האוניברסיטה העברית בירושלים

THE HEBREW UNIVERSITY OF JERUSALEM

A DISCOUNTED STOCHASTIC GAME WITH NO STATIONARY NASH EQUILIBRIUM

By

YEHUDA (JOHN) LEVY

Discussion Paper # 596R   May 2012

מרכז לחקר הרציונליות

CENTER FOR THE STUDY OF RATIONALITY

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

A Discounted Stochastic Game with No Stationary Nash Equilibrium*

Yehuda Levy†

May 16, 2012

Abstract

We present an example of a discounted stochastic game with a continuum of states, finitely many players and actions, and deterministic transitions, that possesses no measurable stationary equilibria, or even stationary approximate equilibria. The example is robust to perturbations of the payoffs, the transitions, and the discount factor, and hence gives a strong nonexistence result for stationary equilibria. The example is a game of perfect information, and hence it also does not possess stationary extensive-form correlated equilibrium. Markovian equilibria are also shown not to exist in appropriate perturbations of our example.

Keywords: Stochastic Game, Discounting, Stationary Equilibrium

1 Introduction

1.1 Background

The question of the existence of stationary equilibria in discounted stochastic games with uncountable state spaces has received much attention. The increasing usefulness of these games in modeling economic situations, combined with the simplicity and universality of stationary strategies, has made equilibrium existence and characterization results a very active area of research. However, it was unknown whether the general models of such games did indeed possess stationary equilibria, which was known to be true in the case of discrete state spaces.

Stochastic games were introduced by Shapley (1953). In a stochastic game, players play in stages. At each stage, the game is in one of the available states,

*Research supported in part by Israel Science Foundation grants 1123/06 and 1596/10.

†Center for the Study of Rationality, and Department of Mathematics, The Hebrew University of Jerusalem, 91904 Jerusalem, Israel.
Tel.: (972-2)-6584135, Fax: (972-2)-6513681, E-mail: john.calculus@gmail.com

and each player chooses an action from the action spaces in that state. The actions chosen then determine a probability distribution on the state space which is used to determine the state at the next stage.

At every stage a player receives a payoff that is a function of the current state and the actions of all the players. The β-discounted game, which is the focus of this paper,[1] is the one in which each player receives the β-discounted sum of the stream of his stage payoffs.

A strategy for a player is a rule by which a player chooses, at each stage, a mixed action from his action space in that state. In the most general setup, one allows players to choose their actions contingent on the choices of actions and states up to the present time. A particular class of strategies, the *stationary strategies*, in which a player makes his decision based only on the current state, has been particularly studied in games with discounted payoffs.

There are two main reasons for this focus. First of all, stationary strategies are the natural class of strategies for the discounted payoff evaluation, as sub-games that are defined by different histories but begin at the same state are strategically equivalent: players will have the same preferences over plays in one sub-game as in the other. The view that strategies should only dependent on payoff-relevant data in the discounted game is highlighted by Maskin and Tirole (2001), who use it to motivate the development of the concept of *Markov Perfect Equilibria*. Hellwig and Leininger (1998), call this view the *subgame-consistency principle*, i.e., "the behaviour principle according to which a player's behaviour in strategically equivalent subgames should be the same, regardless of the different paths by which these subgames might be reached." Harsanyi and Selten (1988, p. 73) argue similarly that "invariance with respect to isomorphisms" is "an indispensable requirement for any rational theory of equilibrium point selection that is based on strategic considerations exclusively." Hence, stationarity of the equilibrium strategies is, in fact, dictated by considerations of rationality and consistency.

The other main reason for focusing on the class of stationary strategies, as discussed in Duffie et al. (1994), is because of their simplicity and easy implementation. This class is clearly the simplest class of strategies that can yield equilibria in many general classes of games. Furthermore, other equilibria seem to require sophisticated coordinate. To quote Guesnerie and Woodford (1992, Section 3), "An equilibrium which does not display minimal regularity through time - maybe stationarity - is unlikely to generate the coordination between agents that it assumes."

Shapley (1953) proved the existence of stationary optimal strategies for each

[1] This was the method of evaluation of the stream of payoffs used by Shapley (1953) when he introduced the concept of stochastic games.

player in the β-discounted two-player zero-sum stochastic games with finitely many states and actions. This was generalized to the existence of equilibrium in stationary strategies in n-person stochastic games with finitely many states and actions, independently, by Fink (1964), Takahashi (1964), Rogers (1969), and Sobel (1971). These results were extended to the case of countably infinite state space[2] by Parthasarathy (1973).

A framework for stochastic games with a continuum of states was introduced by Maitra and Parthasarathy (1970). They worked with the zero-sum case, and extended Shapley's result. (For further extensions in the zero-sum case, see Nowak (1985a), and the references there.) The earliest attempt at generalization to games with a continuum of states in the non-zero-sum case seems to be Sobel (1973), who gave a proof for the existence of β-discounted equilibria in stationary strategies; however, his proof is flawed (this was already pointed out by Federgruen (1978).)

Himmelberg et. al. (1976), began what would become a long series of equilibrium-existence results for discounted stochastic games with uncountable state spaces; they proved the existence of stationary equilibria under very strong separability conditions on the payoffs and transitions and finite action spaces. In another noteworthy result, Parthasarathy and Sinha, (1989), proved the existence of stationary equilibria for finite action sets and state-independent transitions; this result was strengthened significantly by Nowak (2002), to the case of compact action sets and transitions that are combinations of a given finite set of measures. Nowak (1985b), gave a condition[3] under which stationary ε-equilibrium will exist; the condition, which we will refer to as the *absolute continuity condition*, henceforth ACC, has subsequently been assumed in many works. It means that there exists a fixed measure ν on the state space such that any transition measures arising in any state and for any action profile are absolutely continuous w.r.t. ν; see Definition 2.1.1.

Alongside these developments, there began to emerge a literature of models in economics that were themselves particular cases of discounted stochastic games.[4] A continuum of states is the natural framework for many of these applications, which usually work with continuous variables - resources, stocks, population percentages, share percentages, etc. Levhari and Mirman (1980), presented a model of resource extraction. Their work led to a series of papers on resource extraction and capital accumulation; Amir (1996), combines many of these works into a single stochastic games framework and presents a general

[2]See also Federgruen (1978) and Whitt (1980); the latter allows for infinite action spaces.

[3]This paper also imposed additional integrability conditions on the Radon-Nikodym derivates of the transitions w.r.t. this distribution; however, these are satisfied automatically for games with finite action spaces, and it is not difficult to show that when the action spaces are compact and the continuity conditions there hold, the game can be appropriately approximate by a game with finite action spaces.

[4]Even though they were often not recognized as such at the time.

stationary equilibrium existence result. Monahan and Sobel (1994), present a dynamic oligopoly model, in which market share changes over time. Rosenthal (1982) and Beggs and Klemperer (1992), present models of repeated duopolistic competition. Nowak (2007) gives methods for finding stationary equilibria in several classes of games arising in economics. Many other examples and elaborations can be found in Amir (2003).

There have also been works dealing with wider classes of strategies. Rieder (1979) proved the existence[5] of ε-equilibria in the class of *Markovian*[6] strategies. Incorrect proofs[7] appeared in Amir (1991) and Chakrabarti (1999) claiming to prove the existence of Markovian equilibria for games which satisfy ACC. A significant breakthrough came when Mertens and Parthasarathy (1987) proved the existence of history-dependent equilibria; their result was given alternate proofs by Solan[8] (1998) and Maitra and Sudderth[9] (2007). Dutta and Sundaram (1998) provide an excellent survey of many of these results.

Another direction that emerged was the study of extensive-form correlated equilibria[10] Under ACC, Nowak and Raghavan (1992) demonstrated the existence of stationary extensive-form correlated equilibria.[11] Under conditions stronger than ACC, Duffie et al (1994) proved the existence of extensive-form correlated equilibria[12] that induces an ergodic process on the state space. Although not directly related to stochastic games, a parallel result that employs many of the same techniques was presented by Harris et al. (1995) where an extensive-form correlated equilibrium was proven to exist for infinite-extensive form games with continuous payoffs.

1.2 Results

The purpose of this paper is to show, by example, that discounted stochastic games with uncountable state space need not possess equilibria in stationary

[5]Rieder (1979) allows for the player set to be countably infinite, and action sets to be infinite.

[6]This class is more general than the class of stationary strategies; *Markovian strategies* may depend on both the current state and number of stages played so far. These are also often referred to as *feedback strategies* or *closed-loop strategies*; see, e.g., Basar and Olsder (1999).

[7]The error in Amir (1991) appears on p. 157. When choosing for each state s, a subsequence $\{v_l(s)\}$ of the sequence $\{v_n(s)\}$ of n-stage equilibrium selection, the indices of the subsequence may depend on s. For an explanation of the error in Chakrabarti (1999), see Remark 7 of Nowak (2007).

[8]Solan (1998) assumes the ACC.

[9]Maitra and Sudderth (2007) also present results for more general continuous payoff evaluations.

[10]For the general notion of extensive-form correlated equilibrium, see Forges (1986).

[11]That is, existence of a stationary equilibrium when a stationary correlation device is allowed to be added; their result has result been generalized by Duggan (2011) to allow for more general correlation devices, in the form of state-dependent noise.

[12]In consistency with economic literature, they refer to these as *sunspot equilibria*.

strategies, even when the action sets are finite (and state-invariant), and the player set is finite. In fact, our example has stronger implications: for each $\beta > 0$, we build a game with the following properties:

- For $\varepsilon > 0$ small enough, stationary ε-equilibria do not exist.
- Stationary extensive-form correlated equilibria do not exist.
- The game is a game of perfect information,[13] with deterministic transitions.
- For $\varepsilon > 0$ small enough, if payoffs, transitions,[14] and discount factor are perturbed less than ε the resulting game still does not possess stationary ε-equilibria.
- The game does not possess sub-game[15] perfect Markov equilibria.
- For any $\varepsilon > 0$, there is a perturbation of our example of less than ε which does not possess Markov equilibria.[16]

A few remarks are in order. First of all, we see that the case of atomic transitions contrasts sharply then with games satisfying ACC. Indeed, as Nowak (1985b) showed, in such a case ε-equilibria do exist in stationary strategies. Furthermore, as Nowak and Raghavan (1992) showed, games satisfying ACC do possess stationary extensive-form correlated equilibria; in games of perfect information, it is immediate to construct from these (uncorrelated) stationary equilibria.

We also mention an example from Harris et al. (1995) of a two-stage extensive-form game without an equilibrium. Duggan (2011) pointed out this example can be modeled as a stochastic game. However, as Duggan (2011) also shows, in representing this example as a stochastic game, one has to allow for transitions that are not strongly[17] continuous on the infinite action spaces. Therefore, this example does not fit into the models that are usually studied in works establishing equilibrium existence results.

As listed, our example is robust to small enough perturbations of the transitions (as well as the payoffs and discount factor). This robustness not only strengthens the counterexample, it also has implications pertaining to the delicate matter of non-measurable equilibria. As we discuss in Section 6.1, for games with purely atomic transitions, the expected payoff can be defined even

[13]At each stage, there is only one player whose action has any affect on the transitions or payoffs.

[14]In the total variation norm.

[15]A strategy profile is a sub-game perfect equilibrium if it represents a Nash equilibrium of every sub-game of the original game.

[16]Markov equilibria which are not sub-game perfect do exist in the unperturbed version of the example.

[17]I.e., norm-continuous.

if the players use nonmeasurable strategies. Furthermore, in nonmeasurable strategies, stationary equilibria always exists in purely atomic games. We will elaborate in that section why nonmeasurable strategies are not only undesirable but are also pathological and ill defined; nevertheless, one may insist on the validity of nonmeasurable equilibria. However, in generic perturbations, payoffs of the game under nonmeasurable strategies becomes mathematically meaningless, and hence so do equilibria. Hence, it is important that the purely atomic transitions can be perturbed in our example.

Finally, we are also able provide an answer concerning Markovian equilibria. Although it is still not known whether Markovian equilibria must exist under the ACC assumption - as remarked in Section 1.1, incorrect proofs have appeared - we show that appropriate perturbations of our example do not posess Markovian equilibria.

1.3 Organization of Paper

In Section 2 we present the formal stochastic game model and a precise statement of the some of the properties which our example satisfies. In Section 3, we construct the stochastic game and present some observations pertaining to equilibria in it or in its perturbations. Section 4 proves that no stationary equilibria exist. Section 5 shows how the arguments of Sections 3.4 and 4 can be modified to show that equilibria need not exist in Markovian strategies. Section 6 contains some remarks and acknowledgements.

2 Model and Results

2.1 Stochastic Game Model

The components of a discounted stochastic game[18] with a continuum of states and finitely[19] many actions are the following:

- A standard Borel[20] space Ω of states.
- A finite set $\mathcal{P}$ of players.
- A finite set of actions I^p for each $p \in \mathcal{P}$. Denote $\overline{I} = \prod_{p \in \mathcal{P}} I^p$
- A discount factor $\beta \in (0,1)$.
- A bounded payoff function $r : \Omega \times \overline{I} \to \mathbb{R}^{\mathcal{P}}$, which is Borel-measurable.
- A transition function $q : \Omega \times \overline{I} \to \Delta(\Omega)$, which is Borel-measurable (where $\Delta(\Omega)$, the space of regular Borel probability measures on Ω, possesses the Borel structure induced from the topology of narrow convergence).

[18]Much of this introduction to the model follows Nowak (2003).

[19]This is a particular case of the general model, which allows for compact actions spaces that are state-dependent; see, e.g., Mertens and Parthasarathy (1987).

[20]That is, a space that is homeomorphic to a Borel subset of a complete, metrizable space.

The game is played in discrete time. If $z \in \Omega$ is a state at some stage of the game and the players select an $a \in \overline{I}$, then $q(z,a)$ is the probability distribution of the next state of the game. A *behavioral strategy* for a player is a Borel-measurable mapping that associates with each given history a probability distribution on the set of actions available to him. A *stationary strategy* for Player p is a behavioral strategy that depends only on the current state; equivalently, it is a Borel-measurable mapping that associates with each state $s \in \Omega$ a probability distribution on the set I^p.

Let $H^\infty = (S \times \overline{I})^{\mathbb{N}}$ be the space of all infinite histories of the game, endowed with the product σ-algebra. For any profile of strategies $\sigma = (\sigma^p)_{p \in \mathcal{P}}$ of the players and every initial state $z_1 = z \in \Omega$, a probability measure P_z^σ and a stochastic process $(z_n, a_n)_{n \in \mathbb{N}}$ are defined on H^∞ in a canonical way, where the random variables z_n, a_n describe the state and the action profile chosen by the players, respectively, in the n-th stage of the game (see, e.g., Chapter 7 in Bertsekas and Shreve (1996)). The expected payoff vector under σ, in the game starting from state z, is:

$$\gamma_\sigma(z) = E_z^\sigma\big(\sum_{n=1}^{\infty} \beta^{n-1} r(z_n, a_n)\big) \tag{2.1}$$

Let Σ^p denote the set of behavioral strategies for Player $p \in \mathcal{P}$, and $\Sigma = \prod_{p \in \mathcal{P}} \Sigma^p$. A profile $\sigma \in \Sigma$ will be called a Nash equilibrium if

$$\gamma_\sigma^p(z) \geq \gamma_{(\tau, \sigma^{-p})}^p(z), \ \forall p \in \mathcal{P}, \forall z \in \Omega, \forall \tau \in \Sigma^p \tag{2.2}$$

and it will be called an ε-equilibrium if

$$\gamma_\sigma^p(z) \geq \gamma_{(\tau, \sigma^{-p})}^p(z) - \varepsilon, \ \forall p \in \mathcal{P}, \forall z \in \Omega, \forall \tau \in \Sigma^p \tag{2.3}$$

Denote, for $\sigma \in \Sigma$ that is stationary, for a state $z \in \Omega$, and for a mixed action profile $a \in \prod_{p \in \mathcal{P}} \Delta(I^P)$,

$$X_\sigma(z,a) := r(z,a) + \beta \int_\Omega \gamma_\sigma(t) dq(z,a)(t) \tag{2.4}$$

For $\sigma \in \Sigma$ that is a stationary strategy profile, it is easily shown that (2.2) implies[21] that

$$X_\sigma^p(z, \sigma(z)) \geq X_\sigma^p(z, (b, \sigma^{-p}(z))), \ \forall p \in \mathcal{P}, \forall z \in \Omega, b \in I^p \tag{2.5}$$

i.e., that for all z, $\sigma(z)$ is an equilibrium in the game with payoff $X_\sigma(z, \cdot)$, and that (2.3) implies[22] that

$$X_\sigma^p(z, \sigma(z)) \geq X_\sigma^p(z, (b, \sigma^{-p}(z))) - \varepsilon, \ \forall p \in \mathcal{P}, \forall z \in \Omega, b \in I^p \tag{2.6}$$

[21] In fact, they are equivalent; both directions follow from standard dynamic programming results.

[22] In this case, they are not quite equivalent; (2.6) implies that σ is an $\frac{\varepsilon}{1-\beta}$-equilibrium.

i.e., that for all z, $\sigma(z)$ is an ε-equilibrium in the game with payoff $X_\sigma(z,\cdot)$.

For reference, we define:

Definition 2.1.1. *The game above is said to satisfy the Absolute Continuity Condition (ACC) if there is $\nu \in \Delta(\Omega)$ such that for all $z \in \Omega$, $a \in \overline{I}$, $q(z,a)$ is absolutely continuous w.r.t. ν.*

2.2 The Result

For any given $\beta \in (0,1)$, we will construct a stochastic game $(\Omega, \mathcal{P}, (I^p), \beta, r, q)$ (with purely atomic transitions and perfect information) that does not possess a stationary (measurable[23]) equilibrium.

In fact, we will deduce a stronger result for the game:

There exist $\varepsilon > 0$ (which depends on β) such that if $r' : \Omega \times \overline{I} \to \mathbb{R}^{\mathcal{P}}$ and $q' : \Omega \times \overline{I} \to \Delta(\Omega)$ satisfy the measurability conditions given in the model of Section 2.1, and also satisfy[24]

$$||r'(z,a) - r(z,a)||_\infty < \varepsilon,\ ||q'(z,a) - q(z,a)|| < \varepsilon,\ \forall z \in \Omega,\ \forall a \in \overline{I}$$

and

$$|\beta' - \beta| < \varepsilon$$

then the game $(\Omega, \mathcal{P}, (I^p), \beta', r', q')$ does not possess a stationary ε-equilibrium.

3 The Example

Henceforth, let $\beta \in (0,1)$ be a fixed discount factor, let $Y = \{-1,1\}^\omega$, where $\omega = \{0,1,2,\ldots\}$, let T denote the left shift-operator operator[25] on Y, and let μ denote the Lebesgue measure on Y.

3.1 An Informal Description of the Construction

We begin by giving an informal description of the primary properties of the game.

We begin this description by allowing a countable set of players - or, if you will, one player in each generation $n \in \mathbb{N}$. The state space will be $\mathbb{N} \times Y$, along with a "quitting state" $\overline{0}$; all payoffs are zero after the first transition to the quitting state. The transition from a state (n,y) will either be to state

[23] We point out the measurability in light of Section 6.1. The state space will be a finite product of Cantor sets, and the measurability we refer to is with respect to the Lebesgue σ-algebra.
[24] The latter distance is the total variation norm.
[25] That is, $(Tx)_n = x_{n+1}$.

$(n+1, T(y))$ or to $\overline{0}$. In a state $(n, *)$, only Player n's action has any affect on either payoffs or transitions; we can think of him as the only "active" player. Player n receives payoffs both when he is active, in state (n, y), and in the following state, $(n+1, T(y))$ (if the game has not quit). This is reminiscent of models of overlapping generation games: each player can be imagined as being alive for two generations. In the first generation, he is "young" and takes an action, and receives some resulting payoff. In the second generation, he is "old"; he does not take an action but he does receive a payoff as a result of the "young" player's action.

Each player can play either L or R. The factor that affects the structure of the payoff and transition in state (n, y) is the 0-th bit of y, denoted $\kappa(y)$. The key is that we define the payoff and transitions such that if Player $n+1$ would play one particular action with high probability in state $(n+1, T(y))$, then Player n in state (n, y) will want to match Player $n+1$'s expected action if $\kappa(y) = 1$, and will want to mismatch it if $\kappa(y) = -1$. Furthermore, we arrange the structure such that regardless of Player $n+1$'s mixed action in state $(n+1, T(y))$, Player n will not be indifferent between his own actions in both of the possible states preceding $(n+1, T(y))$.

The modification to finitely many players is done simply: we just have the generations repeat themselves periodically, with some period M; the state space becomes $\{0, \ldots, M-1\} \times Y \cup \{\overline{0}\}$, with the generation-counter running via addition modulo M. If M is chosen large enough - it will depend on the discount factor - each Player will make a decision based only on the payoffs of the current and following stages when he is called to play; the payoffs from his next "reincarnation", M stages later, will be negligible and will not affect his decision.

3.2 Preliminary Notations and Variables

We fix $\varepsilon, \rho, \delta > 0$, and $M \in \mathbb{N}$, as follows:

- $\delta < \frac{1}{40}$.
- $\rho = \frac{1}{\beta}$.
- Let $D(\beta, \rho) > 0$ be defined as in Proposition 7.1.1 of the Appendix. We choose ε small enough[26] so that it satisfies

$$\varepsilon < \min\left[\beta, 1-\beta,\ \frac{\delta}{2} \cdot \frac{1-\beta}{D(\beta,\rho)},\ \frac{\delta}{20}\right]$$
$$\left|\frac{1}{1-\beta} - \frac{1}{1-\beta \pm \varepsilon}\right| \leq \frac{1}{1-\beta}$$

- $M > 1$, such that $\rho \cdot \sum_{j=M}^{\infty} \beta^j < \frac{\delta}{2}$.

[26] The requirements $\varepsilon < \beta, 1-\beta$ are only so that $\beta \pm \varepsilon \in (0,1)$, hence the second inequality is well-defined.

If q is a mixed action over an action space A and $a \in A$, then $q[a]$ denotes the probability that q chooses A.

3.3 Construction

We will construct the game $(\Omega, \mathcal{P}, (I^p), \beta, r, q)$. For M defined in Section 3.2, denote $Z = \omega_M \times Y$, where $\omega_M = \{0, \ldots, M-1\}$. The state space will be $\Omega = Z \cup \{\overline{0}\}$, where $\overline{0}$ is an absorbing[27] state with payoff 0 for all players.

The set of players in the game will be $\mathcal{P} = \omega_M$. Each player's action set is $I = \{L, R\}$. For $n \in \omega_M$, let

$$n^{\pm} = (n \pm 1)_{mod\ M} \in \omega_M$$

Define $S : Z \to Z$ by $S(n, y) = (n^+, T(y))$. Also for $z = (n, y) \in Z$, we denote:

$$\kappa(z) = y_0 \tag{3.1}$$

$$n(z) = n \tag{3.2}$$

$$n^{\pm}(z) = n^{\pm} \tag{3.3}$$

The game is a game of perfect information: that is, for each[28] $z \in Z$, there is only one player, $n(z)$, whose action has any effect on payoffs or transitions. Fix a state $z \in Z$:

- Only $n(z)$ and $n^-(z)$ receive non-zero payoffs in state z. That is, if $p \notin \{n(z), n^-(z)\}$, then $r^p(z, \cdot) \equiv 0$.
- The payoff to players $n(z), n^-(z)$, and the next state z', are all determined only by the action of Player $n(z)$ and are given by the following rules:

If $\kappa(z) = 1$:

$a^{n(z)} =$	L	R
$r^{n(z)}(z, a) =$	0	0.3
$r^{n^-(z)}(z, a) =$	$\frac{1}{\beta}$	0
$z' =$	$S(z)$	$\overline{0}$

If $\kappa(z) = -1$:

$a^{n(z)} =$	L	R
$r^{n(z)}(z, a) =$	0.7	0
$r^{n^-(z)}(z, a) =$	$\frac{1}{\beta}$	0
$z' =$	$\overline{0}$	$S(z)$

Remark 3.3.1. Observe that $||r||_\infty \leq \rho = \frac{1}{\beta}$.

3.4 Observations and Characterization of Equilibria

Assume a fixed perturbation $(\Omega, \mathcal{P}, \{L, R\}^{\mathcal{P}}, \beta', r', q')$ of the game defined in Section 3.3, satisfying the conditions of Section 2.2 with ε given in Section 3.2. Fix a stationary ε-equilibrium profile σ of this game. Recall the notation γ_σ and

[27] A state $z \in \Omega$ is called an absorbing state of $q(z \mid z, a) = 1$ for all action profiles a.

[28] In $\overline{0}$, no player's action has any effect.

X_σ from Section 2.1; we will apply these notations to the unperturbed game, and $\gamma'_\sigma, X'_\sigma$ to the perturbed game.[29] The following lemma follows from Lemma 7.1.1 of the Appendix, the choice of ε, and Remark 3.3.1.

Lemma 3.4.1. *For any behavioral strategy τ,*

$$||X_\tau(\cdot) - X'_\tau(\cdot)||_\infty < \frac{\delta}{2}$$

For $p \in \mathcal{P}$ and $z \in Z \subseteq \Omega$, $\sigma^p(z)$ will denote the probability distribution on $\{L, R\}$ induced by Player p's mixed action in state z. Recall the definition of $\kappa(z)$ from (3.1), and denote further that:

$$\ell(z) = \sigma^{n(z)}(z)[L] \tag{3.4}$$

Definition 3.4.2. *A state $z \in Z$ will be called L-quasi-pure (resp. R-quasi-pure) if $\ell(z) > 1 - \delta$ (resp. $\ell(z) < \delta$). If z is either L- or R-quasi-pure, we may simply refer to z as being quasi-pure.*

We will study the relationship between $\sigma(S(z))$ and $\sigma(z)$.

Lemma 3.4.3. *For any stationary strategy τ and any $z \in Z$, we have*

$$||X^{n(z)}_\tau(z, \cdot) - K(z)||_\infty < \frac{\delta}{2}$$

where

$$K(z) = r^{n(z)}(z, \cdot) + \beta q(z' = S(z) \mid z, \cdot) \cdot r^{n(z)}\big(S(z), \tau(S(z))\big)$$

In other words, up to an error of $\frac{\delta}{2}$, the payoff (in the unperturbed game) to Player $p = n(z)$ is influenced only by the payoff of the current stage and the following stage. The lemma follows immediately from the requirement of M in Section 3.2 and from the definition of the payoffs.

Notation 3.4.4. *Let $\langle \alpha_L; \alpha_R \rangle$, for $\alpha_L, \alpha_R \in \mathbb{R}$, denote the single-player normal form game - a.k.a., the decision - that gives payoff α_L (resp. α_R) if the player plays L (resp. R).*

Lemmas 3.4.1 and 3.4.3 show that the decision facing player $n(z)$ at state z is δ-close to the decision $K(z)$. Then we have

$$K(z) = \begin{cases} \langle \ell(S(z)); \frac{3}{10} \rangle & \text{if } \kappa(z) = 1 \\ \langle \frac{7}{10}; \ell(S(z)) \rangle & \text{if } \kappa(z) = -1 \end{cases} \tag{3.5}$$

[29]In defining γ'_σ by the definition analog to (2.1), one needs to use the expectation operator defined by the dynamics in the perturbed game - that is, via the strategy σ and the transition kernel q' - as well as the discount factor β' and the payoff function r'. A similar remark holds for the definition of X'_σ via the analog of (2.4)

Probability of Playing Left in Best Reply to $K(z)$ as Function of $q = \ell(S(z))$:	
$\kappa(z) = 1$	$\kappa(z) = -1$
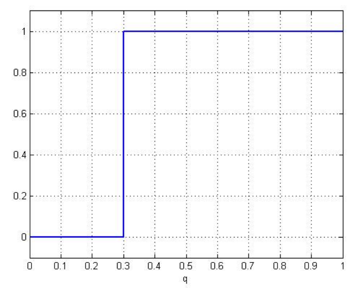	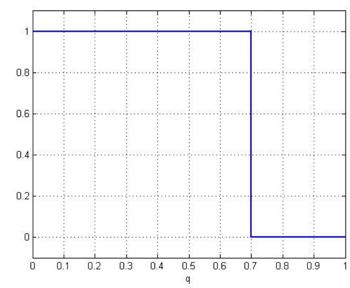

We formally derive the necessary implications in the following lemma:

Lemma 3.4.5. *Let $z \in Z$. If $\kappa(z) = 1$,*

$$\ell(S(z)) < \frac{1}{5} \Longrightarrow \ell(z) < \delta,\ \ell(S(z)) > \frac{2}{5} \Longrightarrow \ell(z) > 1 - \delta \tag{3.6}$$

and if $\kappa(z) = -1$,

$$\ell(S(z)) < \frac{3}{5} \Longrightarrow \ell(z) > 1 - \delta,\ \ell(S(z)) > \frac{4}{5} \Longrightarrow \ell(z) < \delta \tag{3.7}$$

Proof. Let

$$Y(z)(\cdot) = \langle (X'_\sigma)^{n(z)}(z, L, \sigma^{-n(z)}(z)); (X'_\sigma)^{n(z)}(z, R, \sigma^{-n(z)}(z)) \rangle$$

where we have used Notation 3.4.4. In other words, $Y(z)$ is the decision facing Player $n(z)$, when the payoffs are given by $(X'_\sigma)^{n(z)}(z, \cdot)$, and the other players are restricted[30] to playing by $\sigma^{-n}(z)$. Lemmas 3.4.1, 3.4.3, and equality (3.5) show that

$$||Y(z) - \langle \ell(S(z)); \frac{3}{10} \rangle||_\infty < \delta, \text{ if } \kappa(z) = 1 \tag{3.8}$$

$$||Y(z) - \langle \frac{7}{10}; \ell(S(z)) \rangle||_\infty < \delta, \text{ if } \kappa(z) = -1 \tag{3.9}$$

We carry out the proof of the Lemma for the case $\kappa(z) = 1$; the other case follows similarly. If $\ell(S(z)) < \frac{1}{5}$, then

$$Y(z)[L] \leq \ell(S(z)) + \delta < 0.2 + \delta,\ 0.3 - \delta \leq Y(z)[R]$$

Therefore,

$$Y(z)[R] - Y(z)[L] \geq \frac{1}{10} - 2\delta \geq \frac{1}{20}$$

[30]Of course, the way we have defined the payoffs, it is irrelevant how the others play in state z.

The criteria (2.6) (for the perturbed game) implies that playing L with probability $\ell(z)$ is an ε-best-reply in $Y(z)$, and hence $\ell(z) < 20\varepsilon \leq \delta$. On the other hand, if $\ell(S(z)) > \frac{2}{5}$,

$$Y(z)[L] \geq \ell(S(z)) > 0.4 - \delta,\ 0.3 + \delta \geq Y(z)[R]$$

Therefore,

$$Y(z)[L] - Y(z)[R] \geq \frac{1}{10} - 2\delta \geq \frac{1}{20}$$

and we similarly derive that in this case, $\ell(z) > 1 - 20\varepsilon \geq 1 - \delta$. □

We now get to the two aspects of the approximate equilibria that we really need, summed up in the following two lemmas:

Lemma 3.4.6. *If $S(z)$ is quasi-pure in σ, then so is z. If the former is a-quasi-pure ($a \in \{L, R\}$), then the latter is as well if and only if $\kappa(z) = 1$.*

Proof. The lemma follows by repeated use of Lemma 3.4.5.

- If $S(z)$ is L-quasi-pure and $\kappa(z) = 1$, then $\ell(S(z)) < \delta < \frac{1}{5}$, so $\ell(z) < \delta$.
- If $S(z)$ is L-quasi-pure and $\kappa(z) = -1$, then $\ell(S(z)) < \delta < \frac{3}{5}$, so $\ell(z) > 1 - \delta$.
- If $S(z)$ is R-quasi-pure and $\kappa(z) = 1$, then $\ell(S(z)) > 1 - \delta > \frac{2}{5}$, so $\ell(z) > 1 - \delta$.
- If $S(z)$ is R-quasi-pure and $\kappa(z) = -1$, then $\ell(S(z)) > 1 - \delta > \frac{4}{5}$, so $\ell(z) < \delta$.

□

Lemma 3.4.7. *For any $z \in Z$, at least one of the two states in $S^{-1}(z)$ is quasi-pure. (Note that this is so even if z is not quasi-pure.)*

Proof. We must have at least one of the following two inequalities:

$$\ell(S(z)) > \frac{2}{5},\ \ell(S(z)) < \frac{3}{5}$$

Suppose that the left inequality holds. Lemma 3.4.5 then shows that if $z' \in S^{-1}(z)$ with $\kappa(z') = 1$, then $\ell(z') > 1 - \delta$ and hence z' is L-quasi-pure. In the case of the right inequality, we deduce similarly that if $z'' \in S^{-1}(z)$ with $\kappa(z'') = -1$, then z'' is also L-quasi-pure.

□

4 Nonexistence of Stationary Equilibria

Recall that μ is the Lebesgue-measure on Y, and let λ be the uniform measure on ω_M; let $\nu = \lambda \times \mu$. Assume that σ is a stationary ε-equilibrium, as in Section 3.4, measurable w.r.t. ν. We shall use Lemmas 3.4.6 and 3.4.7 to show that σ cannot be a (ν-measurable[31]) equilibrium. Assume, to the contrary, that it is.

Lemma 4.0.8. *Let $\Xi = \{z \in Z \mid z$ is not quasi-pure$\}$. Then $\nu(\Xi) = 0$.*

Proof. By assumption, Ξ is ν-measurable. Lemma 3.4.6 implies that

$$S(\Xi) \subseteq \Xi \tag{4.1}$$

Let $\iota : Z \to Z$ be the involution defined such that $\iota(n, y)$ is obtained from (n, y) by changing only the 0-th bit of y. Lemma 3.4.7 then implies that

$$\Xi \cap \iota(\Xi) = \emptyset \tag{4.2}$$

Furthermore, for any $B \subseteq Z$,

$$S^{-1}(S(B)) = B \cup \iota(B) \tag{4.3}$$

S and ι are both ν-preserving.[32] Also observe that $S(\Xi)$ is ν-measurable.[33] Hence (4.1), (4.2), and (4.3) imply that

$$2\nu(\Xi) = \nu(\Xi) + \nu(\iota(\Xi)) = \nu(S^{-1}(S(\Xi))) = \nu(S(\Xi)) \leq \nu(\Xi)$$

Hence, $\nu(\Xi) = 0$. □

Define the map $g : Z \to \{-1, 1\}$ by $g(z) = 1$ if and only if z is L-quasi-pure. Denote for all $y, y' \in Y$, $D(y, y') = \{j \in \omega \mid y_j \neq y'_j\}$, and if $D(y, y')$ is finite, $N(y, y') = \#D(y, y')$, $M(y, y') = \max D(y, y')$.

Lemma 4.0.9. *For each $n \in \omega_M$, μ-almost every $y \in Y$, we have*

$$g(n, y) = (-1)^{N(y,y')} g(n, y'), \ \forall y' \in Y \text{ s.t. } N(y, y') < \infty \tag{4.4}$$

Proof. By Lemma 3.4.6 and Lemma 4.0.8, we see that for almost every $z = (n, y) \in Z$, $g(z) = y_0 \cdot g(S(z))$, and hence for all k and a.e. z,

$$g(z) = y_0 \cdot \cdots \cdot y_{k-1} \cdot g(S^k(z))$$

If $N(y, y') < \infty$, $z = (n, y)$, $z' = (n, y')$, then $S^{M(y,y')}(z) = S^{M(y,y')}(z')$ and $N(y, y') = \prod_{j \leq M(y,y')} \frac{y_j}{y'_j}$; hence the result follows. □

---

[31] Where ν is also a measure on Ω via inclusion.

[32] Recall that a mapping ψ on a measure space (Ω, λ) is *measure-preserving* if $\lambda(\psi^{-1}(A)) = \lambda(A)$ for all λ-measurable $A \subseteq \Omega$. ι is clearly ν-preserving, and the map $n \to n^+$ in ω_M is clearly λ-preserving; that shifts are Lebesgue-measure preserving, and that the product of measure-preserving systems are also measure-preserving, are standard results in ergodic theory.

[33] This is easy to establish in the case that $\Xi \subseteq \{1\} \times Y$ or $\Xi \subseteq \{-1\} \times Y$, and the general case follows.

Proposition 4.0.10. *There does not exist a μ-measurable function $f : Y \to \{-1, 1\}$, such that for a.e. $y \in Y$,*

$$f(y) = (-1)^{N(y,y')} f(y'),\ \forall y' \in Y \text{ s.t. } N(y, y') < \infty \tag{4.5}$$

Proposition 4.0.10 contradicts 4.0.9, and therefore completes our proof that there are no stationary equilibria. Before the proof, we recall several notions: Let S_ω denote the set of permutations π on ω such that $\exists N \in \omega, \forall n > N, \pi(n) = n$. S_ω acts on Y by $(\pi(y))_n = y_{\pi^{-1}(n)}$. A transposition (on ω) is an element π of S_ω for which there are $i, j \in \omega$ with $\pi(i) = j$ and $\pi(j) = i$, and $\pi(k) = k$ for all $k \neq i, j$. It is well known that every element of S_ω is a composition of finitely many transpositions. We also denote by $\chi_j : Y \to Y$ the involution which changes only the j-th bit of the sequence.

Proof. Suppose that we did have such an f. Denote

$$L = \{y \in Y \mid f(y) = 1\}$$

Note that $\mu(L) = \frac{1}{2}$: first, note that $f(y) = -f(\chi_0(y))$ for μ-a.e. y. Hence, for a.e. y, exactly one of the following options holds: $y \in L$ or $\chi_0(y) \in L$ (equivalently, $y \in \chi_0(L)$). Hence $\mu(\chi_0(L) \cap L) = 0$, $\mu(\chi_0(L) \cup L) = 1$.

On the other hand, let $\pi \in S_\omega$ and $y \in Y$ for which (4.5) holds. We contend that $y \in L$ if and only if $\pi(y) \in L$; it's enough to check this in the case that π is a transposition. We have either $\pi(y) = y$ or $\pi(y) = \chi_i \circ \chi_j(y)$, so $N(\pi(y), y) \in \{0, 2\}$. Therefore, $\mu(\pi(L)\Delta L) = 0$ for all $\pi \in S_\omega$, where Δ denotes the symmetric difference of sets. By the Hewitt-Savage zero-one law, $\mu(L) = 0$ or $\mu(L) = 1$, a contradiction.

□

5 Markovian Strategies

5.1 The Concept and Dynamic Programming

A class of strategies that is larger than the stationary strategies is the *Markovian strategies*. A Markovian strategy σ^p for a player $p \in \mathcal{P}$ is a sequence, $\sigma^p = (\sigma_1^p, \sigma_2^p, \ldots)$, where for each $m \in \mathbb{N}$, σ_m^p is a measurable map $\Omega \to \Delta(I^p)$. In other words, a Markovian strategy is a behavioral strategy in which a player's action can depend on the current stage of the game and the current state. We will show that our example does not possess subgame perfect Markovian equilibria.[34] Furthermore, we will show that there are arbitrarily small perturbations of our example that do not possess Markovian equilibria.[35]

[34] A similar argument can show that our example does not posses a Markovian subgame perfect ε-equilibrium - i.e., a Markovian strategy profile which induces an ε-equilibrium in any subgame - but we will settle for simplicity.

[35] The nonperturbed example possesses Markovian equilibria: for all $p \in \mathcal{P}$, let $\sigma_1^p(z) = R$ for $\kappa(z) = 1$ and $\sigma_1^p(z) = L$ for $\kappa(z) = -1$, let $\sigma_2^p(z) = R$ for all z, and let $\sigma_k^p(z)$ be arbitrary for $k \geq 3$.

We state an analog of (2.5) for Markovian equilibrium, which can also be deduced from standard dynamic programming techniques. We adopt the various notations of Section 2.1. Furthermore, if $\sigma = (\sigma_1, \sigma_2, \ldots)$ is a Markovian strategy profile, let σ_{*m} be the Markovian strategy profile $(\sigma_{m+1}, \sigma_{m+2}, \ldots)$, and we generalize the notation of Section 2.1 by defining for each state $z \in \Omega$, and for a mixed action profile $a \in \prod_{p \in \mathcal{P}} \Delta(I^P)$,

$$X^p_{\sigma_{*m}}(z, a) := r(z, a) + \beta \int_\Omega \gamma_{\sigma_{*m}}(t) dq(z, a)(t)$$

Proposition 5.1.1. *Let $(\Omega, \mathcal{P}, (I^p), \beta, r, q)$ be a discounted stochastic game, and let σ be a Markovian strategy profile.*
(a) If σ is a subgame perfect equilibrium, then for every state $z \in \Omega$ and every $m \in \mathbb{N}$ it holds that for every $z \in \Omega$,

$$X^p_{\sigma_{*m}}(z, \sigma_m(z)) \geq X^p_{\sigma_{*m}}(z, (b, (\sigma_m)^{-p}(z))), \ \forall p \in \mathcal{P}, \forall b \in I^p \tag{5.1}$$

(b) If σ is an equilibrium, then (5.1) holds for every state $z \in \Omega$ and every $m \in \mathbb{N}$ that satisfies

$$\exists z_1 \in \Omega, \ P^{z_1}_\sigma(z_m = z) > 0 \tag{5.2}$$

5.2 Markovian Equilibrium in Our Example

Fix some β. We will show that the game $\Gamma = (\Omega, \mathcal{P}, \{L, R\}^{\mathcal{P}}, \beta, r, q)$ defined in Section 3.3 does not have a subgame perfect Markovian equilibria. At the end of this section we remark how to find pertubations of Γ which do not possess Markovian equilibria.

Assume, by way of contradiction, a fixed measurable subgame perfect Markovian equilibrium profile σ.

$$\ell_m(z) = \sigma_m^{n(z)}(z)[L] \tag{5.3}$$

Definition 5.2.1. *For each $m \in \mathbb{N}$, a state $z \in Z$ will be called (L, m)-pure (resp. (R, m)-pure) if $\ell_m(z) = 1$ (resp. $\ell_m(z) = 0$). If z is either (L, m)- or (R, m)-pure, we may simply refer to z as being m-pure.*

Recall Notation 3.4.4. Denote

$$Y_m(z)(\cdot) = \langle (X_{\sigma_{*m}})^{n(z)}(z, L, \sigma_m^{-n(z)}(z)); (X_{\sigma_{*m}})^{n(z)}(z, R, \sigma_m^{-n(z)}(z)) \rangle$$

Combining Lemma 3.4.1 and Lemma 3.4.3 (also easily modified for Markovian strategies) yields these analogs of (3.8) and (3.9):

$$||Y_m(z) - \langle \ell_{m+1}(S(z)); \frac{3}{10} \rangle||_\infty < \delta, \ \text{if } \kappa(z) = 1 \tag{5.4}$$

$$||Y_m(z) - \langle \frac{7}{10}; \ell_{m+1}(S(z)) \rangle||_\infty < \delta, \ \text{if } \kappa(z) = -1 \tag{5.5}$$

The analog of Lemma 3.4.5, established along the same lines (using Proposition 5.1.1.a), is

Lemma 5.2.2. *Let $z \in Z$, $m \in \mathbb{N}$. If $\kappa(z) = 1$,*

$$\ell_{m+1}(S(z)) < \frac{1}{5} \Longrightarrow \ell_m(z) = 0,\ \ell_{m+1}(S(z)) > \frac{2}{5} \Longrightarrow \ell_m(z) = 1 \tag{5.6}$$

and if $\kappa(z) = -1$,

$$\ell_{m+1}(S(z)) < \frac{3}{5} \Longrightarrow \ell_m(z) = 1,\ \ell_{m+1}(S(z)) > \frac{4}{5} \Longrightarrow \ell_m(z) = 0 \tag{5.7}$$

We therefore deduce the following parallels of Lemmas 3.4.6 and 3.4.7:

Lemma 5.2.3. *If $S(z)$ is $m+1$-pure, then z is m-pure. If the former is $(a, m+1)$-pure ($a \in \{L, R\}$), then the latter is (a, m)-pure if and only if $\kappa(z) = 1$.*

Lemma 5.2.4. *For any $z \in Z$, $m \in \mathbb{N}$, at least one of the two states in $S^{-1}(z)$ is m-pure.*

The analog of Lemma 4.0.8 is:

Lemma 5.2.5. *For each $m \in \mathbb{N}$, let Ξ_m denote the set of states which are not m-pure. Then $\nu(\Xi_m) = 0$ for all $m \in \mathbb{N}$.*

The proof could also be given along the lines of the alternative proof for Lemma 4.0.8 given in Section 7.2. We suffice with mimicking the shorter proof of Lemma 4.0.8 given in Section 4.

Proof. Fix m; Ξ_m, Ξ_{m+1} are assumed ν-measurable. To begin with, note that Lemma 5.2.3 implies that

$$S(\Xi_m) \subseteq \Xi_{m+1} \tag{5.8}$$

Let ι be as in the proof of Lemma 4.0.8. Lemma 5.2.4 then implies that

$$\Xi_m \cap \iota(\Xi_m) = \emptyset \tag{5.9}$$

The ν-invariance of ι and S and the fact that $S(\Xi_m)$ is ν-measurable,[36] together with (5.8), (5.9), and (4.3), imply that

$$2\nu(\Xi_m) = \nu(\Xi_m) + \nu(\iota(\Xi_m)) = \nu(S^{-1}(S(\Xi_m))) = \nu(S(\Xi_m)) \leq \nu(\Xi_{m+1})$$

Inductively, we see that

$$2^k \cdot \nu(\Xi_m) \leq \nu(\Xi_{m+k}) \leq 1,\ \forall k, m \in \mathbb{N}$$

and hence $\nu(\Xi_m) = 0$. □

For each $m \in \mathbb{N}$, define the map $g_m : Z \to \{-1, 1\}$ by $g_m(z) = 1$ if and only if z is (L, m)-pure. Denote for all $y, y' \in Y$, $D(y, y') = \{j \in \omega \mid y_j \neq y'_j\}$, and if $D(y, y')$ is finite, $N(y, y') = \#D(y, y')$, $M(y, y') = \max D(y, y')$.

[36] See the footnotes in the proof of Lemma 4.0.8

Lemma 5.2.6. *For each $n \in \omega_M$, μ-almost every $y \in Y$, we have*

$$g_1(n, y) = (-1)^{N(y,y')} g_1(n, y'),\ \forall y' \in Y \ s.t.\ N(y, y') < \infty \tag{5.10}$$

Proof. By Lemma 5.2.3 and Lemma 5.2.5, we see that for almost every $z = (n, y) \in Z$ and all $m \in \mathbb{N}$, $g_m(z) = y_0 \cdot g_{m+1}(S(z))$, and hence for all k and a.e. z,

$$g_1(z) = y_0 \cdot \dots \cdot y_{k-1} \cdot g_{k+1}(S^k(z))$$

and the proof follows as in the proof of Lemma 4.0.9. □

By Theorem 4.0.10, such g_1 cannot be measurable, which completes our contradiction.

Now, let $\Gamma' = (\Omega, \mathcal{P}, \{L, R\}^{\mathcal{P}}, \beta', r', q')$ be a perturbation of Γ, satisfying the conditions of Section 2.2 with ε given in Section 3.2, and which also satisfies that

$$q(z' = S(z) \mid z, a) > 0,\ \forall z \in Z, a \in \{L, R\}^{\mathcal{P}} \tag{5.11}$$

In this case, in Γ', (5.2) is satisfied for all z, m and one uses this to show that Γ' does not possess Markovian equilibria via the same techniques we have used in this section.

6 Final Remarks and Acknowledgements

6.1 Nonmeasurable Stationary Equilibria

Let $(\Omega, \mathcal{P}, (I^p), \beta, r, q)$ be a stochastic game with purely atomic transitions; that is, for each state $z \in \Omega$ and each action profile $a \in I^{\mathcal{P}}$, $q(z, a)$ is purely atomic. Note that the game defined in Section 3.3 has purely atomic transitions. Define the following relationship on Ω: zEz' if for some $a \in I^{\mathcal{P}}$, $q(z' \mid z, a) > 0$. Let E^* be the reflexive and transitive closure of this relationship.

If σ is a stationary strategy that is not necessarily measurable in a game with purely atomic transitions, then for each $z \in \Omega$, the right-hand side of (2.1) is defined, as it depends only on the countable set of action profiles $(\sigma(z'))_{\{z' \mid zE^*z'\}}$. This brings us to the following result:

Proposition 6.1.1. *Every game with purely atomic transitions has stationary (possibly nonmeasurable) equilibria.*

Proof. The axiom of choice implies the well-ordering theorem: Every set has a well-ordering. Let $<$ be a well-ordering of Ω, and define σ using transfinite recursion: assuming we've already defined σ on $S \subseteq \Omega$, let z be the $<$-least element of $\Omega \backslash S$. Denote

$$E_z = \{z' \in \Omega \mid zE^*z'\},\ G_z = \{z' \in S \mid zE^*z'\},\ F_z = E_z \backslash G_z$$

We define σ on F_z by viewing the stochastic game Γ_z on the set of states E_z, in which the players at states $z' \in G_z$ are constrained to playing[37] $\sigma(z')$, which has already been defined. Since stochastic games with countable state spaces possess stationary equilibria (see Parthasarathy (1973) or Federgruen (1978)) we get a stationary strategy τ for Γ_z. Setting $\sigma(z') = \tau(z')$ for all $z' \in F_z$ (which includes z) completes the induction step. It's easy to show that the resulting strategy σ is, indeed, an equilibrium. □

As such, there are non-Lebesgue-measurable stationary strategies of the game defined in Section 3.3. However, nonmeasurable strategies possess severe disadvantages. Their existence cannot be proven without using the axiom of choice; indeed, it is possible to assume, in addition to the standard Zarmelo-Fraenkel axioms of set theory, axioms under which non-Lebesgue-measurable sets and functions do not exist. In particular, there is no way to explicitly define such strategies. For a good survey explaining and emphasizing these points in a game-theoretical framework, we recommend Section 4 of Zame (2007). A more thorough set-theoretical treatment can be found in, e.g., Section 7 of Herrlich (2006) and the references there.

6.2 Bayesian Games

Simon (2003) presents an example of a Bayesian game with a continuum of states of the world that possess no (measurable) Bayesian equilibrium. Although there does not appear to be a formal relationship between his example and ours, his paper nonetheless served as an inspiration for this and other works. In Simon's example, the space of states of the world is $X = 2^{\mathbb{Z}}$; the common knowledge components are those generated by the shift and the involution given by $(\xi(x))_n = x_{-n}$, and there are three players, each with two actions.

However, ε-Bayesian equilibria do exist in his example (Simon, private communication). Perhaps some of the techniques we have demonstrated here can give a stronger example, one in which even approximate equilibria do not exist.

6.3 Acknowledgements

I am grateful to Abraham Neyman for going over the drafts of this paper and providing countless corrections and suggestions. I wish to thank Eilon Solan for pointing out the work of Bob Simon (2003) to me, and to thank Bob Simon for communications we had. I also thank Ilan Nehama for many helpful conversations.

[37] That is, in a state $z' \in G_z$ in Γ_z, the players' actions have no effect; the transitions and payoffs are as if they had played $\sigma(z')$.

7 Appendix

7.1 Bounding Payoff Differences

Proposition 7.1.1. *Let $\rho > 0$, $0 < \varepsilon \leq 1$. Suppose that $(\Omega, \mathcal{P}, (I^p), \beta, r, q)$, $(\Omega, \mathcal{P}, (I^p), \beta', r', q')$ are two games in the model of Section 2.1 with the same state / player / action spaces, such that*

$$||r'(z,a) - r(z,a)||_\infty < \varepsilon,\ ||q'(z,a) - q(z,a)|| < \varepsilon,\ \forall z \in \Omega,\ \forall a \in \overline{I}$$

and

$$||r||_\infty \leq \rho,\ |\beta' - \beta| < \varepsilon,\ |\frac{1}{1-\beta} - \frac{1}{1-\beta'}| \leq \frac{1}{1-\beta}$$

Let σ be any behavioral strategy profile, and $\gamma_\sigma(z), \gamma'_\sigma(z)$ denote the expected payoffs in these games starting with state z. Then we have

$$||\gamma_\sigma - \gamma'_\sigma||_\infty \leq \frac{\varepsilon \cdot D}{1-\beta}$$

where $D = D(\beta, \rho) := 1 + \frac{2(\beta+1)(\rho+1)}{1-\beta}$.

For notational convenience, we will prove the proposition only in the case where σ is a stationary strategy profile; the general case follows either by making some slight changes, or by observing an auxiliary game in which the state space is the space of all finite histories of the original game.

Proof. Let $\gamma^n_\sigma(z), \gamma'^n_\sigma(z)$ denote the payoffs in these games when only the payoffs through stage n are summed. Clearly, $\gamma^n_\sigma(\cdot) \to \gamma_\sigma(\cdot)$ uniformly, and similarly for $\gamma'_\sigma(\cdot)$. Hence, it's enough to prove that for all $n \in \omega$ (we drop explicit reference to σ),

$$||\gamma^n - \gamma'^n||_\infty \leq D\varepsilon(\sum_{k=0}^{n} \beta^k)$$

This is clear for $n = 1$, as $D \geq 1$. Note that $||r'||_\infty \leq \rho + \varepsilon$, and hence for all n,

$$||\gamma'^n||_\infty \leq \frac{\rho + \varepsilon}{1-\beta'} \leq (\rho+1)\frac{2}{1-\beta}$$

Also, observe that

$$\gamma^{n+1}(z) = r(z, \sigma(z)) + \beta \int_\Omega \gamma^n(t) dq(z, \sigma(z))(dt)$$

and similarly for γ'^n. This leads to

$$\begin{aligned}
||\gamma^{n+1} - \gamma'^{n+1}||_\infty &\leq ||r - r'||_\infty + |\beta - \beta'| \cdot ||\gamma'^n||_\infty \\
&+ \beta(||q - q'||_\infty \cdot ||\gamma'^n||_\infty + ||\gamma^n - \gamma'^n||_\infty) \\
&\leq \varepsilon + \varepsilon \cdot (\rho+1)\frac{2}{1-\beta} + \beta\varepsilon(\rho+1)\frac{2}{1-\beta} + \beta||\gamma^n - \gamma'^n||_\infty \\
&= D\varepsilon + \beta||\gamma^n - \gamma'^n||_\infty
\end{aligned}$$

This calculation then gives the induction step. □

7.2 Alternative Proof of Lemma 4.0.8

First, we introduce notation. For $N \in \mathbb{N}$, let

$$2^N = \{u \in 2^* \mid \forall n \geq N,\ u_n = 1\}$$

and let $2^* \subseteq Y = \{-1,1\}^\omega$ be $2^* = \cup_{N\in\mathbb{N}} 2^N$. We define an action of 2^* on Y by

$$u \cdot y = (u_0 y_0, u_1 y_1, u_2 y_2, \ldots),\ \forall u \in 2^*, y \in Y$$

Since $2^* \subseteq Y$, it's easy to see that 2^* acts on itself, inducing a group action. We also denote $e = (1,1,1,1\ldots)$, and for $y \in 2^N$, we identify y with its image

$$(y_0, \ldots, y_{N-1}, 1, 1, 1, \ldots) \in 2^*$$

which we sometimes denote (y, e).

Now, we can get to the proof. We claim that for all $N \in \mathbb{N}$, it holds for all $u \in 2_0^N := 2^N \backslash \{e\}$ and all $z \in \Xi$, that $u \cdot z \notin \Xi$, where $u \cdot (n, y) := (n, u \cdot y)$. We do this by induction on N: for $N = 1$, if $z \in \Xi$, it follows that $(-1, e) \cdot z \notin \Xi$ from Lemma 3.4.7, since $z, (-1, e) \cdot z \in S^{-1}(S(z))$ and z is not quasi-pure.

Assuming that we have proved the claim for N, let us prove it for $N + 1$. Let $z \in \Xi$. By (4.1), $S(z), S^{N+1}(z) \in \Xi$. The induction hypothesis tells us that the set

$$S^{-N}(S^{N+1}(z)) \backslash \{S(z)\} = \{u \cdot S(z) \mid u \in 2_0^N\}$$

is disjoint from Ξ. Therefore, it follows from Lemma 3.4.6 that for all $u \in 2_0^N$, neither element of $S^{-1}(u \cdot S(z))$ - that is, neither of the two elements $(1, u) \cdot z$ and $(-1, u) \cdot z$ - are in Ξ. To conclude the induction step, it only remains to show that $(-1, e) \cdot z \notin \Xi$; but this follows from the case $N = 1$.

Now, fix N. For each $u \in 2_0^N$, let $\Xi^u = \{u \cdot z \mid z \in \Xi\}$. The argument above shows that for $u \in 2_0^N$, $\Xi^u \cap \Xi = \emptyset$. If $u \neq v \in 2^N$, then we contend that $\Xi^u \cap \Xi^v = \emptyset$; for if we had $z \in \Xi$ with $u \cdot z = v \cdot z$, then denoting $w \in 2^N$ via $w_j = \frac{u_j}{v_j}$, we would have $w \neq e_N$ but $w \cdot z = z \in \Xi$, a contradiction.

Clearly, for all $u \in 2_0^N$, $\nu(\Xi^u) = \nu(\Xi)$. Hence, denoting $\Xi^N = \cup_{u \in 2_0^N} \Xi^u$, we have

$$1 \geq \nu(\Xi_N) = (2^N - 1)\nu(\Xi)$$

and this is for any N; hence $\nu(\Xi) = 0$.

References

[1] Amir, R. (1991) On stochastic games with uncountable state and action spaces, in *Stochastic Games and Related Topics*, (T. E. S. Ragahavan, T. S. Ferguson, T. Parthasarathy, and O. J. Vrieze, Eds.), Boston, Kluwer Academic, pp. 149–159.

[2] Amir, R. (1996), Continuous Stochastic Games of Capital Accumulation with Convex Transitions, *Games and Economic Behavior*, 15, 111–131.

[3] Amir, R. (2003), Stochastic Games in Economics and Related Fields: An Overview, in *Stochastic Games and Applications, NATO Science Series C, Mathematical and Physical Sciences* (A. Neyman and S. Sorin, Eds.), Vol. 570, Dordrecht, Kluwer Academic Publishers, Chapter 30, pp. 455-470.

[4] Basar, T. and Olsder, G. (1999), *Dynamic Noncooperative Game Theory*, Philadelphia, SIAM Classics.

[5] Beggs, A. and Klemperer, P. (1992), Multi-Period Competition with Switching Costs, *Econometrica*, 60, 651-666.

[6] Bertsekas, D. and Shreve, S. (1996) *Stochastic Optimal Control: The Discrete Time Case*, Belmont, MA, Athena Scientific.

[7] Chakrabarti, S.K. (1999), Markov Equilibria in Discounted Stochastic Games. *Journal of Economic Theory*, 85, 294-327.

[8] Duffie, D., Geanakoplos, J., Mas-Colell, A., and McLennan, A. (1994), Stationary Markov Equilibria, *Econometrica*, 62, 745–781.

[9] Duggan, J. Noisy Stochastic Games, Pre-Print.

[10] Dutta, P. and Sundaram, R. (1998), The Equilibrium Existence Problem in General Markovian Games, in *Organizations with Incomplete Information* (M. Majumdar, Ed.), Cambridge, Cambridge University Press.

[11] Federgruen, A., On N-Person Stochastic Games with Denumerable State Space, *Advances in Applied Probability*, 10, 452–471.

[12] Fink, A.M. (1964), Equilibrium in Stochastic n-Person Games, *Journal of Science of the Hiroshima University*, 28, 89–93.

[13] Forges, F. (1986), An Approach to Communication Equilibria, *Econometrica*, 54, 1375–1385.

[14] Guesnerier, R., and Woodford, M. (1992), Endogenous Fluctuations, *in Advances in Economic Theory*, Vol. II, (J.J. Laffont, Ed.), Cambridge, Cambridge University Press.

[15] Halmos, P. (1956), *Lectures on Ergodic Theory*, New York, Chelsea Publishing Company.

[16] Harris, C., Reny, P., Robsin, A. (1995), The Existence of Subgame-Perfect Equilibrium in Continuous Games with Almost Perfect Information: A Case for Public Randomization, *Econometrica*, 63, 507–544.

[17] Harsanyi,J.C., and Selten, R. (1988), *A General Theory of Equilibrium Selection in Games*, MIT Press.

[18] Hellwig, M. and Leininger, W. (1988), Markov-perfect Equilibrium in Games of Perfect Information, Discussion Paper A-183, University of Bonn.

[19] Herrlich, H. (2006), *Axiom of Choice*. Lecture Notes in Mathematics, New York, Springer.

[20] Himmelberg, C.J., Parthasarathy, T., Raghavan, T.E.S. and Van Vleck, F.S. (1976) , Existence of p-Equilibrium and Optimal Stationary Strategies in Stochastic Games, *Proceedings of the American Mathematical Society*, 60, 245-251.

[21] Levhari, D. and Mirman, L. (1980), The Great Fish War: An Example Using a Dynamic Cournot-Nash Solution, *Bell Journal of Economics* 11, 322-334.

[22] Maitra, A. and Sudderth, W. (2007), Subgame-Perfect Equilibria for Stochastic Games, *Mathematics of Operations Research*, 32, 711–722.

[23] Maitra, A. and Parthasarathy, T. (1970), On Stochastic Games, *Journal of Optimization Theory and Applications*, 5, 289-300.

[24] Maskin, E. and Tirole, J. (2001), Markov Perfect Equilibrium: I. Observable Actions, *Journal of Economic Theory*, 100, 191–219.

[25] Mertens, J.-F. and Parthasarathy, T.E.S. (1987) Equilibria for Discounted Stochastic Games, Paper 8750, Université Catholique de Louvain, Louvain-la-Neuve, Belgium.

[26] Monahan, G. and Sobel, M. (1994), Stochastic Dynamic Market Share Attraction Games, *Games and Economic Behavior*, 6, 130-149.

[27] Nowak, A.S. (1985a), Universally Measurable Strategies in Zero-Sum Stochastic Games, *The Annals of Probability*, 13, 269–287.

[28] Nowak, A.S. (1985b), Existence of Equilibrium Stationary Strategies in Discounted Noncooperative Stochastic Games with Uncountable State Space, *Journal of Optimization Theory and Applications*, 54, 591-602.

[29] Nowak, A. (2003), N-Person Stochastic Games: Extensions of the Finite State Space and Correlation, in *Stochastic Games and Applications, NATO Science Series C, Mathematical and Physical Sciences* (A. Neyman and S. Sorin, Eds.), Vol. 570, Dordrecht, Kluwer Academic Publishers, Chapter 8, pp. 93-106.

[30] Nowak, A.S. (2002) On a New Class of Non-Zero-Sum Discounted Stochastic Games having Stationary Nash Equilibrium Points, *International Journal of Game Theory*, 32, 121–132.

[31] Nowak, A.S. (2007) Stochastic Games in Economics, *Mathematical Methods of Operations Research*, 66, 513–530.

[32] Nowak, A.S. and Raghavan, T.E.S. (1992), Existence of Stationary Correlated Equilibria with Symmetric Information for Discounted Stochastic Games, *Mathematics of Operations Research* , 17, 519-526.

[33] Parthasarathy, T. (1973), Discounted, Positive and Noncooperative Stochastic Games, *International Journal of Game Theory*, 2, 25-37.

[34] Parthasarathy, T. and Sinha, S. (1989) Existence of Stationary Equilibrium Strategies in Non-Zero-Sum Discounted Games with Uncountable State Space and State–Independent Transitions, *International Journal of Game Theory*, 18, 189–194.

[35] Rieder, U. (1979), Equilibrium Plans for Non-Zero-Sum Markov Games, in *Game Theory and Related Topics* (O. Moeschlin and D. Pallaschke, Eds.), Amsterdam, North Holland, pp. 91-102

[36] Rogers, P.D. (1969), Nonzero-Sum Stochastic Games. *Report ORC 69–8*, Operations Research Center, University of California, Berkeley.

[37] Rosenthal, R.W. (1982), A Dynamic Model of Duopoly with Consumer Loyalties, *Journal of Economic Theory*, 27, 69-76.

[38] Shapley, L. (1953), Stochastic Games. *Proceedings of the National Academy of Sciences, USA* 39, 1095–1100.

[39] Simon, R. (2003), Games of Incomplete Information, Ergodic Theory, and the Measurability of Equilibria. *Israel Journal of Mathematics*, 138, 1, 73–92.

[40] Shiryaev, A. (1996), *Probability Theory* (2nd ed.), New York, Springer.

[41] Sobel, M. (1971), Noncooperative Stochastic Games. *Annals of Mathematical Statistics*, 42, 1930–1935.

[42] Sobel, M. (1973), Continuous Stochastic Games. *Journal of Applied Probability*, 10, 597–604.

[43] Solan, E. (1998), Discounted Stochastic Games, *Mathematics of Operations Research*, 23, 1010-1021.

[44] Takahashi, M. (1964), Equilibrium Points of Stochastic Non-Cooperative n–Person Games, *Journal of Science of the Hiroshima University*, 28, 95-99.

[45] Whitt, W. (1980), Representation and Approximation of Noncooperative Sequential Games, *SIAM Journal of Control and Optimization*, 18, 33-48.

[46] Zame, W. (2007), Can Intergenerational Equity be Operationalized?, *Theoretical Economics*, 2, 187–202.