

Games of Incomplete Information, Ergodic Theory, and the Measurability of Equilibria

Robert Samuel Simon

This research was supported by the Center for Rationality and Interactive Decision Theory (Jerusalem), the Department of Mathematics of the Hebrew University (Jerusalem), the Edmund Landau Center for Research in Mathematical Analysis (Jerusalem) sponsored by the Minerva Foundation (Germany), the German Science Foundation (Deutsche Forschungsgemeinschaft), and the Institute of Mathematical Stochastics (Goettingen).

Abstract: We present an example of a one-stage three player game of incomplete information played on a sequence space $\{0,1\}^{\mathbf{Z}}$ such that the players' locally finite beliefs are conditional probabilities of the canonical Bernoulli distribution on $\{0,1\}^{\mathbf{Z}}$, each player has only two moves, the payoff matrix is determined by the 0-coordinate and all three players know that part of the payoff matrix pertaining to their own payoffs. For this example there are many equilibria (assuming the axiom of choice) but none that involve measurable selections of behavior by the players. By measurable we mean with respect to the completion of the canonical probability measure, e.g. all subsets of outer measure zero are measurable. This example demonstrates that the existence of equilibria is as much a philosophical issue as a mathematical one. We consider the double-shift Bernoulli probability space $B^{\mathbf{Z}^2}$, where T_1 is the shift in the first coordinate, T_2 is the shift in the second coordinate, and $x^{i,j}$ is the i, j -coordinate of $x \in B^{\mathbf{Z}^2}$. Let C be a compact and convex set with compact subsets $(A_b \mid b \in B)$ indexed by the set B such that $\bigcap_{b \in B} A_b = \emptyset$. We conjecture that measurable functions $f : B^{\mathbf{Z}^2} \rightarrow C$ can not keep $\frac{1}{4}(f(x) + f(T_1(x)) + f(T_2(x)) + f(T_1 \circ T_2(x)))$ in $A_{x,0,0}$ for all $x \in B^{\mathbf{Z}^2}$ and that the inability of measurable functions to satisfy this property (in expectation) is bounded below by a positive constant dependent on the sets $(A_b \mid b \in B)$. We give an example of a one-stage zero-sum game played on $B^{\mathbf{Z}^2}$ that would not have a value (but would have equilibria!) if this conjecture were valid.

Key words: Bayesian Equilibria, Belief Spaces, Sequence Spaces, Non-measurable Sets

1 Introduction

An equilibrium of a game is a set of strategies, one for each player, such that no player does better by choosing a different strategy, given that the other players do not change their strategies. In a game of incomplete information, what does it mean to do no better by choosing a different strategy? Should one evaluate a player's actions according to the subjective and local beliefs of that player, or should one evaluate according to a probability distribution determined objectively by the game?

When the subjective beliefs of the players are conditionals of a common prior, a central question is whether there is a difference between the equilibria defined according to the subjective local beliefs and the equilibria defined according to a global functional evaluation. The former we call *Bayesian* equilibria and the latter we call *Harsanyi* equilibria. Harsanyi (1967-8) showed that these equilibria are equivalent when the set of all possible situations in the game is finite. Is there such an equivalence when the possible situations of the game are infinitely many?

In this paper we look at games that satisfy the following conditions:

- 1) there is one stage of play,
- 2) there are finitely many players,
- 3) each player has finitely many moves,
- 4) there are finitely many payoff matrices,
- 5) there is a compact space Ω with a Borel probability distribution μ and a finite partition into clopen subsets corresponding to the different payoff matrices; (the payoffs at $x \in \Omega$ are determined by the matrix associated with the partition member containing x),
- 6) at every point $x \in \Omega$ every player j has a discrete probability distribution on Ω with a finite support set $S^j(x)$ containing x such that at all the other points in this finite support set $S^j(x)$ the player j has the same discrete distribution,
- 7) these discrete beliefs of the players change continuously (with respect to the weak topology), and for any player j they are regular conditional probability distributions of μ with respect to the sigma algebra $\mathcal{F}^j := \{B \mid B \text{ is Borel and } x \in B \Leftrightarrow S^j(x) \subseteq B\}$.

Such games we call *ergodic* games.

A strategy for a player in an ergodic game is a function from Ω to the probability simplex of his moves that is constant on every finite support set

that he could believe to be possible. We call a function on Ω *measurable* when it is measurable with respect to the completion of μ , meaning that all sets of outer measure zero are measurable.

We give an example of a three player ergodic game that has no Bayesian equilibrium in measurable strategies (meaning also no Harsanyi equilibrium), yet it has many Bayesian equilibria that are not measurable. This example has the additional property that each player knows that part of the payoff matrix that pertains to his own payoff.

In the context of most multi-agent epistemic logics, a statement concerning the knowledge or belief of a player will correspond to a measurable subset of our sequence space. When applicable, the lack of a measurable equilibrium implies the impossibility for the behavior of the players in equilibrium to be determined by syntactic formulations of knowledge or belief.

Because the existence of non-measurable sets in our probability space can be denied by rejecting the axiom of choice, the existence of Bayesian equilibria for this game could be considered to be a philosophical question.

Do zero-sum ergodic games have values?

The question is difficult to answer because one must define the concept of value for ergodic games. A zero-sum game is a two-person game such that for all possible payoff matrices and all combinations of moves the payoff for one player is the negation of the payoff for the other player. Usually, a zero-sum game is defined to have a value $r \in \mathbf{R}$ (as a payoff for the first player) when for every positive ϵ the first player has a strategy such that no matter what the second player does it guarantees to him an expected payoff of at least $r - \epsilon$, and vice-versa, the second player has a strategy such that no matter what the first player does the payoff to the first player is held down to $r + \epsilon$ or less. The problem with this definition of value is how to define the expected payoff of an ergodic game. If the strategies are measurable, then a payoff can be defined as an expectation over the probability space. But if strategies are not measurable, meaning that the evaluations of the players are strictly local in character, how should we define the expected payoff? Indeed, we will show an example of a Bayesian equilibrium of a zero-sum ergodic game that has an expected payoff for one player but not for the other player. We believe that the only reasonable definition of a value for a zero-sum ergodic game should be in relation to the set of measurable strategies.

We present an example of a zero-sum ergodic game and a conjecture of ergodic theory whose affirmation would show that this game would have

no value in measurable strategies, and yet this game would have Bayesian equilibria.

The conjecture is the following. Let $\Omega = B^{\mathbf{Z}^2}$ be the double-shift space over a finite set B , with the canonical Bernoulli probability distribution μ that gives equal probability to every $b \in B$ in every coordinate position. Let T_1 and T_2 be the measure preserving transformations on Ω associated with shifting the first and second coordinates, respectively. Let C be a convex and compact subset of an Euclidean space with compact subsets $(A_b \mid b \in B)$ such that $\bigcap_{b \in B} A_b = \emptyset$. For every measurable function $f : \Omega \rightarrow C$ let $d_f : \Omega \rightarrow \mathbf{R}$ be the non-negative function defined by $d_f(x)$ being the distance between $\frac{1}{4}(f(x) + f(T_1(x)) + f(T_2(x)) + f(T_1 \circ T_2(x)))$ and the set $A_{x^{0,0}}$, where $x^{0,0} \in B$ is the $0, 0$ coordinate of $x \in \Omega$.

Conjecture: There a positive value $w > 0$ (dependent on the choice of the $(A_b \mid b \in B)$) such that for every choice of measurable function $f : \Omega \rightarrow C$ the expected value of d_f over Ω must exceed w .

Due to Luzin's theorem, one could replace in the conditions for the conjecture the measurability of f with the continuity of f .

Define the measure preserving involutions $\tau_1, \tau_2 : \Omega \rightarrow \Omega$ by $(\tau_1(y))^{i,j} := y^{1-i,j}$ and $(\tau_2(y))^{i,j} := y^{i,1-j}$. For our game theoretic purposes, it would suffice if the conjecture were true with the additional assumption that the function f is τ_1 and τ_2 invariant, meaning for all $x \in \Omega$ that $f(\tau_1(x)) = f(x)$ and $f(\tau_2(x)) = f(x)$.

As we shall see, the conjecture is false if formulated for the usual single-shift space $\{0, 1\}^{\mathbf{Z}}$ (with d_f defined with respect to the distance between $\frac{1}{2}(f(x) + f(T(x)))$ and the appropriate compact subset).

The basic ideas in this paper belong to ergodic theory. One chooses a sequence space that allows for finitely many measure preserving involutions σ_i whose orbits are almost everywhere dense in the space. ($\sigma_i \circ \sigma_i$ is the identity, and for almost every x the subset of all the $\sigma_{i_1} \circ \sigma_{i_2} \circ \dots \circ \sigma_{i_n}(x)$ is dense in the space.) To each player is associated a finite subset of involutions that commute with each other, and for every point x , this player believes that only the points in the finite orbit of his commuting involutions are possible. The entries in the payoff matrix are determined by the topological position in the sequence space, in particular, by some coordinate position. The conditions defining Bayesian equilibria pertain to the orbits of the involutions σ_i , and these orbits do not contract the space in a measurable way.

In the next section we define Bayesian and Harsanyi equilibria, prove that the countability of a belief structure implies the existence of Bayesian equilibria regardless of whether there exists a common prior, and we show that a Harsanyi equilibrium for an ergodic game generates a measurable Bayesian equilibrium for that game. In the third section we present the example of an ergodic game without measurable Bayesian equilibria. In the fourth section we discuss the conjecture and present an example of a zero-sum ergodic game that would have no value in measurable strategies if the conjecture were true. In conclusion, we explore fundamental aspects of zero-sum games.

2 Bayesian and Harsanyi Equilibria

$\Delta(A)$ will stand for the set of regular Borel probability distributions of A , where A is a topological space (and if A is finite then we give A the discrete topology and $\Delta(A)$ is a simplex embedded in an Euclidean space). The distance in an Euclidean space, including in a simplex embedded in an Euclidean space, will be the Euclidean distance. Throughout this paper we will assume the axiom of choice. Sometimes a player will be referred to as *he* and sometimes as *she*.

2.1 Mertens-Zamir spaces

Ergodic games are a special case of games played on probability spaces that satisfy the Mertens-Zamir definition for a *belief space*.

A Mertens-Zamir belief space (Mertens and Zamir, 1985) is a tuple $(S, X, \psi, N, (t^j \mid j \in N))$, where X is a compact parameter set, S is a compact set, ψ is a continuous map from S to X , N is a finite set of players, for every $j \in N$ $t^j : S \rightarrow \Delta(S)$ is a continuous function (with respect to the weak topology), and for every player j and every pair of points $s, s' \in S$ if $s' \in \text{support}(t^j(s))$ then $t^j(s) = t^j(s')$.

Define a *cell* of a Mertens-Zamir belief space to be a minimal set C with the property that at every point y in C every player's support set for the point y is contained in C (without the requirement that C must be compact).

Of special interest is the definition of mutual consistency for Mertens-Zamir belief spaces. For every player $j \in N$ and define \mathcal{T}^j to be the smallest

Borel field of subsets of S such that the function t^j is measurable. A probability distribution μ on S is defined to be *consistent* if for every Borel subset $A \subseteq S$ we have that $\mu(A) = \int t^j(y)(A) d\mu(y)$. Mertens and Zamir (1985) showed that consistency is equivalent to the stronger statement that for every $B \in \mathcal{T}^j$ and Borel subset $A \subseteq S$ we have $\mu(A \cap B) = \int_B t^j(y)(A) d\mu(y)$. We prefer to use the Borel field \mathcal{F}^j defined by $\mathcal{F}^j := \{B \mid B \text{ is Borel and } x \in B \Leftrightarrow \text{support}(t^j(x)) \subseteq B\}$, the largest Borel field on which a Borel measurable strategy for Player j can be defined. For the relation between the cells of a Mertens-Zamir belief space and the common prior, see Simon (2000).

For Mertens-Zamir belief spaces with a finite parameter space X it is easy to construct games that are played on the space S . Let $\psi : S \rightarrow X$ be the continuous function, N the finite player set, and for each $j \in N$ $t^j : S \rightarrow \Delta(S)$ the subjective beliefs of Player j . For each player j , there is a finite action set A^j with $n^j := |A^j|$. There are $|X|$ different $n^1 \times \dots \times n^{|N|}$ matrices $(Q_x \mid x \in X)$; every entry of every matrix is a vector payoff for the players in \mathbf{R}^N . Nature chooses a point in S according to the common prior μ , which means also that a parameter in X is chosen through the function ψ . The players choose moves in their respective A^j independently, and after the choices are made the payoff to the players is the vector entry in Q_x corresponding to nature's choice of the parameter in X and the moves of the players. A Bayesian equilibrium for a point z in the belief space is an $|N|$ -set of functions $(f^j \mid j \in N)$, each f^j from the cell that contains z to $\Delta(A^j)$, the simplex of mixed strategies, with the following properties for every player $j \in N$

- 1) f^j is constant within all support sets of Player j ,
- 2) for all $j' \neq j$ within the support set of $t^{j'}(z)$ the function f^j is $t^{j'}(z)$ measurable, and
- 3) within the support set of $t^j(z)$ Player j can do no better than $f^j(z) \in \Delta(A^j)$ in response to the other functions $f^{j'}, j' \neq j$, as evaluated by $t^j(z)$.

When the $|N|$ -set of functions is a Bayesian equilibrium for all points in a cell, then we call it a *cellular* equilibrium. A Bayesian equilibrium for the whole space is a collection of *cellular* equilibria, once for each cell.

If additionally the Mertens-Zamir belief space has a consistent common prior μ , we define a Harsanyi equilibrium to be a set of functions $(f^j : S \rightarrow \Delta(A^j) \mid j \in N)$, each f^j *measurable* with respect to the Borel field \mathcal{F}^j , such that no player can attain a higher expected payoff as evaluated by μ by

choosing another such measurable function, (given that the strategies of the other players do not change).

Notice that ergodic games are special cases of games played on Mertens-Zamir spaces with consistent priors. From now on one will use the terminology of Mertens-Zamir spaces to describe an ergodic game, in particular the t^j to represent the belief of Player j ; we will write support $(t^j(x))$ instead of $S^j(x)$. We will adopt the term *cell*, though for all our examples a cell will be an orbit of a non-abelian group acting on the probability space. (A confusion with the orbits of a shift transformation T is avoided.) In an ergodic game, the local measurability condition 2) of a Bayesian equilibrium disappears, leaving Bayesian equilibria with no measurability conditions at all.

2.2 Locally Finite Games

Now we define *locally finite* games, a class including the ergodic games, but without the assumption of a common prior or a compact space.

Let S be a finite or countably infinite set. There is a finite or countably infinite collection \mathcal{J} and a map $\nu : \mathcal{J} \rightarrow 2^S$ such that for every $J \in \mathcal{J}$ $\nu(J)$ is a subset of S and for every $s \in S$ the point s is contained in $\nu(J)$ for only finitely many members J of \mathcal{J} . A member of \mathcal{J} is a player of our game. For every $J \in \mathcal{J}$ there is a finite set A^J of moves and a sigma-additive probability distribution p^J on the set $\nu(J)$ such that the support of p^J is the set $\nu(J)$. For every $s \in S$ define $\mathcal{J}(s)$ to be the finite subset of \mathcal{J} whose projections by ν contain the point s . For every $s \in S$ there is a payoff matrix Q_s of size $\times_{J \in \mathcal{J}(s)} |A^J|$ with entries in $\mathbf{R}^{\mathcal{J}(s)}$; for every choice of $a \in \prod_{J \in \mathcal{J}(s)} A^J$ there are corresponding payoffs for all the players in $\mathcal{J}(s)$.

A strategy of a player $J \in \mathcal{J}$ is a member of $\Delta(A^J)$. A Bayesian equilibrium for a locally finite game is a set of strategies $(f^J \in \Delta(A^J) \mid J \in \mathcal{J})$, such that for every $J \in \mathcal{J}$ the player J cannot get a higher expected payoff by choosing a move in A^J not in the support of f^J , given that the other players remain with $(f^K \in \Delta(A^K) \mid K \neq J, K \in \mathcal{J})$, with the expected payoff of a move calculated according to the distribution p^J and the expected payoffs from the matrices Q_s at each of the states $s \in \nu(J)$,

Proposition 1: If for every player $J \in \mathcal{J}$ in a locally finite game there is a uniform bound M^J on the absolute value of all payoffs to J in the set $\nu(J)$, then there exists a Bayesian equilibrium.

Proof: Let the set S be enumerated by $S = \{s_1, s_2, s_3, \dots\}$. For every $i = 1, 2, \dots$ define the game Γ_i to be that played on the set $S_i := \{s_1, \dots, s_i\}$ with the players $\mathcal{J}_i := \{J \in \mathcal{J} \mid \nu(J) \cap S_i \neq \emptyset\}$ and the probability distributions p_i^J on S_i (for the players in \mathcal{J}_i) induced by $p_i^J(B) := p^J(B)/p^J(S_i)$ for all $B \subseteq S_i$. Since each S_i and \mathcal{J}_i are finite, there exists a Nash equilibrium $\tilde{f}_i := (f_i^J \in \Delta(A^J) \mid J \in \mathcal{J}_i)$ to the game Γ_i for every $i \geq 1$ (Nash, 1950). Assigning any distribution to f_i^J when $\nu(J)$ has an empty intersection with S_i , we have a sequence f_i in the set

$$\tilde{A} := \prod_{J \in \mathcal{J}} \Delta(A^J).$$

We give \tilde{A} the product topology. Due to Tychanov's Theorem, \tilde{A} is compact, and we can assume that there exists a convergent subsequence f_{i_n} of the f_i converging to $f \in \tilde{A}$. Redefine the sequence $(f_i \mid i = 1, 2, \dots)$ so that the new f_n is the old f_{i_n} .

We aim to show that f defines a Bayesian equilibrium of the original game. Fix an $\epsilon > 0$. We will show that f is an ϵ -Bayesian equilibrium.

Let Player $J \in \mathcal{J}$ be given. Let i_0 be so large that $p^J(S_{i_0}) \geq 1 - \epsilon/4M^J$. Let δ be the smallest positive probability by which Player J chooses some move with the strategy f^J . Let N_i be the finite cardinality of the set \mathcal{J}_i . Now choose an $i_1 > i_0$ so large that $l > i_1$ implies that f_l^K is within $\min(\epsilon/4M^J N_{i_0}, \delta/2)$ of f^K for all the $K \in \mathcal{J}_i$, including J . Given that all other players $K \neq J$ stay with their strategies f^K , we need only show that the differences in expected payoffs for Player J by choosing different moves in support (f^J) do not exceed $\epsilon/2$, and furthermore that Player J cannot obtain more than $\epsilon/2$ by choosing a move outside of support (f^J) . Because support $(f^J) \subseteq \text{support}(f_l^J)$ for any $l > i_1$, both claims are easy to confirm. \square

2.3 Measurable Equilibria

The following proof was explained to me by J.-F. Mertens.

Proposition 2: Any Harsanyi equilibrium of an ergodic game will generate a measurable Bayesian equilibrium of that game.

Proof: For all $j \in N$ assume that $f^j : \Omega \rightarrow \Delta(A^j)$ are Harsanyi equilibria, meaning also that they are Borel measurable functions (measurable

with respect to the Borel fields \mathcal{F}^j , respectively). Let μ be the common prior probability distribution on Ω . For all $j \in N$, all moves $a \in A^j$, and all positive integers m define $W_0^j(a, m)$ to be the subset of Ω such that Player j can obtain a payoff of at least $1/m$ more by choosing the move a instead of the distribution determined by f^j , as evaluated by the subjective probability distribution t^j . Since f^k for $k \neq j$ is Borel measurable and the t^j are continuous in the weak topology, the sets $W_0^j(a, m)$ are also Borel. If $\mu(W_0^j(a, m))$ were positive for some j , a , and m , then the $(f^j \mid j \in N)$ would not have been a Harsanyi equilibrium. We define W_0 to be $\cup_{j,a,m} W_0^j(a, m)$. Because this union is countable, the set W_0 is also Borel of measure zero. Now for all $l \geq 1$ define inductively the sets $W_l^j := \{x \in \Omega \mid t^j(x)(W_{l-1}) > 0\}$ and $W_l := \cup_{j \in N} W_l^j$. We claim for all l and j that W_l^j is in \mathcal{F}^j , W_l is Borel, and $\mu(W_l) = 0$. We proceed by induction, assuming the claim for $l - 1$. That W_{l-1} is Borel implies that W_l^j is in \mathcal{F}^j , and this implies that $W_l = \cup_{j \in N} W_l^j$ is Borel. Due to the formula $\mu(W_{l-1} \cap W_l^j) = \int_{W_l^j} t^j(y)(W_{l-1}) d\mu(y)$ and the fact that $t^j(y)(W_{l-1}) > 0$ for all $y \in W_l^j$, $\mu(W_l^j) > 0$ would imply that $\mu(W_{l-1}) \geq \mu(W_{l-1} \cap W_l^j) > 0$, a contradiction. Therefore we can assume also that $\mu(W_l) = 0$.

Define $W := \cup_{l=0}^{\infty} W_l$. We have two important properties, that W is Borel with $\mu(W) = 0$, and also from the structure of ergodic games that W is the union of cells. We can alter our Harsanyi equilibrium. We keep the original functions on $\Omega \setminus W$, and for all the cells in the set W we introduce any Bayesian equilibria obtained from Proposition 1. The result is a measurable Bayesian equilibrium that we seek.

3 Many Bayesian Equilibria, none Measurable

3.1 The Example

Let a and b be two distinct states of nature. Define Ω to be $\{a, b\}^{\mathbf{Z}}$, where \mathbf{Z} is the set of integers, including both the positive and the negative. The 0-coordinate of a point in Ω determines the state of nature, so that if $y \in \Omega$ and $y^0 = a$ then a is the state of nature at the point y . We define μ to be the canonical Bernoulli distribution on Ω , giving equal probability independently

to a and b in all coordinate positions.

There will be three players, Players One, Two, and Three. Let $\sigma : \Omega \rightarrow \Omega$ be the measure preserving involution defined by $(\sigma(y))^i := y^{-i}$, where x^i is the i th coordinate of $x \in \Omega$. σ is the reflection of the doubly infinite sequence about the position zero. Let $\tau : \Omega \rightarrow \Omega$ be the measure preserving involution defined by $(\tau(y))^i := y^{1-i}$. It follows that $T := \tau \circ \sigma$ is the usual Bernoulli shift operator $(T(y))^i = y^{i-1}$. The beliefs of the players are determined as follows: at any point $y \in \Omega$ Player One considers only y and $\sigma(y)$ to be possible, and with equal probability; (if $\sigma(x) = x$ then Player One believes in x with full probability). At any x , both Player Two and Player Three believe that only x and $\tau(x)$ are possible, and with equal probability. Because the involutions σ and τ are continuous functions, our game is an ergodic game. The cells will be the orbits of the involutions σ and τ . Player Two and Three have the same beliefs. Player One always knows the state of nature, a or b , but not always what the other players might know. Such a game is called a game of incomplete information *on one and a half sides* (Sorin and Zamir 1985), though this term was invented with reference to infinitely repeated zero-sum games.

Each player has two moves, L and R , standing for left and right. The game is a variation of the well known game of matching pennies, with Player One playing against both Player Two and Player Three. Player Two and Three want to match the pennies, Player One wants to have a mismatch. The differences to the conventional game of matching pennies has two aspects. First, Player Two has a special relationship to the move L , her favorite move, and likewise Player Three has a special relationship to her favorite move R . Second, if both Players Two and Three choose their favorite moves, then the payoff to Player One is dependent on the state of nature.

By playing L , Player One at $x \in \Omega$ receives
 -1 if Player Two and Player Three both choose L ,
 1 if Player Two and Three both choose R ,
 0 if Player Two chooses R and Player Three chooses L , and
 $\delta_a(x^0)$ if Player Two chooses L and Player Three chooses R . δ stands for the Kroniker delta, which means that $\delta_a = 1$ if the state of nature is a and $\delta_a = 0$ if the state of nature is b .

By playing R , Player One at $x \in \Omega$ receives
 1 if Player Two and Player Three both choose L ,

-1 if Player Two and Three both choose R ,
 0 if Player Two chooses R and Player Three chooses L , and
 $\delta_b(x^0)$ if Player Two chooses L and Player Three chooses R .

By playing L , Player Two receives (independently of the state of nature)
 2 if Player One and Player Three choose L ,
 1 if Player One chooses L and Player Three chooses R ,
 -1 if Player One chooses R .

By playing R , Player Two receives
 -1 if Player One chooses L , and
 1 if Player One chooses R .

By playing L , Player Three receives
 1 if Player One chooses L , and
 -1 if Player One chooses R .

By playing R , Player Three receives
 -1 if Player One chooses L ,
 1 if Player One chooses R and Player Two chooses L , and
 2 if Player One and Player Two choose R .

For all Players $i = 1, 2, 3$ let $f^i : \Omega \rightarrow [0, 1]$ be the behavior strategy for Player i , with $f^i(x)$ representing the probability at $x \in \Omega$ that Player i chooses the move L . The only a-priori requirement placed on the behavior strategies $(f^i \mid i = 1, 2, 3)$ is that for all $x \in \Omega$ $f^1(x) = f^1(\sigma(x))$ and $f^j(x) = f^j(\tau(x))$ for either $j = 2, 3$.

Given any Player $j \in \{1, 2, 3\}$ and two functions $(f^i \mid i \neq j)$ for the other two players, one can calculate the expected payoff to Player j at every $x \in \Omega$ for choosing L and for choosing R ; for Player One one must average the expected payoffs of the moves with respect to $(f^i(x) \mid i \neq 1)$ and $(f^i(\sigma(x)) \mid i \neq 1)$, and for Players Two and Three one must do the same with respect to $(f^i(x) \mid i \neq j)$ and $(f^i(\tau(x)) \mid i \neq j)$ for $j = 2, 3$. The three functions $(f^i \mid i = 1, 2, 3)$ are a Bayesian equilibrium for the game if and only if at every appropriate pair of points for every j if Player j chooses the move $D \in \{L, R\}$ with positive probability according to f^j then the other move $E \in \{L, R\}$, $E \neq D$ must deliver to Player j no higher an expected payoff than does the move D .

Define a pair of strategies $(f^2, f^3) \in [0, 1]^2$ for Players Two and Three to be a *ballanced* pair if and only if

$f^2 = 1$ and $f^3 < 1$ or
 $f^3 = 0$ and $f^2 > 0$.

The ballanced pairs are those for which at least one of the two players chooses her favorite move with certainty and the other player is not choosing the same move with certainty.

The two pairs of strategies $(f^2, f^3) \in [0, 1]^2$ such that $f^2 = f^3 = 0$ or $f^2 = f^3 = 1$ will be called the *coordinated* pairs. A strategy $f^1 \in [0, 1]$ of Player One that takes on the value of 0 or 1 will be called a *pure* strategy. A strategy of Player One in $[0, 1]$ that is not pure will be called *mixed*.

Lemma 1: In any Bayesian equilibrium behavior Players Two and Three are using only pairs that are ballanced or coordinated.

Proof: Suppose for the sake of contradiction that $f^2(x) < 1$ and $f^3(x) > 0$. If $f^1(x) + f^1(\tau(x)) \geq 1$ then Player Two would prefer to choose L ($f^2 = 1$) at the pair x and $\tau(x)$, and if $f^1(x) + f^1(\tau(x)) \leq 1$ then Player Three would prefer to choose R ($f^3 = 0$) at the pair x and $\tau(x)$. \square

Although the strategies available a-priori to Players Two and Three are two dimensional, in Bayesian equilibrium only a one-dimensional subset will be used. This will allow us to perceive the game much like a two-by-two game played between two players. Notice that if both Players Two and Three choose their favorite moves with certainty in a Bayesian equilibrium then it is impossible for either player to prefer her favorite move over the other move (since otherwise the other player would also prefer her non-favorite move). For this reason we perceive the pair $(f^2 = 1, f^3 = 0)$ as ballanced, although strictly speaking both players are choosing pure strategies.

3.2 Strategy of the Proof

We suppose for the sake of contradiction that the three functions $f^i : \Omega \rightarrow [0, 1]$ for $i = 1, 2, 3$ are measurable behavior strategies in Bayesian equilibrium. We will show that this leads to a contradiction. From Proposition 1 there would exist non-measurable Bayesian equilibria.

Our strategy is the following. We will divide Ω into two parts. The first part will be the set of $x \in \Omega$ where Player One is using a mixed strategy and Players Two and Three are using a ballanced strategy; the second part will be its compliment. The measurability assumption will imply that both parts

of Ω are of measure zero.

From ergodic theory, we will need a few well known results.

First, we need the Birkhoff Ergodic Theorem. If a transformation T on a probability space Ω is measure preserving and f is an integrable function on Ω then the Birkhoff Ergodic Theorem states that $\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} f(T^i(x))$ converges almost everywhere to an integrable function f^* such that the integral of f^* over Ω is equal to that of f (Theorem 1.14, Walters, 1982).

Second, we need the property called *ergodic*. A measure preserving transformation T of a probability space with a probability measure μ is ergodic if the only measurable sets B with the property $T^{-1}(B) = B$ satisfy $\mu(B) = 0$ or $\mu(B) = 1$. A measure preserving transformation is ergodic if and only if the only measurable sets B with $\mu(T^{-1}(B) \Delta B) = 0$ are those with $\mu(B) = 0$ or $\mu(B) = 1$ (Theorem 1.5, Walters, 1982), where here Δ stands for the symmetric difference. It is well known that the Bernoulli shift operator of our example is ergodic.

The cells partition the space Ω . Because the σ and τ are measure preserving, the smallest union of cells containing any measure zero subset of Ω is also of measure zero. Notice that for almost all points $x \in \Omega$ that the points $T^k(x)$ and $T^k \circ \tau(x)$ for all $k \in \mathbf{Z}$ are distinct and comprise the cell containing x . Define such points and their cells to be *doubly infinite*.

For any finite $k \geq 1$ any measurably defined behavior that occurs at most k times in any doubly infinite cell must occur only in a set of measure zero, since otherwise the distinctness of the points $T^k(x)$ and $T^k \circ \tau(x)$ for almost all $x \in \Omega$ would imply that the space Ω has infinite measure.

Define a point $x \in \Omega$ to be *normal* if and only if Player One uses a mixed strategy at x and Players Two and Three use a ballanced strategy at x . Define a cell to be normal if and only if it contains at least one normal point. The other cells will be called *abnormal*.

3.3 Normal cells

We define a homeomorphism between the ballanced pairs and the real numbers. Define $k : (0, 1] \rightarrow \mathbf{N}_0 = \{0, 1, 2, \dots\}$ and $s : (0, 1] \rightarrow [0, 1]$ by $r = 3^{-k(r)}(1 - 2s(r)/3)$ for $k(r)$ being the last number such that $3^{-k(r)} \geq r$. Define W to be the set of ballanced pairs, $W := \{(f^2, f^3) \mid f^2 = 1 \text{ or } f^3 = 0\} \setminus \{(0, 0), (1, 1)\}$. Define the homeomorphism $\phi : W \rightarrow \mathbf{R}$ by

$\phi(1, 0) := 0$,
 $\phi(1, t) := k(1 - t) + s(1 - t)$, and
 $\phi(t, 0) := -s(t) - k(t)$.

Lemma 2: If x is doubly infinite, $f^1(x)$ is mixed and the strategy pair $(f^2(x), f^3(x))$ is ballanced, then there is one and only one value for $f^1(\tau x)$ that preserves the Bayesian equilibrium property for Players Two and Three and there is one and only one ballanced pair $(f^2(\sigma(x)), f^3(\sigma(x)))$ that preserves the Bayesian equilibrium property for Player One, namely $f^1(\tau(x)) = 1 - f^1(x)$ and

$$(f^2(\sigma(x)), f^3(\sigma(x))) = \phi^{-1}(-1^{\delta_b(x^0)} - \phi(f^2(x), f^3(x))).$$

Proof: By symmetry we assume that $f^2(x) = 1$. The condition that $f^1(\tau(x)) = 1 - f^1(x)$ follows by the indifference of Player Three (and for both players in the case that $f^3(x) = 0$, as discussed above).

Again, assuming $f^2(x) = 1$, we must divide the argument into three cases. Case A is that of $0 \leq f^3(x) \leq 2/3$ and $x^0 = a$, Case B is that of $x^0 = b$, and Case C is that of $2/3 \leq f^3(x) < 1$ and $x^0 = a$. By Lemma 1 we need only consider the ballanced pairs and the coordinated pairs. Notice that increasing values for ϕ in both cases of $x^0 = a$ and $x^0 = b$ imply an increasing preference by Player One for the move R , so that if we find one ballanced pair that delivers indifference to Player One then we have found the only such ballanced pair, and furthermore neither coordinated pair could deliver such an indifference.

Case A: Consider the strategy pair $g^2 = 1$ and $g^3 = 2/3 - f^3(x)$, which satisfy $\phi(f^2(x), f^3(x)) + \phi(g^2, g^3) = 1$. If Player One chooses L , he can expect a payoff of $\frac{1}{2}(-f^3(x) + 1 - f^3(x) - g^3 + 1 - g^3) = 1/3$. By playing R he can expect a payoff of $\frac{1}{2}(f^3(x) + g^3) = 1/3$.

Case B: Consider the strategy pair $g^3 = 0$ and $g^2 = (1 - f^3(x))/3$, which satisfy $\phi(f^2(x), f^3(x)) + \phi(g^2, g^3) = -1$. If Player One chooses L , he can expect a payoff of $\frac{1}{2}(-f^3(x) + 1 - g^2) = (1 - f^3(x))/3$. If Player One chooses R , he can expect a payoff of $\frac{1}{2}(f^3(x) + 1 - f^3(x) - (1 - g^2) + g^2) = (1 - f^3(x))/3$.

Case C: Consider the strategy pair $g^3 = 0$ and $g^2 = 3(1 - f^3(x))$, which satisfy $\phi(f^2(x), f^3(x)) + \phi(g^2, g^3) = 1$. If Player One chooses L , he can expect a payoff of $\frac{1}{2}(-f^3(x) + 1 - f^3(x) + g^2 + 1 - g^2) = 1 - f^3(x)$. By playing R he can expect a payoff of $\frac{1}{2}(f^3(x) - (1 - g^2)) = 1 - f^3(x)$. \square

Lemma 3: The measure of the union of all normal cells must be zero or one.

Proof: By Lemma 2 a doubly infinite cell is normal (contains a normal point) if and if all the points in the cell are normal. This means that the set of normal points is a T -invariant set, and therefore by the ergodicity of T they must be of measure zero or one. \square

Lemma 4: Let $f : \Omega \rightarrow \mathbf{R}$ be a function with $f(Tx) + f(x) = -1^{\delta_b(x^0)}$ for almost all $x \in \Omega$. **Conclusion:** The function f cannot be measurable.

Proof: Suppose for the sake of contradiction that the function f is measurable. This would imply that there exists an $M > 0$ such that the probability that f is in $[-M + 2, M - 2]$ is at least $9/10$.

Now consider the function $g_M : \Omega \rightarrow [0, 1]$ defined by

$$g_M(x) := \liminf_{n \rightarrow \infty} \frac{|\{k \mid |k| \leq n, |f(T^k x)| \geq M\}|}{2n + 1}.$$

Now consider the transformation T^2 . The action of T^2 is that of a random walk, with

$$f(T^2 x) - f(x) = 2 \text{ if } x^0 = b \text{ and } x^{-1} = a,$$

$$f(T^2 x) - f(x) = -2 \text{ if } x^0 = a \text{ and } x^{-1} = b, \text{ and}$$

$$f(T^2 x) = f(x) \text{ if both } x^0 \text{ and } x^{-1} \text{ have the same value.}$$

Therefore for almost all x there is a $k \in \mathbf{Z}$ with $f(T^k x) > M + 1$, which means that $g_M(x) \geq 1/2$ for almost all $x \in \Omega$. But by the Birkhoff Ergodic Theorem, the expected value of g_M must equal the probability in Ω that the function f exceeds the value of M , a contradiction. \square

Proposition 3: The union of all normal cells is of measure zero.

Proof: This result follows directly from Lemmata 2, 3 and 4. \square

3.4 Abnormal cells

Define $x \in \Omega$ to be a *pure* point if and only if $f^i(x) \in \{0, 1\}$ for all $i = 1, 2, 3$ with $f^2(x) = f^3(x)$ (meaning that Player One uses a pure strategy at x and Players Two and Three use a coordinated pair). Define the *wind* direction from a doubly infinite pure point x to be the direction on the cell of the y adjacent to x (equal to either $\sigma(x)$ or $\tau(x)$) such that $y = \sigma(x)$ if all three of

the $f^i(x)$ have the same value in $\{0, 1\}$, and otherwise $y = \tau(x)$ if Player One chooses a different move from that chosen by Players Two and Three. Define the *lee* direction from a doubly infinite pure point to be the direction opposite to the wind direction. Define a pure point x to be a *book end* if additionally either Player One at x has a clear preference for his chosen move or the player among Players Two or Three who is not choosing her favorite move has a clear preference for her chosen and non-favorite move.

Lemma 5:

(1) The behavior of all players on the wind side of a doubly infinite pure point is determined; all of the points on the wind side of a doubly infinite pure point are also pure, alternating in values, and none are book ends.

(2) There can be at most two book ends in any doubly infinite abnormal cell.

(3) If there are no book ends in a doubly infinite abnormal cell, then either Player One or Players Two and Three are performing alternating behavior, meaning that choosing any point x in the cell we have that either the value $f^1(T^n(x)) \in \{0, 1\}$ is determined by the parity of n or the value $f^2(T^n(x)) = f^3(T^n(x)) \in \{0, 1\}$ is determined by the parity of n .

Proof:

(1) By symmetry, we can assume that Players Two and Three choose the move L at x . We consider two cases, (A) that Player One chooses L and (B) that Player One chooses R . In Case (A) the point $\sigma(x)$ (and the transformation T) defines the wind direction. If Players Two and Three did not both choose R at $\sigma(x)$, then it would have been in the interest of Player One to choose R at x . In Case (B) the point $\tau(x)$ (and T^{-1}) defines the wind direction. If Player One did not choose L at the point $\tau(x)$ then it would have been in the interest of Player Three to choose R at x . The rest follows by induction.

(2) For the sake of contradiction, let us assume that there is at least three book ends x in a doubly infinite cell. One of these three book ends must be between the other two. By Part 1, in one of the two directions from this middle book end there are no book ends, a contradiction.

(3) We consider three cases, (A) that there are no pure points in the cell and that Player One chooses a pure strategy at x , (B) that there are no pure

points in the cell and Players Two and Three choose a coordinated pair at x , and (C) x is a pure point in the cell.

Case (A): If Player One is choosing L at x , then the lack of a pure point implies that Players Two and Three are choosing ballanced strategies at both x and $\sigma(x)$. To maintain the indifference of these ballanced choices by at least one of either Players Two or Three, Player One must choose R at Tx and $T^{-1}x$. The rest follows by induction.

Case (B): If Players Two and Three are both choosing L at x , then the lack of a pure point implies that Player One is choosing mixed strategies at both x and $\tau(x)$. To maintain the indifference of Player One, Players Two and Three are both choosing R at both Tx and $T^{-1}x$. The rest follows by induction.

Case (C): Assume that x is a pure point, that Players Two and Three are choosing L at x , and Player One is choosing either L or R at x . If Player One is choosing L at x , then the lack of a book end implies that Player One is choosing R at $T^{-1}(x)$, in the lee direction of x . If Player One is choosing R at x , then the lack of a book end implies that Players Two and Three are both choosing R at $T(x)$, in the lee direction of x . By induction and Part 1 we conclude that the whole cell consists of alternating behavior by all three players. \square

Proposition 4: The union of all the abnormal cells is a measure zero set.

Proof: Due to Part Two of Lemma 5, the measure of the set of all doubly infinite book ends is zero, and therefore also of all the cells containing a book end.

Let us define A to be the set of all points where Player One chooses the pure strategy L , and B the set where Player One chooses the pure strategy R . By assumption both A and B are measurable sets and by Lemma 5 $T(A) = B$ and $T(B) = A$ (modulo sets of measure zero). By the ergodicity of T we must have that $\mu(A) = 0$ or $\mu(A) = 1$. Since A and B are sets of the same measure, we must conclude that both are of measure zero. We proceed the same way with the sets where Players Two and Three both choose either L or R . By Lemma 5 we have exhausted the abnormal cells. \square

Theorem 1: (f^1, f^2, f^3) , an arbitrary Bayesian equilibrium, could not have consisted of measurable functions, (and therefore there could not have

been a Harsanyi equilibrium by Proposition 2).

Proof: The result follows directly from Propositions 3 and 4. \square

4 The Challenge of Approximate Equilibria

4.1 The Example

The following is an example of a zero-sum ergodic game for which there would be no value in measurable strategies if the conjecture (presented in the introduction) were true, though (from Proposition 1) there would exist Bayesian equilibria.

There are eleven states of nature, with $B = \{b_0, \dots, b_{10}\}$. The space is $\Omega := B^{\mathbb{Z}^2}$, with the canonical Bernoulli distribution μ . The $0,0$ -coordinate determines the state of nature, so that if $y \in \Omega$ and $y^{0,0} = b_i$ then the payoffs are determined by the state b_i .

There are two players, Player One and Player Two, and at each state each player has eleven moves. It is a zero-sum game, so that the payoffs presented will be those of the first player. The following payoffs are for the state i , modulo 11, and the moves, also modulo 11.

If Player One chooses the move i , then the payoff is
 -100 if Player Two chooses i ,
 11 if Player Two chooses $i - 1$ (modulo 11), and
 10 if Player Two chooses any other move $k \notin \{i, i - 1\}$.

If Player One chooses the move $k \neq i$, then the payoff is
 -100 if Player Two chooses k , and
 10 if Player Two chooses $l \neq k$, including $l = i$.

Next, consider the measure preserving involutions $\sigma_1, \sigma_2 : \Omega \rightarrow \Omega$ defined by $(\sigma_1(y))^{i,j} := y^{-i,j}$ and $(\sigma_2(y))^{i,j} := y^{i,-j}$. Notice that the σ_1 and σ_2 commute, so that for every $x \in \Omega$ we get at most four distinct points in the orbit of σ_1 and σ_2 applied to x . At every $x \in \Omega$ define $t^1(x)$ so that Player One believes that $x, \sigma_1(x), \sigma_2(x), \sigma_2 \circ \sigma_1(x)$ are given each $1/4$ probability, (with $i/4$ probability given to a point that appears with multiplicity i). As before with the previous example, Player One always knows the state of

nature, but not necessary what Player Two believes.

The measure preserving involutions $\tau_1, \tau_2 : \Omega \rightarrow \Omega$ are defined by $(\tau_1(y))^{i,j} := y^{1-i,j}$ and $(\tau_2(y))^{i,j} := y^{i,1-j}$. Likewise, the τ_1 and τ_2 commute, with at least four possibilities. At every $x \in \Omega$ define $t^2(x)$ so that Player Two believes that $x, \tau_1(x), \tau_2(x), \tau_2 \circ \tau_1(x)$ are given 1/4 probability each, (with $i/4$ probability given to a point that appears with multiplicity i). It follows that $\tau_1 \circ \sigma_1$ is the shift $T_1 : \Omega \rightarrow \Omega$ defined by $T(y)^{i,j} = y^{i-1,j}$ and $\tau_2 \circ \sigma_2$ is the shift $T_2 : \Omega \rightarrow \Omega$ defined by $T(y)^{i,j} = y^{i,j-1}$. Define $\mathcal{F}^1 := \{B \mid B \text{ is Borel and } x \in B \Leftrightarrow \{x, \sigma_1(x), \sigma_2(x), \sigma_1 \circ \sigma_2(x)\} \subseteq B\}$, and define \mathcal{F}^2 correspondingly with τ_1 and τ_2 . As before, for every Player j t^j is a regular conditional probability distribution of μ with respect to the Borel field \mathcal{F}^j .

4.2 Application of the Conjecture

Lemma 6: In Bayesian equilibrium, at every quartet of points considered possible by Player One, she receives an expected payoff of at least 0.

Proof: It suffices to consider the strategy that puts an equal weight of 1/11 to all moves. Regardless of the state of nature and the strategies of Player Two, this strategy of Player One delivers to her a payoff of at least 0. \square

Lemma 7: If Q is a member of \mathcal{F}^1 of positive measure and Player One uses a single move i with a probability of at least $1 - \delta$ at all points of Q then there will be a $Q' \in \mathcal{F}^2$ containing Q where on the average Player Two holds Player One down to a payoff of no more than $-67/4 + 111\delta/4$ and there is a $Q^* \in \mathcal{F}^1$ containing Q' where Player One is obtaining on the average no more than $-4/3 + 37\delta/3$.

Proof: For almost all points $x \in \Omega$, the structure of the cell containing x will be that of the two dimensional lattice \mathbf{Z}^2 . To see this, we define the actions of $\sigma_1, \sigma_2, \tau_1, \tau_2$ on the set $D := \mathbf{Z}^2$ in the following way.

If i is even, then $\sigma_1(i, j) = (i + 1, j)$ and $\tau_1(i, j) = (i - 1, j)$,
if i is odd, then $\sigma_1(i, j) = (i - 1, j)$ and $\tau_1(i, j) = (i + 1, j)$,
if j is even, then $\sigma_2(i, j) = (i, j + 1)$ and $\tau_2(i, j) = (i, j - 1)$,
if j is odd, then $\sigma_2(i, j) = (i, j - 1)$ and $\tau_2(i, j) = (i, j + 1)$.

We have that $T_1(i, j) := \tau_1 \circ \sigma_1(i, j) = (i + 2, j)$ if i is even and equal to $(i - 2, j)$ if i is odd, and that $T_2(i, j) := \tau_2 \circ \sigma_2(i, j) = (i, j + 2)$ if j is even

and $(i, j - 2)$ if j is odd.

To confirm that this lattice structure represents correctly the generic cell and the actions of the involutions, we need to know two things, that points identified by D came from identical points in Ω , and that generically distinct points in D come from distinct points of Ω . For the former, if we start at any point in D and choose two pathways defined by our involutions that lead to the same point in D then we need to know that these involutions induce the same action on Ω . We must check that moving from (i, j) to $(i + 1, j + 1)$ (and also the corresponding paths to $(i + 1, j - 1)$, $(i - 1, j + 1)$, or $(i - 1, j - 1)$) through either $(i + 1, j)$ or $(i, j + 1)$ correspond to the same action on Ω . This is confirmed by some pairs of involutions commuting, namely σ_1 with σ_2 , τ_1 with τ_2 , τ_1 with σ_2 , and σ_1 with τ_2 . For the latter, we notice for almost all $x \in \Omega$ that the points $T_1^m \circ T_2^n(x)$, $T_1^m \circ T_2^n \circ \tau_1(x)$, $T_1^m \circ T_2^n \circ \tau_2(x)$ and $T_1^m \circ T_2^n \circ \tau_1 \circ \tau_2(x)$ for all $m, n \in \mathbf{Z}$ are distinct and comprise the orbit of x . (The probability that any given pair are equal is a set of measure zero, and the rest follows by sigma additivity.)

Now let us assume that at some point $x \in \Omega$ Player One puts a weight of at least $1 - \delta$ on the move $k \in \{0, 1, \dots, 10\}$. Without loss of generality, let us assume that x is represented by the point $(0, 0)$ in D . This means also that a weight of at least $1 - \delta$ is put on the move k at the points $\sigma_1(x)$, $\sigma_2(x)$, $\sigma_1 \circ \sigma_2(x)$, which are now mapped to $(0, 1)$, $(1, 0)$ and $(1, 1)$ in D . Since at all four of these points Player Two could respond with the move k , the average payoff for Player One at the points mapped to $\{(i, j) \mid -1 \leq i, j \leq 2\}$ cannot exceed $\frac{1}{16}(12 \cdot 11 - 4 \cdot 100)(1 - \delta) + 11\delta = -67/4 + 111\delta/4$. Now consider the 36 different positions in D defined by $\{(i, j) \mid -2 \leq i, j \leq 3\}$. The average payoff for Player One on this set cannot exceed $\frac{1}{36}(32 \cdot 11 - 4 \cdot 100)(1 - \delta) + 11\delta = -4/3 + 37\delta/3$. \square

For every $k \in \{0, \dots, 10\}$ and distinct pair $m, n \in \{0, \dots, 10\}$ define $E_k^{(m,n)}$ to be the subset of mixed strategies of Player Two in $\Delta(\{0, \dots, 10\})$ against which the moves m and n both maximize Player One's payoff at the state b_k (meaning also that Player One is indifferent between the moves m and n). Define $E_k := \cup_{m \neq n} E_k^{(m,n)}$.

Lemma 8: The intersection $\cap_{k \in \{0, \dots, 10\}} E_k$ is empty, and for every strategy $g \in \Delta(\{0, \dots, 10\})$ that is a positive distance of $\delta > 0$ away from the set E_k there is a move of Player One that is preferred to all other moves at the state

k by a value of at least $\delta/5$.

Proof: Let $q \geq 0$ be the least weight given by g to any move in $\{0, \dots, 10\}$. If $q = 1/11$ and g gives equal weight to all moves, then at any state b_k the move k will be preferred to all others by Player One. On the other hand, if $q < 1/11$, then choose any $k \in \{0, \dots, 10\}$ such that g gives strictly more than q to the move $k - 1$ (modulo 11) and gives exactly q to the move k (modulo 11). At the state b_k the move k will be preferred by Player One to all other moves.

The last claim follows directly from the fact that 1 is the smallest positive difference between any of the payoffs and the diameter of the simplex is strictly smaller than 5. \square

Theorem 2: If the conjecture is valid, then there is a zero-sum ergodic game with no value in measurable functions.

Proof: Define the variable d_g in relation to the subsets E_i as stated by the conjecture. By the conjecture (and Lemma 8) the expected value of d_g is at least $w > 0$ for all measurable functions g . Without loss of generality we will assume that $w \leq 1$. We will show that there cannot be a value $v \in \mathbf{R}$ and strategy functions (f, g) for Players One and Two, measurable in \mathcal{F}^1 and \mathcal{F}^2 respectively, such that the combination (f, g) delivers an expected payoff of v to Player One, f guarantees a payoff of at least $v - w^2/10,000$ against all measurable strategies of Player Two and g holds Player One down to a payoff of no more than $v + w^2/10,000$ against all measurable strategies of Player One.

Let Q be the measurable subset in \mathcal{F}^1 where according to f Player One uses a single move with probability of at least $1 - 1/90$. By Lemma 7 there is an even larger set $Q^* \in \mathcal{F}^1$ where Player One is getting an average payoff of no more than $-4/3 + 37/270 < -1$. Since by Lemma 6 Player One can guarantee a payoff of at least 0 in the set Q^* , by the approximate optimality of g we must assume that the measure of Q^* (and also Q) is no greater than $w^2/10,000$.

Now let us look at a point x such that $x^{0,0} = b_k$ and $\frac{1}{4}(g(x) + g(T_1(x)) + g(T_2(x)) + g(T_1 \circ T_2(x)))$ is a distance of positive $r > 0$ from the set E_k . We notice that $T_1(x) = \tau_1 \circ \sigma_1(x)$, $T_2(x) = \tau_2 \circ \sigma_2(x)$, and $T_1 \circ T_2(x) = \tau_2 \circ \tau_1 \circ \sigma_1 \circ \sigma_2(x)$. Since the behavior of Player Two remains constant within the finite orbit of the τ_1 and τ_2 , we see that $\frac{1}{4}(g(x) + g(\sigma_1(x)) + g(\sigma_2(x)) + g(\sigma_1 \circ \sigma_2(x)))$

is a distance of positive r from the set E_k . By Lemma 8, Player One has a single move at x that delivers to her $r/5$ more than any other move. By the conjecture the distance $r = w/2$ is exceeded in a member of \mathcal{F}^1 of measure at least $w/10$ (since the diameter of the simplex is less than 5). Since the set Q where Player One chooses one move with frequency greater than $1 - 1/90$ is so small, we conclude that there is a set $P \in \mathcal{F}^1$ of measure at least $w/11$ where the distance to the appropriate E_k (the k with $b_k = x^{0,0}$) is at least $w/2$ and Player One does not use any move with more than $1 - 1/90$ probability. But now we can conclude by Lemma 8 that Player One can improve her expected payoff by at least $\frac{1}{90} \frac{w}{10}$ in this subset of measure at least $\frac{w}{11}$ by choosing some single move. Due to the measurability of g , a decision concerning which move is sufficient to obtain a payoff gain of $\frac{1}{90} \frac{w}{10}$ can be done in a measurable way. \square

4.3 Single-Shift Spaces

Now we look at how the conjecture fails for single-shift spaces.

Let us consider the following example. $C = [0, 2]$, $A_1 = \{1/3, 4/3\}$, $A_2 = \{2/3, 5/3\}$, and $\Omega = \{1, 2\}^{\mathbf{Z}}$. We assume for all $x \in \Omega$ that $\frac{1}{2}(f(x) + f(T(x))) \in A_{x^0}$. Modulo one, the determination of $f(x)$ for any $x \in \Omega$ will determine $f(T^k x)$ for all $k \in \mathbf{Z}$. From now on we will view everything in terms of modulo one. By the axiom of choice, we have a function f defined on all of Ω . First we show that f cannot be measurable, and second that measurable functions f can however approximate the requirement $\frac{1}{2}(f(x) + f(T(x))) \in A_{x^0}$.

If f were measurable, by Luzin's Theorem, we should be able to approximate our function f in probability with continuous functions, meaning that for every $\epsilon > 0$ there exists a continuous function $g : \{1, 2\}^{\mathbf{Z}} \rightarrow [0, 1]$ (modulo 1) such that f and g differ only on a set of measure ϵ . Let ϵ be $1/100$. Again by Luzin's Theorem we know that there is a cylinder set E and a value $r \in [0, 1]$ (modulo 1) such that within the set E the probability (conditioned on membership in E) that the value of f is further than ϵ from r is less than ϵ . Without loss of generality, we can assume that our cylinder set is defined by the choice of coordinates $y^{-k}, \dots, y^0, \dots, y^{k-1}$ for some positive integer k . We define the value q (modulo 1) by it being the functional value of $T^{-k}(x)$ (modulo 1) for any x with $f(x) = r$ and define the value s to be the corresponding value of the $T^k(x)$ (modulo 1). As long as $x \in E$ is mapped to

within ϵ of r , $T^{-k}(x)$ and $T^k(x)$ (modulo 1) will be mapped to within ϵ of q and s , respectively. The occurrences of the cylinder set E in a typical orbit of T will be separated by many different sequences of intermediate values for the 0-coordinate. By our assumption, these separating sequences must, with probability at least $1 - 2\epsilon$, connect values within ϵ of s to values within ϵ of q , an absurdity as long as ϵ is small compared to the distance between $1/3$ and $2/3$.

Looking at the same example, for any $0 < \delta < 1/3$ let us require instead that $|\frac{1}{2}(f(x) + f(T(x))) - 1/3| \leq \delta$ (modulo 1) if $x^0 = 1$ and $|\frac{1}{2}(f(x) + f(T(x))) - 2/3| \leq \delta$ (modulo 1) if $x^0 = 2$. Let N be larger than $\frac{3}{\delta}$. Call an orbit *typical* if there are infinitely many points x with $x^0 = 2$ and also infinitely many points x of the orbit such that $(T^n(x))^0 = 1$ for all $-N \leq n \leq N$. The latter kind of point we call a *center* point. We will define a map $f : \Omega \rightarrow [0, 1]$ (modulo 1) on the typical orbits in the following way. Every center point gets mapped to $1/3$ (modulo 1). For every center point x let $m(x) > N$ be the first number such that $T^{m(x)+N}(x)$ is a center point and $(T^k(x))^0 = 2$ for some $k < m(x)$. ($m(x)$ will be finite for all center points of typical orbits). For all points of the form $y = T^k(x)$ for $0 \leq k \leq m(x)$, determine $f(y)$ inductively according to the rule $\frac{1}{2}(f(y) + f(T(y))) = 1/3$ (modulo 1) if $y^0 = 1$ and $= 2/3$ (modulo 1) if $y^0 = 2$. For the row $y = T^{m(x)}(x), T^{m(x)+1}(x), \dots, T^{m(x)+N}(x)$ we define f so that $|\frac{1}{2}(f(y) + f(T(y))) - 1/3| < \delta$ (modulo 1) and $f(T^{m(x)+N}(x)) = 1/3$, (with $T^{m(x)+N}(x)$ the next center point after $T^{m(x)}(x)$). We claim that these motions back to $1/3$ can be done in a deterministic and measurable way with respect to the starting point $T^{m(x)}(x)$. Because the definition of f on the typical orbits can be broken down according to the distances to the next center points, the definition of f as restricted to the typical orbits would be measurable in character. But since almost all points belong to typical orbits, that would be sufficient for the measurability of f .

We believe that careful analysis of the example in Section 3 would reveal a similar mechanism to generate measurable ϵ -Bayesian equilibria for every positive ϵ , and we suspect that there exist measurable approximate Bayesian equilibria for all ergodic games played on Bernoulli single-shift spaces. On the other hand we see no such similar mechanism for getting measurable approximate Bayesian equilibria for games played on double-shift spaces or more exotic probability spaces with measure preserving transformations.

5 Conclusion: Zero-sum games

We are interested in zero-sum ergodic games because of the possibilities for violating conventional game theory. For the finite dimensional matrix games for which the original min-max theorem was proven (Von Neumann 1928), the value of the game is the expected payoff from all equilibrium strategies. Equilibrium strategies are also called optimal strategies, because if one exchanges the strategies in two pairs of equilibria, one obtains again an equilibrium.

Even if a zero-sum ergodic game does has measurable Bayesian equilibria, the non-measurable varieties can play havoc with the conventional understanding of value, expected payoff, equilibrium, and optimality. Let us look at a very simple example.

Let Ω be $\{a, b\}^{\mathbf{Z}}$ with the same belief structure as the example in Section 3, but with only two players, Player One and Player Two. We assume that there is only one payoff matrix, that corresponding to the game of *matching pennies*. Both players have only two moves, 0 and 1, and if the sum of their moves is even then Player One receives 1 and if this sum is odd than Player One receives -1 . Let $f^j : \Omega \rightarrow [0, 1]$ for $j = 1, 2$ represent the probability that Player j chooses the move 0.

We will consider two pairs of Bayesian equilibrium strategies for the players, (f_1^1, f_1^2) and (f_2^1, f_2^2) . For both pairs the payoff will not be a measurable function on Ω , however for the first pair for each player there will be an expected payoff of 0 at all his pairs of states. For the second pair, however, Player Two will have an expected payoff of 0 at all his pairs of states, but the payoff to Player One will have no measurable interpretation. Exchanging these pairs of strategies will not create a new Bayesian equilibrium.

To define the pair (f_1^1, f_1^2) choose any doubly infinite $x \in \Omega$ and define $f_1^1(x) = f_1^2(x) = 1$, but also $f_1^1(\tau(x)) = f_1^1(\sigma \circ \tau(x)) = 0$ and $f_1^2(\sigma(x)) = f_1^2(\tau \circ \sigma(x)) = 0$. Continue defining the strategies f_1^1 and f_1^2 on the cell containing x in an alternating way so that at all y in the cell we have $f_1^1(y) + f_1^1(\tau(y)) = 1$ and $f_1^2(y) + f_1^2(\sigma(y)) = 1$, with the values for f_1^1 and f_1^2 always either 0 or 1. This defines an equilibrium on the cell, and repeat this process for all the doubly infinite cells.

For x being any doubly infinite point with $f_1^1(x) = f_1^2(x) = 1$, for all points $y = T^k x$ with $k \geq 1$ define $f_2^2(y)$ to be the same as $f_1^2(y)$, but for all $y = T^{-k}(x)$ with $k \geq 0$ switch the values, meaning that $f_2^2(y) + f_1^1(y) = 1$.

For any y on this same cell, define f_2^1 by $f_2^1(y) + f_1^1(y) = 1$. Notice that f_2^1 and f_2^2 are also in Bayesian equilibrium, but with the difference that for every doubly infinite cell there is one special pair of points considered possible by Player One, namely our chosen x and $\sigma(x)$, where Player One expects a payoff of 1 instead of 0.

The pair (f_1^1, f_2^2) is not in Bayesian equilibrium, because Player One is very dissatisfied in his choice at the point x . One could object: Player One's dissatisfaction at only one point in the cell means that he is satisfied almost everywhere. This objection is flawed, for two reasons. First, since an adjacent point on a cell is evaluated to have a probability of one-half by some player, once the equilibrium property is destroyed at one point, it could be destroyed inductively for the entire containing cell. Second, since such behavior is repeated for almost all cells, the extreme dissatisfaction of Player One would not be limited to a set of measure zero. A set of points that intersect every cell in one and only one point is a non-measurable set, meaning that it has positive outer measure. It follows by the same argument that with the Bayesian equilibrium (f_2^1, f_2^2) there is no expected payoff for Player One.

One could try to define the value of a zero-sum ergodic game as some kind of limiting average payoff on the cells. Although the payoff function may not be measurable, perhaps this limiting average payoff is measurable and well behaved (e.g. the conclusion of the Birkhoff Ergodic Theorem without the conditions being satisfied). The problem with this approach is that it is uncertain what kind of interpretation should be given to such a limiting average. The measure preserving transformations in our situation do not represent the flow of time, the usual justification for ignoring finite subsets of a sequence. As argued above, we don't understand why within a cell the players should consider arbitrary finite subsets to be insignificant.

Acknowledgements: The belief space of the first example, without the payoff matrices, is a variation of a belief space given in Simon (2000), but its precise formulation in this paper is due to Robert Aumann. The proof of Proposition 2 was based on a conversation with J.-F. Mertens. Discussions with Robert Aumann, Zbigniew Nitecki, and Benjamin Weiss were very helpful.

6 References

- HARSANYI, J. C.**, (1967-8). “Games With Incomplete Information Played by Bayesian Players: Parts I, II, III”, *Management Science* 14 (3,5,7).
- MERTENS, J.-F. and ZAMIR, S.**, (1985). “Formulation of Bayesian Analysis for Games with Incomplete Information”, *International Journal of Game Theory*, Vol. 14, Issue 1, pp. 1-29.
- NASH, J.**, (1950). “Equilibrium Points in n -person Games”, *Proceedings of the National Academy of Sciences*, 36, pp. 48-49.
- VON NEUMANN, J.**, (1928). “Zur Theorie der Gesellschaftsspiele”, *Mathematische Annalen* 100, pp. 295-320.
- SIMON, R.** (2000). “The Common Prior Assumption in Belief Spaces: An Example”, Discussion Paper 228, (December 2000), Center for Rationality and Interactive Decision Theory, Hebrew University, Jerusalem.
- SORIN, S. and ZAMIR, S.** (1985). “A Two-Person Game with Lack of Information on One and One-Half Sides,” *Mathematics of Operations Research* 10, 17-23.
- WALTERS, P.**, (1982). *An Introduction to Ergodic Theory*. Graduate Texts in Mathematics 79, Springer Verlag.