

A GAME WITH NO BAYESIAN APPROXIMATE EQUILIBRIA

ZIV HELLMAN

ABSTRACT. Simon (2003) presented an example of a 3-player Bayesian game over a continuum of states with no Bayesian equilibria but left open the question of whether it is possible to approximate Bayesian equilibria in such games. We present an example of a Bayesian game with two players, two actions and a continuum of states that possesses no Bayesian approximate equilibria, thus resolving the question. As a side benefit we also have for the first time an example of a 2-player Bayesian game with no Bayesian equilibria and an example of a strategic-form game with no approximate Nash equilibria. To achieve this we show a close relationship between strategic considerations in overlapping generations games and certain Bayesian games and then make use of an example by Y. Levy of an overlapping generations game with no stationary equilibria.

1. INTRODUCTION

It is safe to say that it is impossible to imagine modern economic modelling and game theory without the theory of Bayesian games. This would not be the case without one of the seminal contributions of Harsányi (1967), namely the analysis of Bayesian games for studying games of incomplete information, which included showing that every finite Bayesian game (finite number of players, finite actions, finite states of the world) has a Bayes-Nash, or Bayesian, equilibrium. The fact that modellers could safely assume the existence of at least one equilibrium was undoubtedly an element

Date: This version: November 2012 .

School of Mathematical Sciences, Tel Aviv University, *email:* ziv.hellman@mail.huji.ac.il. This research was supported in part by the European Research Council under the European Commission's Seventh Framework Programme (FP7/20072013)/ERC grant agreement no. 249159, and by Israel Science Fund Grants 538/11 and 212/09. Special thanks are due to Yehuda Levy for suggesting that I study this topic and participating in many long conversations on it. Amanda Friedenberg recommended an important restructuring of the paper that has been implemented. I am also grateful to Sergiu Hart and Abraham Neyman for comments on improving the presentation and their general encouragement, to Robert Simon for extensive email correspondence and to Rann Smorodinsky for catching a mistake in an earlier version of this paper.

in the widespread acceptance of Bayesian games in modelling a wide range of economic situations.

Limiting attention to finite games alone, however, is not sufficient for capturing the full range of possible economic models. A large number of such models must make use of uncountably many states to represent quantities. Examples include models in which prices (as in models of auctions or bargaining, such as that of Chatterjee and Samuelson (1983) for example) are the main state variables, or in which the main variables are profits and outputs in market models (for example Radner (1980)), continuous time points, accumulated wealth or stocks, population percentages, share percentages and so forth.

The question of the existence or non-existence¹ of Bayesian equilibria in games with uncountably many states is therefore a significant matter. It remained, however, an open question for many years, until Simon (2003) gave a negative answer by presenting an example of a three-player Bayesian game over a continuum of states with no measurable Bayesian equilibrium.

That important result left in its wake (at least) two open questions: (1) are there examples² of games that have no Bayesian ε -equilibria?; (2) are there examples of two-player games that have no Bayesian equilibria? In particular, a negative answer to the first question would imply that modellers could always assume that Bayesian equilibria can be approximated as closely as desired in games with uncountably many states, thus significantly weakening Simon (2003)'s result,

We resolving these open questions here by showing that the answers to both are 'yes' via the construction of a two-player Bayesian³ game with no Bayesian ε -equilibria. As a side-benefit, the example also shows that there there exist strategic-form games with a continuum of players and no

¹ By the existence of an equilibrium we mean the existence of a measurable equilibrium. There are several reasons for restricting attention to measurable strategies (and hence measurable equilibria); to consider just two reasons, if a strategy is not measurable it cannot be constructed explicitly, and the payoffs of non-measurable strategies haven't got well-defined expected values. Measurability has in fact been included as a basic requirement in the definition of an equilibrium over uncountable spaces since the earliest literature on the subject (see Schmeidler (1973) for one such example). We therefore throughout this paper use the term 'existence of an equilibrium' as synonymous with 'existence of a measurable equilibrium' without further qualification.

² In private communication, Robert Simon has indicated that the example he presented in Simon (2003) does admit ε -equilibria for every $\varepsilon > 0$.

³ The game constructed here is not only a Bayesian game, it is an *ergodic game* as defined in Simon (2003).

Nash ε -equilibria⁴ and that there exist two-player Bayesian games with no Harsányi equilibria (meaning *ex ante* Nash equilibria over the common prior of a Bayesian game), which had also been open questions.

To achieve this, we make use of a result that is interesting in its own right. In Theorem 1 we relate stationary equilibria of overlapping generations games, which form a subclass of the class of stochastic games, to Bayesian equilibria of certain Bayesian games. We thus have a bridge between strategic considerations in these very different classes of games, even though on one end of this bridge we have a dynamic time-dependent game with complete information and on the other end there is a static game with incomplete information.

We then achieve our goal of identifying a Bayesian game with no Bayesian approximate equilibria because we using the result in Theorem 1 to construct a Bayesian game based on an overlapping generations game with no stationary approximate equilibria presented in Levy (2012).

The significance of counter-examples to the existence of equilibria and approximate equilibria such as the example here (and those in Simon (2003) and Levy (2012)) is that they serve as a sharp warning signal to modellers: although you routinely assume the existence of equilibria when you work with finite games, you cannot automatically do so in games with an uncountable number of states. Matters are not so simple.

Consider, for one example, an extensively-used approach to dealing with a Bayesian game with a finite but large number of states is to analyse instead a similar game with a continuum of states. Myerson (1997), for example, informs readers of Chapter 2 of his textbook on game theory, when referring to Bayesian games, that “it is often easier to analyze examples with infinite type sets than those with large finite type sets.” Given this, it is important for modellers working with Bayesian games with uncountably many states to keep in mind that they cannot blindly rely on the well-known results in finite games guaranteeing the existence of equilibria and approximate equilibria.

⁴ This result does not contradict the result in Schmeidler (1973), which assumes that no deviation from equilibrium undertaken by a finite number (or even a measure zero set) of players can affect payoffs; we do not assume that here. Sion and Wolfe (1957) presents an example of a finite-player game with no equilibrium, but the example there assumes each player has a continuum of actions while we assume that each player has a finite action space.

2. PRELIMINARIES AND NOTATION

2.1. Knowledge Spaces.

A *space of states* is a pair (Ω, \mathcal{B}) composed of a *set of states* Ω and a σ -field \mathcal{B} of measurable subsets (*events*) of Ω .

Sets of *players* will be denoted here by I . Each player $i \in I$ will usually have an associated finite set of *actions* labelled A_i . Furthermore, denote $\bar{A} := \prod_{i \in I} A_i$ for the set of action profiles.

If we suppose that each player i has an associated partition Π_i of Ω , where for each state $\omega \in \Omega$ the element in Π_i that contains ω , $\Pi_i(\omega)$, is a measurable set, then $(\Omega, \mathcal{B}, (\Pi_i)_{i \in I})$ is a *knowledge space*. Denote by Γ_i the sub- σ -algebra of \mathcal{B} generated by Π_i .

The meet of the partition profile $(\Pi_i)_{i \in I}$ of the players is the finest partition that is coarser than the partition of each player. Each element of the meet is called a *common knowledge component*. For $\omega \in \Omega$, denote by $C(\omega)$ the common knowledge component containing ω .

2.2. Types and Priors.

Given a knowledge space $(\Omega, \mathcal{B}, (\Pi_i)_{i \in I})$, a *type function* t_i of player $i \in I$ is a function $t_i : \Omega \rightarrow \Delta(\Omega)$ from states to probability measures over (Ω, \mathcal{B}) such that the mapping $t_i(\cdot)$ satisfies:

- (1) $t_i(\omega)(E)$ is measurable for any fixed event E ,
- (2) $t_i(\omega)(\Pi_i(\omega)) = 1$,
- (3) $t_i(\omega) = t_i(\omega')$ for all $\omega' \in \Pi_i(\omega)$.

For each ω , $t_i(\omega)$ is called player i 's type at ω . Therefore, a quintuple $(\Omega, \mathcal{B}, I, (\Pi_i)_{i \in I}, (t_i)_{i \in I})$, where each t_i is a type function, is a *type space*.

A probability measure μ_i over (Ω, \mathcal{B}) is a *prior* for a type function t_i if for each event V

$$\mu_i(V) = \int_{\Omega} t_i(\omega)(V) d\mu_i(\omega). \quad (2.1)$$

When μ_i is a prior for t_i we also say that t_i is *induced* by μ_i and Γ_i , the sub- σ -algebra Γ_i of \mathcal{B} generated by the partition elements of t_i . In addition, t_i satisfies the condition of being a regular conditional probability of μ_i given Γ_i , meaning that for each $A \in \Gamma_i$ and $B \in \mathcal{B}$,

$$\mu_i(A \cap B) = \int_A t_i(x)(B) d\mu_i(x) \quad (2.2)$$

A probability measure μ that is a prior for each of the players' type function in a type space is a *common prior*.

2.3. Games.

2.3.1. *Stochastic Games.* A *stochastic game* is given by

$$(\Omega, \mathcal{B}, I, (A_i)_{i \in I}, r, \beta, q),$$

where Ω is space of states along with a given σ -field \mathcal{B} , I is a set of players, each player $i \in I$ has a finite set of actions A_i , r is a bounded measurable payoff function $r : \Omega \times \bar{A} \rightarrow \mathbb{R}^I$, β is a discount factor $\beta \in (0, 1)$ and q is a measurable transition function $q : \Omega \times \bar{A} \rightarrow \Delta\Omega$.

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 $a \in \bar{A}$, then $q(z, a)$ is the probability distribution of the next state of the game. A *behavioural strategy* for a player is a measurable mapping that associates each given history with a probability distribution on the set of actions available to him. A *stationary strategy* for a player i is a behavioural strategy that depends only on the current state; equivalently, it is a measurable mapping that associates each state $s \in \Omega$ with a probability distribution on the set A_i .

Let Σ^i denote the set of stationary strategies of player i , and further denote $\Sigma := \prod_{i \in I} \Sigma^i$. An element $\sigma \in \Sigma$ is a profile of stationary strategies. For each initial state $z \in \Omega$, a profile of stationary strategies σ determines an expected payoff $\gamma_\sigma^i(z)$ for each player in a standard manner. A profile $\sigma \in \Sigma$ is then called a *stationary ε -equilibrium*, for $\varepsilon \geq 0$, if for all $i \in I$, all $z \in \Omega$ and all $\tau \in \Sigma^i$,

$$\gamma_\sigma^i(z) \geq \gamma_{\tau, \sigma_{-i}}^i(z) - \varepsilon.$$

A stationary 0-equilibrium is usually simply called a stationary equilibrium, and when we use the term stationary equilibrium without further qualification we will mean a stationary 0-equilibrium.

2.3.2. *Overlapping Generations Games.* A *Levy overlapping generations game* (which we will simply call an overlapping generations game, or OGG, for short) is a stochastic game satisfying the following description. The set of players is countable; we will identify it with the set of integers \mathbb{N} . Let (Y, \mathcal{D}, μ_1) be a standard Borel space; it follows without loss of generality that we can assume (Y, \mathcal{D}, μ_1) is Borel isomorphic to the interval $[0, 1]$, with μ_1 the push-forward of the Lebesgue measure under that isomorphism. The state space of the overlapping generations game will be $\mathbb{N} \times Y$. We will furthermore assume the existence of a μ_1 measure-preserving⁵ mapping $f : Y \rightarrow Y$.

⁵ The measure-preserving assumption can be weakened but at the cost of less readable proofs in the paper.

Each player n has a finite action set A_n . In addition, there is a subset $Q_n \subset A_n$ such that if player n chooses an action $q \in Q_n$ then this effects a transition to a ‘quitting state’ $\bar{0}$. Alternatively, if player n chooses an action $a \in A_n \setminus Q_n$ then the transition from state (n, y) is deterministically effected to state $(n + 1, f(y))$, independently of the specific action a that is chosen.

In a state $(n, *)$, only player n ’s action has any effect on the payoffs; we think of him as the only ‘active’ player. Player n receives payoffs both when he is active, in state (n, y) , and in the subsequent state $(n + 1, f(y))$. Choosing to transition to the quitting state is akin to ‘blowing up the world’; if the quitting state $\bar{0}$ is ever attained following the actions of a player n then the payoffs to all players k for $k > n$ are zero.

This is an overlapping generations model because each player can be imagined as being alive for two generations. In the first generation, he is ‘young’ and takes an action, thus receiving some resulting payoff. In the second generation, he is ‘old’; he does not take any action but he does receive a payoff as a result of the ‘young’ player’s action.

The class of Levy overlapping generations games will be denoted by \mathcal{L} .

2.3.3. *Bayesian Games.* A Bayesian game is given by

$$(\Omega, \mathcal{B}, I, (\Pi_i)_{i \in I}, (t_i)_{i \in I}, (A_i)_{i \in I}, (r_i)_{i \in I}),$$

where $(\Omega, \mathcal{B}, I, (\Pi_i)_{i \in I}, (t_i)_{i \in I})$ is a type space, and for each i , A_i is player i ’s action set and r_i is a bounded measurable payoff function $r_i : \Omega \times \bar{A} \rightarrow \mathbb{R}$.

A Bayesian ε -equilibrium, for $\varepsilon \geq 0$, is a profile of strategies $\sigma = (\sigma^i)_{i \in I}$, where a strategy for player i is a mapping $\Omega \rightarrow \Delta(A_i)$ that is Γ_i -measurable, such that for each $i \in I$, each atom V of Γ_i , and each $x \in \Delta(A_i)$,

$$\int_V r_i(\omega, \sigma(\omega)) dt_i(\omega) \geq \int_V r_i(\omega, x, \sigma^{-i}(\omega)) dt_i(\omega) - \varepsilon$$

2.3.4. *Agents Games.*

Recall the definition of the agents game \mathbf{K} associated with a Bayesian game \mathbf{B} : \mathbf{K} is a strategic-form game whose set of players, which is a measurable space, has a (measurable) bijection η with the the set of all the types of all the players in \mathbf{B} . The action set of each player θ in \mathbf{K} is equal to the action set of the player j in \mathbf{B} associated with $\eta(\theta)$, and the payoff to player θ for an action profile is the corresponding expected payoff of j at $\eta(\theta)$. Every strategy $\hat{\psi}$ of \mathbf{B} is naturally associated in this way with a strategy ψ in \mathbf{K} .

The analysis of the equilibria of a Bayesian game \mathbf{B} can be accomplished by analysing the associated strategic-form game \mathbf{K} in the sense that, for any

$\varepsilon \geq 0$, if the strategy $\hat{\psi}$ is a (measurable) Bayesian ε -equilibria in \mathbf{B} then ψ in \mathbf{K} is a (measurable) Nash ε -equilibria.

3. OVERLAPPING GENERATIONS GAMES AND BAYESIAN GAMES

3.1. Main Theorems.

Theorem 1. *Let $\mathbf{L} \in \mathcal{L}$ be an overlapping generations game. Then there exists a Bayesian game \mathbf{B} such that there is an injective mapping from the the set of Bayesian ε -equilibria of \mathbf{B} to the set of stationary ε -equilibria of \mathbf{L} , for all $\varepsilon \geq 0$.*

The proof of Theorem 1 is in the appendix.

Theorem 2. *There exists a two-player Bayesian game with a continuum of states that admits no Bayesian ε -equilibria.*

Proof. Levy (2012) presents an overlapping generations game $\mathbf{L} \in \mathcal{L}$ that admits no stationary ε -equilibria. Theorem 1 then implies that the corresponding Bayesian game \mathbf{B} , as constructed in the proof of that theorem, admits no Bayesian ε -equilibria. ■

Corollary 3.1. *There is a two-player Bayesian game that admits no Bayesian equilibria.*

Proof. This follows by setting $\varepsilon = 0$ in Theorem 2. ■

Corollary 3.2. *There is a strategic-form game with a continuum number of players that admits no Nash ε -equilibrium.*

Proof. This follows by constructing the agents game associated with the Bayesian game of Theorem 2. ■

3.2. An Example of a Game with no Approximate Bayesian Equilibria.

For completeness, we present here an explicit construction of a game $\tilde{\mathbf{B}}$ with no approximate Bayesian equilibria, as guaranteed by Theorem 2. For ε , fix the value $0 < \varepsilon < \frac{1}{100}$. Let $Y := \{-1, 1\}^{\mathbb{Z}_{\geq 0}}$. The state space Ω of $\tilde{\mathbf{B}}$ is $\mathbb{N} \times Y$, with Y endowed with the standard Lebesgue measure.

The measure-preserving mapping is the Bernoulli shift $S : Y \rightarrow Y$, which maps each sequence $(x_0, x_1, x_2, \dots) \in Y$ to (x_1, x_2, x_3, \dots) , hence the mapping $h : \Omega \rightarrow \Omega$ of the proof of Theorem 1 is in this case

$$h(n, x_0, x_1, x_2, \dots) = (n + 1, x_1, x_2, x_3, \dots)$$

In addition, we have an operator $\iota : Y \rightarrow Y$ that is defined by $\iota(n, x_0, x_1, x_2, \dots) = (n, -1 \cdot x_0, x_1, x_2, \dots)$

There are two players, Player E and Player O . The action sets of the players are identical, namely the set $\{U, D\}$. The partitions of the players are as detailed in the proof of Theorem 1: if $\omega = (m, y)$ is an even state then

$$\begin{aligned}\Pi_E(\omega) &= \{\omega\} \cup h^{-1}(\omega) \\ \Pi_O(\omega) &= \{\omega, \iota(\omega), (m + 1, S(y))\}.\end{aligned}$$

If $\omega = (n, y)$ is an odd state with $n > 1$ then

$$\begin{aligned}\Pi_O(\omega) &= \{\omega\} \cup h^{-1}(\omega) \\ \Pi_E(\omega) &= \{\omega, \iota(\omega), (n + 1, S(y))\}.\end{aligned}$$

If $\omega = (1, y)$ then

$$\begin{aligned}\Pi_O(\omega) &= \{\omega\} \\ \Pi_E(\omega) &= \{\omega, \iota(\omega), (2, S(y))\}.\end{aligned}$$

The type functions t_E and t_O are simple: for any pair of states ω and ω' such that $\omega' \in \Pi_E(\omega)$, $t_E(\omega)(\{\omega'\}) = 1/3$ and similarly for any pair of states ω and ω' such that $\omega' \in \Pi_O(\omega)$, $t_O(\omega)(\{\omega'\}) = 1/3$.

The payoff functions r_E and r_O are as follows. $r_E(\omega, \cdot, \cdot) = 0$ for all odd states ω , and $r_O(\omega, \cdot, \cdot) = 0$ for all even states. If ω is an even state, then $r_E(\omega, a_1, a_2)$ is given by Table 1 with Player E the row player and Player O the column player. If ω is an odd state, then $r_O(\omega, a_1, a_2)$ is given by Table 1 with Player O the row player and Player E the column player.

If $x_0 = 1$:	If $x_0 = -1$:																		
<table border="1" style="border-collapse: collapse; margin: auto;"> <tr><td></td><td style="text-align: center;">U</td><td style="text-align: center;">D</td></tr> <tr><td style="text-align: center;">U</td><td style="text-align: center;">1</td><td style="text-align: center;">0</td></tr> <tr><td style="text-align: center;">D</td><td style="text-align: center;">0.3</td><td style="text-align: center;">0.3</td></tr> </table>		U	D	U	1	0	D	0.3	0.3	<table border="1" style="border-collapse: collapse; margin: auto;"> <tr><td></td><td style="text-align: center;">U</td><td style="text-align: center;">D</td></tr> <tr><td style="text-align: center;">U</td><td style="text-align: center;">0.7</td><td style="text-align: center;">0.7</td></tr> <tr><td style="text-align: center;">D</td><td style="text-align: center;">1</td><td style="text-align: center;">0</td></tr> </table>		U	D	U	0.7	0.7	D	1	0
	U	D																	
U	1	0																	
D	0.3	0.3																	
	U	D																	
U	0.7	0.7																	
D	1	0																	

TABLE 1. The payoff matrix.

This completes the description of a Bayesian game with no Bayesian ε -equilibrium.

3.3. Existence of Non-Measurable Equilibrium.

Throughout this paper, we have used the phrase ‘there is no equilibrium’ to mean ‘there is no measurable equilibrium’. We present here an explicit construction of a Bayesian equilibrium in the game $\tilde{\mathbf{B}}$ of Section 3.2; this equilibrium must, of course, be non-measurable⁶ given the previous results here.

⁶ The example in Simon (2003) also admits a non-measurable equilibrium, but that paper does not include an explicit construction of such an equilibrium.

Definition 3.1. Let $M \subset \Omega$ be a non-empty subset of the set of states $\Omega := \mathbb{N} \times Y$ of the game $\tilde{\mathbf{B}}$ of Section 3.2. Then an element ω is an *immediate neighbour* of M if (i) $\omega \notin M$ and (ii) for some $\omega' \in M$, $\omega = h(\omega')$ or $\omega \in h^{-1}(\omega')$.

Proposition 3.1. *For each state $\omega \in \Omega$ in the game $\tilde{\mathbf{B}}$, there exists a Bayesian equilibrium $\psi_\omega = (\psi_E, \psi_O)$ in the game restricted to $C(\omega)$, the common knowledge component of ω .*

The proof of Proposition 3.1 is in the appendix.

Corollary 3.3. *Let ψ_ω be defined simultaneously over all common knowledge components of $\tilde{\mathbf{B}}$ as in Proposition 3.1. Then ψ_ω is a non-measurable Bayesian equilibrium of $\tilde{\mathbf{B}}$.*

4. ROBUSTNESS TO PERTURBATIONS

For $\delta > 0$, an δ -perturbation of a Bayesian game \mathbf{B} is a Bayesian game \mathbf{B}' over the same type space and action sets, with a set of payoff functions v_i^ω satisfying $\|v_i^\omega - u_i^\omega\|_\infty < \delta$ for all $i \in I$. The example in Levy (2012) is robust to perturbations of the payoff functions of the overlapping generations game. It then follows immediately from Theorem 1 that the Bayesian game example presented in Section 3 is also robust to sufficiently small perturbations.

A similar result holds for sufficiently small perturbations of the posterior probabilities defining the types t_E and t_O .

5. HARSÁNYI ε -EQUILIBRIA

An Harsányi ε -equilibrium of a Bayesian game with a common prior μ is a profile of mixed strategies $\Psi = (\Psi_i)_{i \in I}$ such that for each player i and any unilateral deviation strategy $\hat{\Psi}_i$,

$$\int_{\Omega} u_i^\omega(\Psi(\omega)) d\mu(\omega) \geq \int_{\Omega} u_i^\omega(\hat{\Psi}_i(\omega), \Psi_{-i}(\omega)) d\mu(\omega) - \varepsilon.$$

Simon (2003) shows that the existence of a measurable 0-Harsányi equilibrium implies the existence of a measurable 0-Bayesian equilibrium. However, this result is known not to hold for $\varepsilon > 0$. The example in this paper does not, therefore, imply that there is no Harsányi ε -equilibrium. We leave the open question of whether or not there are examples of games that have no measurable Harsányi ε -equilibria for future research.

Furthermore, in the example in Section 3 there is incomplete information on both sides: if we define the parity of the state and the value of x_0 to comprise the state of nature at a state of the world ω then neither player knows the true state of nature. Because of this ignorance of the state of nature and the way the payoff matrices are defined in Table 1, neither player ever knows the true payoffs.

This contrasts with the example in Simon (2003), where there is incomplete information on one and a half sides (that is, one player always knows the true payoff relevant data but not always what the other players might know) and by construction players know their own payoffs at each state. It is presently unknown whether an example can be constructed of a game in which players know their payoffs at each state but the game has no Bayesian ε -equilibrium.

6. APPENDIX: PROOFS

Proof of Theorem 1. The proof first uses the properties of \mathbf{L} to construct an agents game \mathbf{K} and then uses that to construct the desired Bayesian game \mathbf{B} .

Step 1. Denote $\Omega := \mathbb{N} \times Y$ for the set of states of \mathbf{L} and let $h : \Omega \rightarrow \Omega$ be defined as $h(n, y) := (n + 1, f(y))$. For each state $\omega = (n, y)$, define $\alpha(\omega) = n$, and we will call player n the *active player* at state ω .

Recall that by definition in an OGG at each individual time only the actions of the active player n has any effect on the payoffs. However, in coming to choose his action, the active player n must also consider the action of player $n + 1$, because that too can influence his payoff.

More explicitly, let $\omega = (n, y)$ be a state with active player $\alpha(\omega) = n$. Let $\bar{a} = (a_1, a_2, \dots) \in \bar{A}$ be an action profile. Then the payoff function $r^k(\omega, \bar{a}) = 0$ if $k \neq n, n + 1$. Furthermore, $r^n(\omega, \bar{a})$ depends only on y, a_{n-1}, a_n and a_{n+1} (it depends on a_{n-1} because if the previous player chooses a transition to quitting state $\bar{0}$ then player n receives a zero payoff). If the state is not $\bar{0}$ then when it is player n 's turn to be active, $r^n(\omega, \bar{a})$ can be considered as being composed of two components: $p^n(y, a_n)$, which depends only on y and $a_n \in A_n$, and $p_+^n(f(y), a_n, a_{n+1})$, which depends only on $f(y)$, the action $a_{n+1} \in A_{n+1}$ chosen by player $n + 1$ and a_n (if player n chooses to quit).

The total payoff to player n is thus

$$r^n(\omega, \bar{a}) = p^n(y, a_n) + p_+^n(f(y), a_n, a_{n+1}).$$

In other words, player n strategically is playing a strategic-form game against player $n + 1$, which can be represented as a matrix M^ω , with player n the row player and player $n + 1$ the column player:

$$M^\omega(a_n, a_{n+1}) = \begin{cases} p^n(y, a_n) & \text{if } a_n \in Q_n \\ p^n(y, a_n) + p_+^n(f(y), a_n, a_{n+1}) & \text{if } a_n \in A_n \setminus Q_n. \end{cases}$$

In these terms, if $\sigma(\omega)$ (where $\sigma : \omega \mapsto \Delta(A_{\alpha(\omega)})$ is a measurable function) is an ε -best response, from the perspective of the row player, in the strategic-form game given by M^ω for every $\omega \in \Omega$ then σ is a stationary ε -equilibrium of \mathbf{L} .

Step 2. We next use the matrices M^ω to construct an agents game \mathbf{K} . As \mathbf{K} is a strategic-form game, we need to specify the set of agents, their action sets and their payoff functions.

The set of agents of \mathbf{K} is Ω , the set of states of \mathbf{L} . The action set of agent $\omega \in \Omega$ is $A_{\alpha(\omega)}$, that is, the action set of the active player at ω in the game \mathbf{L} .

In \mathbf{K} all the agents simultaneously choose an action from their action sets. A (pure) action profile is therefore given by an element of the set $\prod_{\omega \in \Omega} A_{\alpha(i)}$. Given an action profile \bar{a} in this set, the payoff to agent $\omega = (n, y)$ is determined solely by the action that he has chosen and the action chosen by agent $h(\omega) = (n+1, f(y))$ and is explicitly given by the matrix M^ω with agent ω the row player and agent $h(\omega)$ the column player.

A measurable function $\psi : \omega \mapsto \Delta(A_{\alpha(\omega)})$ is a Nash ε -equilibrium of \mathbf{K} if and only if $\psi(\omega)$ is an ε -best response, from the perspective of the row player in M^ω , for every $\omega \in \Omega$. It then immediately follows by construction that if ψ is a Nash ε -equilibrium of \mathbf{K} then ψ is at the same time also a stationary ε -equilibrium of \mathbf{L} .

Step 3. What remains is constructing a Bayesian game \mathbf{B} whose agents game is \mathbf{K} . The state space of \mathbf{B} is Ω . States of the form $(n, *)$ for n odd will be called odd states, and states of the form $(m, *)$ for m even will be called even states. There are two players in \mathbf{B} , labelled Player O (for odd) and Player E (for even).

Recall that Y is endowed with a probability measure μ_1 . Let μ_2 be the probability measure over \mathbb{N} that assigns probability $1/2^n$ to each $n \in \mathbb{N}$. We will use $\mu := \mu_1 \times \mu_2$ as a probability measure over Ω , which will serve as the common prior of \mathbf{B} .

We next define the partitions Π_O and Π_E of the players. If $\omega = (m, y)$ is an even state then

$$\Pi_E(\omega) = \{\omega\} \cup h^{-1}(\omega)$$

$$\Pi_O(\omega) = \{\omega\} \cup h(\omega) \cup h^{-1}(h(\omega))$$

Note that if ω is an even state then $\Pi_E(\omega)$ contains only one even state (because all the states in $h^{-1}(\omega)$ are odd) but $\Pi_O(\omega)$ may contain many even states (in $h^{-1}(h(\omega))$) but only one odd state, $h(\omega)$.

If $\omega = (n, y)$ is an odd state with $n > 1$ then

$$\Pi_O(\omega) = \{\omega\} \cup h^{-1}(\omega)$$

$$\Pi_E(\omega) = \{\omega\} \cup h(\omega) \cup h^{-1}(h(\omega))$$

In a manner similar to the above, in this case $\Pi_O(\omega)$ contains only one odd state (because all the states in $h^{-1}(\omega)$ are even) but $\Pi_E(\omega)$ may contain many odd states (in $h^{-1}(h(\omega))$) but only one even state, $h(\omega)$.

The odd states of the form $(1, *)$ are treated separately only because h^{-1} is not defined over them: if $\omega = (1, y)$ then

$$\Pi_O(\omega) = \{\omega\}$$

$$\Pi_E(\omega) = \{\omega\} \cup h(\omega) \cup h^{-1}(h(\omega)).$$

The type function t_E for **B** is the function assigning probability $1/3$ to the sole even element of every partition element of Π_E and uniformly assigning probability $2/3$ over the set of the odd elements of Π_E . Similarly, the type function t_O is the function assigning probability $1/3$ to the sole odd element of every partition element of Π_O and uniformly assigning probability $2/3$ over the set of the even elements of Π_E . Lemma 6.1 shows that these functions are indeed the type functions induced by μ and the σ -algebras generated by Π_E and Π_O .

Finally, we define the payoff functions. Let ρ_O be Player O 's payoff function and ρ_E be Player E 's payoff function.

- For all even states ω , $\rho_O(\omega, \cdot) = 0$ and similarly for all odd states ω , $\rho_E(\omega, \cdot) = 0$.
In words, at even (respectively, odd) states, Player O (respectively, Player E) gets payoff 0 *no matter what* actions are played by him or the other player; strategically, Player E (respectively, Player O) can ignore the odd (respectively, even) states.
- For all odd states ω , $\rho_O(\omega, \cdot)$ is determined by the matrix $3M^\omega$, where Player O is row and Player E is column. This means that Player O cannot strategically ignore the odd states.
- For all even states ω , $\rho_E(\omega, \cdot)$ is also determined by the matrix $3M^\omega$, but now Player E is row and Player O is column. As above, Player E cannot strategically ignore the even states.

This completes the construction of the Bayesian game **B**.

All that remains is showing that **K** is the agents game corresponding to **B**. A type of Player E always contains only one even state while a type of Player O contains only one odd state. We can therefore *uniquely* identify each type of Player E by the even state it contains and each type of Player O by the odd state it contains. Hence we can consider Player E 's agents to be the set of even states and by the same reasoning Player O 's agents are the set of odd states.

Formally, there is a bijection η between Ω and the collection of types $\{\Pi_E(\omega)\}_{\omega \in \Omega} \cup \{\Pi_O(\omega)\}_{\omega \in \Omega}$ as follows:

$$\eta(\omega) = \begin{cases} \Pi_E(\omega) & \text{if } \omega \text{ is even} \\ \Pi_O(\omega) & \text{if } \omega \text{ is odd.} \end{cases}$$

The payoff functions of these agents under this bijection are precisely given by the matrices M^ω . To see this, consider for example Player E 's perspective when the true state of the world is any state in $\Pi_E(\omega)$, where ω is an even state. By construction, ω is the only even state state in $\Pi_E(\omega)$. Since Player E 's action choice can lead to a non-zero payoff for him only at ω , for calculating his expected

payoff at $\Pi_E(\omega)$ he needs only to plan a best reply to Player O 's actions at ω . Furthermore, since $t_E(\omega) = 1/3$, Player E 's payoff is determined by the matrix $1/3 \cdot 3M^\omega = M^\omega$. However, from Player O 's perception, if the true state is ω then the only state in $\Pi_O(\omega)$ that is payoff-relevant to her is $h(\omega)$. Hence in the agents game corresponding to \mathbf{B} agent ω is playing the matrix M^ω against agent $h(\omega)$; this is exactly the description of \mathbf{K} .

It is here that we see explicitly the connection between the strategic reasoning in \mathbf{B} and \mathbf{K} . In the dynamic game \mathbf{K} , at time 1 player 1 is informed that the state is $(1, y)$ and he best replies to the perceived action of player 2 at state $(2, f(y))$, who in turn best replies to the perceived action of player 3 at state $(3, f^2(y))$ and so on. In the static Bayesian game \mathbf{B} , if the true state is $(1, y)$ Player O knows this and best replies to the perceived actions of Player E at state $(2, f(y))$, who in turn best replies to the perceived action of Player O at state $(3, f^2(y))$ and so on.

Letting $\hat{\psi} = (\hat{\psi}_1, \hat{\psi}_2)$ be a strategy profile $\hat{\psi}$ in \mathbf{B} , define a strategy profile ψ in \mathbf{K} by

$$\psi(\omega) = \begin{cases} \hat{\psi}_1(\omega) & \text{if } \omega \text{ is odd} \\ \hat{\psi}_2(\omega) & \text{if } \omega \text{ is even.} \end{cases}$$

The proof is now complete: if $\hat{\psi}$ is a Bayesian ε -equilibrium of \mathbf{B} then ψ is a Nash ε -equilibrium of \mathbf{K} and therefore ψ is also a stationary ε -equilibrium of \mathbf{L} . ■

Lemma 6.1. *The functions t_E and t_O as defined in the proof of Theorem 1 satisfy the conditions for being type functions with μ as their common prior.*

Proof of Lemma 6.1. By construction, for $i \in \{O, E\}$, $t_i(\omega)(\Pi_i(\omega)) = 1$ for all ω and $t_i(\omega) = t_i(\omega')$ for $\omega' \in \Pi_i(\omega)$. Two more items need to be checked: that for each event A , $t_i(\omega)(A)$ is measurable and that $\mu(A) = \int_{\Omega} t_i(\omega)(A) d\mu(\omega)$. We will prove these for $i = E$, with the proof for $i = O$ conducted similarly.

For the rest of this proof, denote by $1_A(\omega)$ the indicator function that returns 1 if $\omega \in A$ and 0 if $\omega \notin A$. Furthermore, ‘even’ will refer to the set of even states of Ω and ‘odd’ will refer to the set of odd states of Ω . Fix an event A . Then:

$$t_E(\omega)(A) = \begin{cases} \frac{2}{3} \int_{h^{-1}(h(\omega))} 1_A(\omega') d\mu_1(\omega') + \frac{1}{3} 1_A(h(\omega)) & \text{if } \omega \in \text{odd} \\ \frac{2}{3} \int_{h^{-1}(\omega)} 1_A(\omega') d\mu_1(\omega') + \frac{1}{3} 1_A(\omega) & \text{if } \omega \in \text{even.} \end{cases}$$

Note there is no problem integrating here only with respect to μ_1 rather than μ , since for any $\omega = (n, y)$, $h^{-1}(h(\omega)) \subset n \times Y$ and $h^{-1}(\omega) \subset (n-1) \times Y$. We conclude that $t_E(\omega)(A)$ is measurable.

Next, let $B \subseteq \Omega$ be saturated. If $A \subseteq \text{even}$ then:

$$\begin{aligned} & \int_{\Omega} 1_B(\omega) t_E(\omega)(A) d\mu(\omega) \\ &= \int_{\text{even}} 1_B(\omega) \left(\frac{1}{3} 1_A(\omega) \right) d\mu(\omega) + 2 \int_{\text{odd}} 1_B(h(\omega)) \left(\frac{1}{3} 1_A(h(\omega)) \right) d\mu(\omega) \end{aligned} \quad (6.1)$$

$$= \frac{1}{3} \int_{\Omega} 1_B(\omega) 1_A(\omega) d\mu(\omega) + \frac{2}{3} \int_{\Omega} 1_B(h(\omega)) 1_A(h(\omega)) d\mu(\omega) \quad (6.2)$$

$$\begin{aligned} &= \int_{\Omega} 1_B(\omega) 1_A(\omega) d\mu(\omega) \quad (6.3) \\ &= \mu(A \cap B). \end{aligned}$$

Note that $1_B(\omega) = 1_B(h(\omega))$ because B is saturated, hence the equality between Equations (6.2) and (6.3). The integration in the right term in Equation (6.1) is multiplied by 2 because μ assigns half as much measure to the image of a set under the mapping h than it does to the original set. Furthermore, since $A \subseteq \text{even}$, $1_A(\omega) = 0$ for $\omega \in \text{odd}$ and $1_A(h(\omega)) = 0$ for $\omega \in \text{even}$, justifying the move to integrating both terms of Equation 6.2 over Ω .

If $A \subseteq \text{odd}$ then:

$$\begin{aligned} & \int_{\Omega} 1_B(\omega) t_E(\omega)(A) d\mu(\omega) \\ &= \int_{\text{odd}} 1_B(\omega) \left(\frac{2}{3} 1_A(\omega) \right) d\mu(\omega) + \int_{\text{even}} 1_B(\omega) \left(\int_{h^{-1}(\omega)} \frac{2}{3} 1_A(\omega') d\mu(\omega') \right) d\mu(\omega) \\ &= \frac{2}{3} \int_{\text{odd}} 1_B(\omega) 1_A(\omega) d\mu(\omega) + \frac{1}{2} \frac{2}{3} \int_{\text{odd}} 1_B(\omega) 1_A(\omega) d\mu(\omega) \end{aligned} \quad (6.4)$$

$$= \frac{2}{3} \int_{\Omega} 1_B(\omega) 1_A(\omega) d\mu(\omega) + \frac{1}{3} \int_{\Omega} 1_B(\omega) 1_A(\omega) d\mu(\omega) \quad (6.5)$$

$$\begin{aligned} &= \int_{\Omega} 1_B(\omega) 1_A(\omega) d\mu(\omega) \quad (6.6) \\ &= \mu(A \cap B). \end{aligned}$$

This is sufficient for completing the proof. ■

Proof of Proposition 3.1. In this proof, we will call action U *the opposite* of action D and D the opposite of action U .

Note that, given the construction of $\tilde{\mathbf{B}}$ in Section 3.2, $C(\omega)$ is either finite or countably infinite for any ω , and for almost all ω it is countably infinite. Furthermore, it is straightforward inductively to build an enumeration $\omega_1, \omega_2, \dots$ of $C(\omega)$ (such that $C(\omega) = \bigcup \omega_n$) by setting $\omega_1 = \omega$, $M_1 = \omega$, then given M_n for all $n \geq 1$ letting ω_{n+1} be an immediate neighbour of M_n and finally setting $M_{n+1} = M_n \cup \omega_{n+1}$.

Define an equilibrium by induction along this enumeration of Ω , as follows. Without loss of generality suppose that ω_1 is an odd state. Let $\psi_O(\omega_1) \in \{U, D\}$ be selected arbitrarily.

Suppose inductively that ψ_O has been defined for all odd states in the sequence $\omega_1, \omega_2, \dots, \omega_n$ and ψ_E has been defined for all even states in $\omega_1, \omega_2, \dots, \omega_n$. ω_{n+1} is an immediate neighbour of $M_n = \omega_1 \cup \omega_2 \dots \cup \omega_n$. By definition, there is an $\omega' = (k, x_0, x_1, x_2, \dots) \in M_n$ such that either $\omega = h(\omega')$ or $\omega \in h^{-1}(\omega')$.

Suppose that k is odd (the same reasoning would hold if it is even). If $\omega_{n+1} = h(\omega')$, then $\omega_{n+1} = (k + 1, x_1, x_2, x_3, \dots)$. If $x_1 = 1$, then set $\psi_E(\omega_{n+1}) = \psi_O(\omega')$. If $x_1 = -1$, then set $\psi_E(\omega_{n+1})$ to be the opposite of $\psi_O(\omega')$.

If $\omega_{n+1} \in h^{-1}(\omega')$, then $\omega_{n+1} = (k - 1, x_{-1}, x_0, x_1, x_2, x_3, \dots)$. If $x_{-1} = 1$, then set $\psi_E(\omega_{n+1}) = \psi_O(\omega')$. If $x_{-1} = -1$, then set $\psi_E(\omega_{n+1})$ to be the opposite of $\psi_O(\omega')$. This completes the definition of ψ_E and ψ_O over all of $C(\omega)$.

The matrices in Table 1 indicate that for any state $\omega' = (k, x_0, x_1, x_2, \dots)$, if $x_0 = 1$ the best reply of the active player at ω' is to match the action of the player at $h(\omega')$, while if $x_0 = -1$ the best reply of the active player at ω' is to mis-match the action of the player at $h(\omega')$. By construction, ψ_E and ψ_O ensure that each player in the game restricted to $C(\omega)$ is best-replying in this way. Hence they form a Bayesian equilibrium. ■

REFERENCES

- Chatterjee, K., and W. Samuelson (1973), Bargaining Under Incomplete Information, *Operations Research*, 5, 835–851.
- Harsányi, J. C. (1967), Games with Incomplete Information Played by Bayesian Players, *Management Science*, 14, 159–182.
- Levy, Y., (2012), A Discounted Stochastic Game with No Stationary Nash Equilibrium, *Discussion Paper 596*, The Centre for the Study of Rationality, Hebrew University of Jerusalem.
- Myerson, R. (1997), *Game Theory: Analysis of Conflict*, Cambridge, Harvard University Press.
- Radner, R. (1980), Collusive Behavior in Noncooperative Epsilon-Equilibria of Oligopolies with Long but Finite Lives, *Journal of Economic Theory*, 22, 136–154.
- Schmeidler, D. (1973), Equilibrium Points of Nonatomic Games, *Journal of Statistical Physics*, 7 (4), 295–300.
- Simon, R. S., (2003), Games of Incomplete Information, Ergodic Theory and the Measurability of Equilibria, *Israel Journal of Mathematics*, (138), 73–92.
- Sion, M., and P. Wolfe (1957), On a Game with No Value, *Annals of Mathematical Studies*, (39), 299–306.