notions for the approximations of expectations under the measures in Σ .

Definition 1 Let $\epsilon \geq 0$. For a product measure $\sigma \in \Sigma$ and a bounded Borel-measurable function $f : X \rightarrow \mathbb{R}$, a point $x \in X$ is said to be a weak ϵ -approximation of the expectation of f under σ if there is $n \in \mathbb{N}$ and a Borel probability measure τ_i on X_i for each $i < n$, such that:

$$|\mathbb{E}_{\tau_1 \otimes \dots \otimes \tau_{n-1} \otimes x_n \otimes x_{n+1} \otimes \dots}[f] - \mathbb{E}_\sigma[f]| \leq \epsilon. \quad (1)$$

Intuitively, Eq. (1) means that the expectation changes by at most ϵ if we replace σ by the measure that is the Dirac measure on x_i in each coordinate $i \geq n$ and τ_i in each coordinate $i < n$.

Definition 2 Let $\epsilon \geq 0$. For a product measure $\sigma \in \Sigma$ and a bounded Borel-measurable function $f : X \rightarrow \mathbb{R}$, a point $x \in X$ is said to be a strong ϵ -approximation of the expectation of f under σ if there is $n \in \mathbb{N}$ such that:

$$|\mathbb{E}_{\sigma_1 \otimes \dots \otimes \sigma_{n-1} \otimes x_n \otimes x_{n+1} \otimes \dots}[f] - \mathbb{E}_\sigma[f]| \leq \epsilon. \quad (2)$$

Thus, Eq. (2) means that the expectation changes by at most ϵ if we replace σ by the Dirac measure on x_i in each coordinate $i \geq n$. Obviously, if $x \in X$ is a strong ϵ -approximation of the expectation of f under σ , then x is also a weak ϵ -approximation.

Let $W_\epsilon(\sigma, f) \subseteq X$ and $S_\epsilon(\sigma, f) \subseteq X$ denote the sets of weak and respectively strong ϵ -approximations of the expectation of f under σ . Thus, $W_\epsilon(\sigma, f) \subseteq S_\epsilon(\sigma, f)$. It follows from the following lemma that both $W_\epsilon(\sigma, f)$ and $S_\epsilon(\sigma, f)$ are universally measurable, and in particular, both sets are measurable with respect to the measure σ .

Lemma 3 *For each $\epsilon \geq 0$, each product measure $\sigma \in \Sigma$, and each bounded Borel-measurable function $f : X \rightarrow \mathbb{R}$, the set $W_\epsilon(\sigma, f) \subseteq X$ is analytic and the set $S_\epsilon(\sigma, f) \subseteq X$ is Borel.*

Proof: For each $n \in \mathbb{N}$, we equip Σ_n with the topology of weak convergence, and $\Sigma^n := \times_{i=1}^n \Sigma_i$ with the product topology. Note that Σ^n is a Polish space. For each $n \in \mathbb{N}$, the function $X \times \Sigma^n \rightarrow \mathbb{R}$ given by $(x, \tau_1, \dots, \tau_n) \mapsto \mathbb{E}_{\tau_1 \otimes \dots \otimes \tau_{n-1} \otimes \tau_n \otimes x_{n+1} \otimes x_{n+2} \otimes \dots} [f]$ is Borel-measurable (this follows by applying Aliprantis and Border [1, Theorem 15.8], Ressel [11, Theorem 2(i)], and Aliprantis and Border [1, Theorem 15.13]). Hence the set $B^n \subseteq X \times \Sigma^n$ consisting of the points $(x, \tau_1, \dots, \tau_n)$ satisfying Eq. (1) is a Borel set. Let A^n be the projection of B^n onto the first coordinate. Then A^n is analytic. Since $W_\epsilon(\sigma, f) = \cup_{n \in \mathbb{N}} A^n$, the first result follows.

The fact that $S_\epsilon(\sigma, f)$ is a Borel set follows since the map $X \rightarrow \mathbb{R}$ given by $x \mapsto \mathbb{E}_{\sigma_1 \otimes \dots \otimes \sigma_{n-1} \otimes x_n \otimes x_{n+1} \otimes \dots} [f]$ is Borel-measurable for each $n \in \mathbb{N}$. \square

3 Results

Theorem 4 *For each product measure $\sigma \in \Sigma$ and each bounded Borel-measurable function $f : X \rightarrow \mathbb{R}$, the set $W_0(\sigma, f)$ has σ -measure 1.*

Proof: Fix a product measure $\sigma \in \Sigma$ and a bounded Borel-measurable function $f : X \rightarrow \mathbb{R}$.

Let \sim be the following tail equivalence relation: for $x, y \in X$ let $x \sim y$ if there is some $n \in \mathbb{N}$ such that $x_i = y_i$ for all $i \geq n$. Let \mathcal{Z} denote the collection of all equivalence classes of \sim . For $Z \in \mathcal{Z}$, let I_Z be the convex hull of $f(Z) = \{f(x) : x \in Z\}$; in particular, I_Z is an interval, with or without its endpoints.

Let $r = \mathbb{E}_\sigma[f]$. Define

$$\begin{aligned} Z_- &= \bigcup \{Z \in \mathcal{Z} : I_Z \subseteq (-\infty, r)\}, \\ Z_0 &= \bigcup \{Z \in \mathcal{Z} : I_Z \ni r\}, \\ Z_+ &= \bigcup \{Z \in \mathcal{Z} : I_Z \subseteq (r, +\infty)\}. \end{aligned}$$

The statement of Theorem 4 follows from Claims 1 and 2 below.

CLAIM 0: Z_- and Z_+ are co-analytic subsets of X .

It follows from Claim 0 that Z_0 is analytic, and hence Z_- , Z_0 , and Z_+ are measurable with respect to σ .

PROOF OF CLAIM 0: Let us write \sim to denote the set $\{(x, y) \in X \times X : x \sim y\}$. Then \sim is a Borel subset of $X \times X$, as can be seen by rewriting it as

$$\bigcup_{n \in \mathbb{N}} \bigcap_{i \geq n} \{(x, y) \in X \times X : x_i = y_i\}.$$

The set $X \setminus Z_-$ is a projection of the Borel set

$$\sim \bigcap \{(x, y) \in X \times X : f(y) \geq r\}$$

onto its first coordinate, and therefore analytic. The argument that $X \setminus Z_+$ is analytic is similar.

CLAIM 1: $Z_0 \subseteq W_0(\sigma, f)$.

PROOF OF CLAIM 1: Take any $Z \in \mathcal{L}$ with $r \in I_Z$. We show that $Z \subseteq W_0(\sigma, f)$.

As $r \in I_Z$, there exist points $x, y \in Z$ such that $f(x) \leq r \leq f(y)$. Let $n \in \mathbb{N}$ be such that $x_i = y_i$ for each $i > n$.

For each $k = 1, \dots, n+1$, let $z(k) = (y_1, \dots, y_{k-1}, x_k, x_{k+1}, \dots)$. In particular, $z(1) = x$ and $z(n+1) = y$, and $z(k) \in Z$ for each $k = 1, \dots, n+1$. Because $z(1) = x$ and $z(n+1) = y$, the interval $[f(x), f(y)]$ is covered by the union of the intervals $[f(z(k)), f(z(k+1))]$ with $k = 1, \dots, n$. Consequently, there is $k \in \{1, \dots, n\}$ such that $r \in [f(z(k)), f(z(k+1))]$. Note that $z(k)$ and $z(k+1)$ only differ in coordinate k .

Let $\alpha \in [0, 1]$ be such that $\alpha \cdot f(z(k)) + (1 - \alpha) \cdot f(z(k+1)) = r$. Define the measure τ_k on X_k as follows: τ_k places probability α on x_k and probability $1 - \alpha$ on y_k . Then,

$$\mathbb{E}_{y_1 \otimes \dots \otimes y_{k-1} \otimes \tau_k \otimes x_{k+1} \otimes x_{k+2} \otimes \dots} [f] = \alpha \cdot f(z(k)) + (1 - \alpha) f(z(k+1)) = r = \mathbb{E}_\sigma [f].$$

This means that $z(k) = (y_1, \dots, y_{k-1}, x_k, x_{k+1}, \dots)$ is a weak 0-approximation of the expectation of f under σ , so that $z(k) \in W_0(\sigma, f)$. Hence, so is each point in the equivalence class Z of $z(k)$.

CLAIM 2: Z_0 has σ -measure 1.

PROOF OF CLAIM 2: The sets Z_- , Z_0 , and Z_+ are tail. Since σ is a product measure, the σ -measure of each set is either 0 or 1 (see Lemma 10 in the Appendix). Since Z_- , Z_0 , and Z_+ partition X , exactly one of them has σ -measure 1. It cannot be the case that $\sigma(Z_-) = 1$, since then, as $Z_- \subseteq f^{-1}(-\infty, r)$, we would obtain $\mathbb{E}_\sigma [f] < r$, a contradiction. One similarly excludes the possibility that $\sigma(Z_+) = 1$. Hence $\sigma(Z_0) = 1$. \square

Remark 5 The proof of Claim 1 in the proof of Theorem 4 reveals the following statement: For each product measure $\sigma \in \Sigma$ and each bounded Borel-measurable function

$f : X \rightarrow \mathbb{R}$, there exists a weak 0-approximation $x \in W_0(\sigma, f)$ with the following property: there is a coordinate n and a Borel probability measure τ_n on X_n such that

$$\mathbb{E}_{x_1 \otimes \dots \otimes x_{n-1} \otimes \tau_n \otimes x_{n+1} \otimes x_{n+2} \otimes \dots}[f] = \mathbb{E}_\sigma[f].$$

That is, we only need to replace x in a single coordinate by a Borel probability measure.

Theorem 6 *For each product measure $\sigma \in \Sigma$, each bounded Borel-measurable function $f : X \rightarrow \mathbb{R}$, and each $\epsilon > 0$, the set $S_\epsilon(\sigma, f)$ has σ -measure 1.*

Proof: Let \mathcal{G}_n be the sigma-algebra on X generated by the coordinate functions $x_k, k \geq n$. Let $\mathcal{G} = \bigcap_{n \in \mathbb{N}} \mathcal{G}_n$. This is the tail sigma-algebra on X . Consider $g_n = \mathbb{E}_\sigma[f | \mathcal{G}_n]$.

The process $\{g_n\}_{n \in \mathbb{N}}$ is a reverse martingale with respect to the sequence $\{\mathcal{G}_n\}_{n \in \mathbb{N}}$ and the measure σ . That is, $\mathbb{E}_\sigma[g_n | \mathcal{G}_{n+1}] = g_{n+1}$. By Durrett [3, Theorems 5.6.1 and 5.6.2], the sequence g_n converges to a \mathcal{G} -measurable limit g_∞ , almost surely with respect to σ , and $\mathbb{E}_\sigma[g_\infty] = \mathbb{E}_\sigma[g_0] = \mathbb{E}_\sigma[f]$. Since σ is a product measure, Kolmogorov's 0-1 law implies that g_∞ is constant σ -almost surely. Hence σ -almost surely, g_n converges to $\mathbb{E}_\sigma[f]$. This means that σ -almost surely, Eq. (2) holds for large $n \in \mathbb{N}$, as desired. \square

As the following example shows, there are product measures $\sigma \in \Sigma$ and bounded Borel-measurable functions f such that there is no strong 0-approximation of the expectation of f under σ , i.e., $S_0(\sigma, f) = \emptyset$.

Example 7 Let $X_i = \{0, 1\}$ for each $i \in \mathbb{N}$, and let $f(x)$ be 1 if $x_i = 1$ for each $i \in \mathbb{N}$ and 0 otherwise. We identify $\sigma_i \in \Sigma_i$ with $\sigma_i(\{1\})$, the probability of 1. Take any $\sigma \in \Sigma$ such that (a) $0 < \mathbb{E}_\sigma[f]$ and (b) there are infinitely many $i \in \mathbb{N}$ such that $\sigma_i < 1$ (for example $\sigma_i = 1 - 2^{-i}$ for each $i \in \mathbb{N}$). Then, $S_0(\sigma, f) = \emptyset$. Indeed, consider an $x \in X$ and an $n \in \mathbb{N}$. If $x_i = 0$ for some $i \geq n$, then $\mathbb{E}_{\sigma_1 \otimes \dots \otimes \sigma_{n-1} \otimes x_n \otimes x_{n+1} \otimes \dots}[f] = 0$. If, on the other hand, $x_i = 1$ for each $i \geq n$, then

$$\mathbb{E}_\sigma[f] = \prod_{i=1}^{\infty} \sigma_i < \prod_{i=1}^{n-1} \sigma_i = \mathbb{E}_{\sigma_1 \otimes \dots \otimes \sigma_{n-1} \otimes x_n \otimes x_{n+1} \otimes \dots}[f].$$

4 Applications in Decision Theory and Game Theory

In this section, we provide two applications of our results: one in Game Theory and one in Decision Theory.

Application in Game Theory. We apply Theorem 6 on strong ϵ -approximations of expectations to prove a result on the minmax values of the players in strategic-form games with an infinite set of players.

We consider a strategic-form game in which the set of players is $\{0\} \cup \mathbb{N}$. Player 0's action set is a non-empty finite set A , whereas the action set of each player $i \in \mathbb{N}$

is a non-empty Polish space X_i . As before, $X = \times_{i \in \mathbb{N}} X_i$. Player 0's payoff function is $u : A \times X \rightarrow \mathbb{R}$, a bounded Borel-measurable function. As before, Σ denotes the set of product probability measures on X , which is the set of strategy profiles of player 0's opponents. Let $\Sigma^* \subseteq \Sigma$ denote the set of strategy profiles $\sigma = \otimes_{i \in \mathbb{N}} \sigma_i$ of player 0's opponents such that σ_i is pure (i.e., a Dirac measure on X_i) for all but finitely many $i \in \mathbb{N}$.

The *minmax value* and the *finitistic minmax value* of player 0 are defined respectively as

$$v = \inf_{\sigma \in \Sigma} \max_{a \in A} \mathbb{E}_{a \otimes \sigma}[u],$$

$$v^* = \inf_{\sigma \in \Sigma^*} \max_{a \in A} \mathbb{E}_{a \otimes \sigma}[u].$$

Intuitively, the minmax value of player 0 is the maximal expected payoff that player 0 can guarantee to obtain, when all other players try to lower her payoff. The minmax value is a fundamental concept, and the expected payoff of a player in a Nash equilibrium can never be below her minmax value.

The finitistic minmax value is a variant where only finitely many opponents can randomize over their actions. As mentioned in the introduction, the concept of the finitistic minmax value played a crucial role in Ashkenazi-Golan, Flesch, Predtetchinski and Solan [2] in their study of games with infinitely many players.

It follows directly from the definitions that $v \leq v^*$. By using Theorem 6, it turns out, perhaps surprisingly, that in fact equality holds.

Theorem 8 $v = v^*$.

Proof: Since $\Sigma^* \subseteq \Sigma$, we always have $v \leq v^*$. To prove Theorem 8 we therefore need to show that $v^* \leq v$. Fix $\epsilon > 0$. We show that $v^* \leq v + 2\epsilon$.

For each $a \in A$, let $f_a = u(a, \cdot) : X \rightarrow \mathbb{R}$. Let $\sigma \in \Sigma$ be ϵ -minmax strategy profile of player 0's opponents; that is, a strategy profile such that $\mathbb{E}_{a \otimes \sigma}[u] = \mathbb{E}_{\sigma}[f_a] \leq v + \epsilon$ for each $a \in A$.

Take any $x \in \bigcap_{a \in A} S_{\epsilon}(f_a, \sigma)$, which is non-empty by Theorem 6 and since A is finite. Using once again the finiteness of A , there exists $n \in \mathbb{N}$ such that Eq. (2) holds for each of the functions f_a , and thus $\mathbb{E}_{\sigma_1 \otimes \dots \otimes \sigma_{n-1} \otimes x_n \otimes x_{n+1} \otimes \dots}[f_a] \leq \mathbb{E}_{\sigma}[f_a] + \epsilon$ for each $a \in A$. This implies that the strategy profile $\tau = \sigma_1 \otimes \dots \otimes \sigma_{n-1} \otimes x_n \otimes x_{n+1} \otimes \dots \in \Sigma^*$ of player 0's opponents is a 2ϵ -minmax strategy profile: $\mathbb{E}_{a \otimes \tau}[u] = \mathbb{E}_{\tau}[f_a] \leq \mathbb{E}_{\sigma}[f_a] + \epsilon \leq v + 2\epsilon$ for each $a \in A$. Therefore, $v^* \leq v + 2\epsilon$, as desired. \square

As the following example shows, the statement of Theorem 8 is not true if player 0's action set A is infinite.

Example 9 Suppose that $A = \mathbb{N} \times \{0, 1\}$ and $X_i = \{0, 1\}$ for each $i \in \mathbb{N}$. Define player 0's payoff function as follows: for $(n, j) \in A$ and $x \in X$, let $u((n, j), x)$ be 1 if $j = x_n$ and

0 otherwise. Intuitively, player 0's action (n, j) can be thought of naming an opponent n and an action j , and if opponent n 's action is indeed j , then player 0 obtains payoff 1.

On the one hand, $v = \frac{1}{2}$, which follows by noticing that if σ_i places equal probability on 0 and 1 for each $i \in \mathbb{N}$ then $\mathbb{E}_{(n,j) \otimes \sigma}[u] = \frac{1}{2}$ for each $(n, j) \in A$. On the other hand, $v^* = 1$, as in any strategy profile $\sigma \in \Sigma^*$ there is a player $n \in \mathbb{N}$ who places probability 1 on one of his actions $j \in \{0, 1\}$, and then $\mathbb{E}_{(n,j) \otimes \sigma}[u] = 1$.

Application in Decision Theory. We apply Theorem 4 to the following class of decision problems with infinite horizon: At each stage $t \in \mathbb{N}$, a decision maker chooses an action x_t from a non-empty Polish space X_t . If the sequence of chosen actions is $(x_1, x_2, \dots) \in X = \times_{t \in \mathbb{N}} X_t$, then her payoff is $u(x_1, x_2, \dots)$, where $u : X \rightarrow \mathbb{R}$ is a bounded Borel-measurable function.

A *Markov strategy* is an element of Σ , which is, as before, the set of product probability measures on X . Intuitively, a Markov strategy $\sigma = (\sigma_t)_{t \in \mathbb{N}} \in \Sigma$ chooses an action at each stage t only depending on the current stage t , and thus not depending on the past actions. Such a strategy is called *eventually pure*, if σ_t is a Dirac measure for all but finitely many stages $t \in \mathbb{N}$.

It follows from Theorem 4 that the set of expected payoffs that the decision maker can obtain by Markov strategies is equal to the set of expected payoffs under eventually pure Markov strategies. In other words, randomizing infinitely many times in a Markov strategy does not give additional expected payoffs to the decision maker.

5 Concluding Remarks

We do not know whether Theorem 4 could be deduced from Theorem 6 or vice versa. The two results appear to be logically unrelated, especially as the proofs use different approaches.

We conclude with an open problem: Can one generalize Theorem 4 to vector-valued (rather than just real-valued) functions? Consider a bounded Borel-measurable function $f : X \rightarrow \mathbb{R}^d$, and a product measure $\sigma \in \Sigma$. Is there a point $x \in X$, an $n \in \mathbb{N}$, and probability measures $\tau_i \in \Sigma_i$ for $i < n$, such that

$$\mathbb{E}_{\tau_1 \otimes \dots \otimes \tau_{n-1} \otimes x_n \otimes x_{n+1} \otimes \dots}[f_i] = \mathbb{E}_\sigma[f_i]$$

for each $i \in \{1, \dots, d\}$? Theorem 4 gives the affirmative answer when $d = 1$. The answer is also a "yes" when f is a tail function (i.e., when f is measurable with respect to the tail sigma-algebra on X). Beyond these special cases, the problem is open. In particular, it is also open if f is continuous.

Appendix

We have applied Kolmogorov's zero-one law to analytic and co-analytic sets. A similar result is stated (without proof) in Rao and Rao [10, Chapter 11, p. 265]. For the sake of

completeness, we state and prove a version of the result suitable for our purposes.

Lemma 10 *If U is a universally measurable tail subset of X and $\sigma \in \Sigma$, then the σ -measure of U can only have the values 0 and 1.*

Proof: We adapt the proof of the Kolmogorov's zero-one law in Shiryaev [12, p.381].

For $n \in \mathbb{N}$ let \mathcal{F}_n be the sigma-algebra on X generated by the first n coordinate functions x_1, \dots, x_n . The Borel sigma-algebra \mathcal{F} on X is the sigma-algebra generated by the set $\cup_{n \in \mathbb{N}} \mathcal{F}_n$. As U is universally measurable, there exists a Borel set B of X such that $\sigma(B \Delta U) = 0$, where Δ denotes the symmetric difference.

By Shiryaev [12, Problem 8, p.69], for every $n \in \mathbb{N}$ there exists a set $B_n \in \mathcal{F}_n$ such that $\sigma(B_n \Delta B) \rightarrow 0$. Since σ is a product measure and U is tail, B_n and U are independent. Therefore,

$$\begin{aligned} \sigma(B_n) &\rightarrow \sigma(B) = \sigma(U) \text{ and} \\ \sigma(B_n) \cdot \sigma(U) &= \sigma(B_n \cap U) = \sigma(B_n \cap B) \rightarrow \sigma(B) = \sigma(U). \end{aligned}$$

It follows that $\sigma(U) = \sigma(U) \cdot \sigma(U)$, implying the result. \square

References

- [1] CD Aliprantis and KC Border (2006): *Infinite Dimensional Analysis A Hitchhiker's Guide*. Springer Berlin, Heidelberg.
- [2] G Ashkenazi-Golan, J Flesch, A Predtetchinski, E Solan (2022): Existence of equilibria in repeated games with long-run payoffs. *Proceedings of the National Academy of Sciences*, 119(11), e2105867119.
- [3] R Durrett (2010): *Probability: Theory and Examples*. Fourth Edition. Cambridge University Press. Cambridge.
- [4] J Flesch, PJ Herings, J Maes, A Predtetchinski (2022): Individual upper semicontinuity and subgame perfect ϵ -equilibria in games with almost perfect information. *Economic Theory*, 73, 695–719.
- [5] TP Hill (1979): On the existence of good Markov strategies. *Transactions of the American Mathematical Society*, 247, 157–176.
- [6] S Kiefer, R Mayr, M Shirmohammadi, P Totzke (2019): Büchi objectives in countable MDPs. *ICALP 2019*: 119:1–14.
- [7] KR Parthasarathy (1967): *Probability measures on metric spaces*. Academic Press, New York.

- [8] B Peleg (1969): Equilibrium points for games with infinitely many players. *Journal of the London Mathematical Society*, 44, 292–294.
- [9] S Rachmilevitch [2020]: A note on discontinuity and approximate equilibria in games with infinitely many players. *Economics Letters* 193, 109–267.
- [10] KPSB Rao and MB Rao [1983]: *Theory of charges: a study of finitely additive measures*. Academic Press. London.
- [11] P Ressel P (1977): Some continuity and measurability results on spaces of measures. *Mathematica Scandinavica*, 40(1): 69–78.
- [12] AN Shiryaev (1996): *Probability*. Second edition. Springer. New York.
- [13] WD Sudderth (2020): Optimal Markov strategies. *Decisions in Economics and Finance*, 43, 43–54.
- [14] M Voorneveld (2010): The possibility of Impossible stairways: Tail events and countable player sets. *Games and Economic Behavior*, 68(1), 403–410.