Borel games with lower-semi-continuous payoffs

Maastricht University

János Flesch, Jeroen Kuipers,
Ayala Mashiah-Yaakovi,
Gijs Schoenmakers, Eilon Solan,
Koos Vrieze

Borel Games with Lower-Semi-Continuous Payoffs

RM/10/041

METEOR

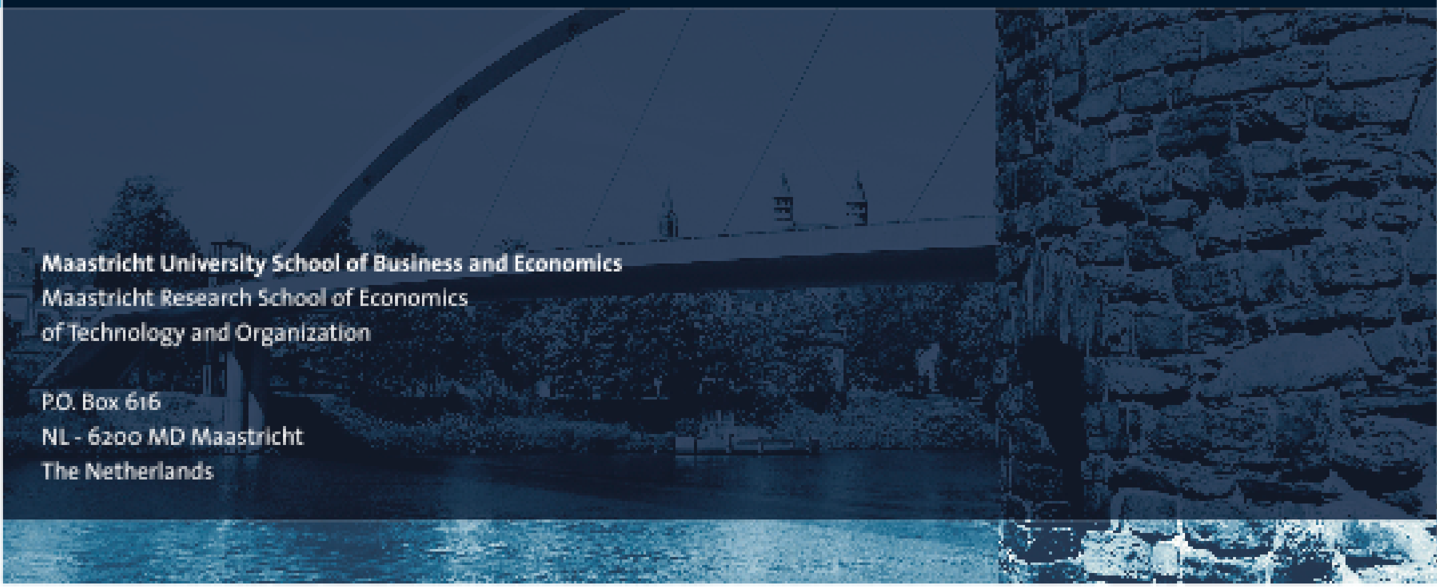

Maastricht University School of Business and Economics
Maastricht Research School of Economics
of Technology and Organization

P.O. Box 616
NL - 6200 MD Maastricht
The Netherlands

Borel Games with Lower-Semi-Continuous Payoffs*†

János Flesch‡, Jeroen Kuipers§, Ayala Mashiah-Yaakovi¶, Gijs Schoenmakers§, Eilon Solan¶ and Koos Vrieze§

March 7, 2010

Abstract

We prove that every multi-player Borel game with bounded and lower-semi-continuous payoffs admits a subgame-perfect ε-equilibrium in pure strategies. This result complements Example 3 in Solan and Vieille (2003), which shows that a subgame-perfect ε-equilibrium in pure strategies need not exist when the payoffs are not lower-semi-continuous. In addition, if the range of payoffs is finite, we characterize in the form of a Folk Theorem the set of all plays and payoffs that are induced by subgame-perfect 0-equilibria in pure strategies.

1 Introduction

A multi-player Borel game is a sequential game with perfect information and without chance moves. The payoff of each player is a function of the infinite sequence of actions that the players choose. Borel games were introduced by Gale and Stewart (1953), who studied two-player zero-sum games where the payoff function is the indicator of some set. In other words, player 1 wins if the play generated by the players is in a given set of plays, and player 2 wins otherwise. Martin (1975) proved that if the winning set of player 1 is Borel measurable, then the game is determined: either player 1 has a winning strategy or player 2 has a winning strategy. This result implies that every two-player

*The research of Mashiah-Yaakovi and Solan was supported by the Israel Science Foundation (grant number 212/09).

†The research of Mashiah-Yaakovi was partially supported by the Farajun Foundation Fellowship.

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

§Dept. of Knowledge Engineering, Maastricht University, P.O. Box 616, 6200MD, Maastricht, The Netherlands. E-mail: kuipers@maastrichtuniversity.nl, gm.schoenmakers@maastrichtuniversity.nl, koos@ictvrieze.nl.

¶School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel. E-mail: ayalam@post.tau.ac.il, eilons@post.tau.ac.il.

zero-sum Borel game has a value, provided the payoff function is bounded and measurable.

Mertens and Neyman (see Mertens, 1987) used the existence of the value in two-player zero-sum Borel games to prove that for every $\varepsilon > 0$, every multi-player non-zero-sum Borel game has a pure ε-equilibrium, provided the payoff functions are bounded and Borel measurable. Roughly, the ε-equilibrium strategies constructed by Mertens and Neyman are as follows: each player i starts by following an $\frac{\varepsilon}{2}$-optimal strategy in an auxiliary two-player zero-sum game G_i, where the payoff is that of player i, player i is the maximizer and the other players try to minimize player i's payoff. This goes on as long as no player deviates. Once some player, say player i, deviates, the other players switch to an $\frac{\varepsilon}{2}$-optimal strategy of the minimizers in the game G_i.

Thus, the players start by generating a play that yields all of them a high payoff, and, if a player deviates, he is punished with a low payoff. This construction has the disadvantage that in the punishment phase, the punishers play without regard to their own payoffs. Therefore, in real-life situations, players may be reluctant to follow the equilibrium strategies constructed by Mertens and Neyman.

To deal with such non-credible threats of punishment, Selten (1965, 1973) introduced the concept of subgame-perfect equilibrium. A strategy vector is a subgame-perfect ε-equilibrium if it induces an ε-equilibrium after any possible finite history of actions. Ummels (2005) proved the existence of a subgame-perfect equilibrium in pure strategies for multi-player Borel games when the payoff function of each player is the indicator of some Borel set. His proof is based on the following recursive construction. First, one identifies all finite histories which are a winning position to at least one of the players; that is, if this finite history occurs, one of the players can ensure that his payoff is 1. After such finite histories, one instructs every winning player to play his winning strategy. This leads to a pruned game where all moves which are excluded by these winning strategies are eliminated. One subsequently identifies winning positions to the players in this new game, and prunes it in a similar way. The process repeats itself, until it reaches a stable state. Ummels proves that a combination of remaining strategies is a subgame-perfect equilibrium of the original game.

In the present paper we show that every multi-player Borel game with bounded and lower-semi-continuous payoffs admits a pure subgame-perfect ε-equilibrium, for every $\varepsilon > 0$. This result complements Example 3 in Solan and Vieille (2003) that shows that when the payoff function of at least one player is not lower-semi-continuous, a pure subgame-perfect ε-equilibrium need not exist. Our proof makes use of a transfinite construction. A different type of transfinite construction was used by Maitra and Sudderth (1993) to prove the existence of the value in a certain class of stochastic games. In Section 4.2 we point at another possible application of our technique.

The determinacy of Borel games has attracted a lot of attention in descriptive set theory (see, e.g., Schilling and Vaught (1983) and Kechris (1995)). A rich literature identifies winning positions for the two players in the class of games that are played on graphs (see Grädel (2004) for a survey). Two-player zero-sum Borel games were used in the computer science literature to study reactive non-terminating programs (see, e.g., Thomas (2002)) and model checking in μ-calculus (see, e.g., Emerson et al. (2001)), and in economics to show that measurable tests are manipulable (Shmaya, 2008).

Our result also relates to the game theoretic literature that studies the existence of a subgame-perfect ε-equilibrium in various classes of infinite games, see, e.g., Solan and Vieille (2003), Mashiah-Yaakovi (2009), Kuipers et al. (2008) or Flesch et al. (2010). In particular, our result generalizes results in Flesch et al. (2010).

The paper is organized as follows. The model and the main result appear in Section 2. Section 3 contains the proof of the main result, and Section 4 concludes with comments.

2 The Model and the Main Result

Definition 1. *An n-player* Borel game *is a quadruple* $(I, A, i, (u^j)_{j\in I})$ *where* $I = \{1, 2, \ldots, n\}$ *is the set of players,* A *is a non-empty set*[1] *of actions,* $i : \bigcup_{t\in\mathbf{N}} A^{t-1} \to I$ *is a function*[2] *that assigns an active player to each finite sequence of actions, and* $u^j : A^{\mathbf{N}} \to \mathbf{R}$ *is the payoff function, for every player* $j \in I$.

A Borel game is a sequential game, where at each stage $t \in \mathbf{N}$, knowing the past history $h_t = (a_1, a_2, \ldots, a_{t-1})$, player $i(h_t)$, the active player at stage t, chooses an action $a_t \in A$. The payoff to each player $j \in I$ is $u^j(a_1, a_2, \ldots)$. The description of the game is common knowledge among the players.

Comment 2. *The assumption that the action set is the same for all players and for all stages is made for simplicity of notations only. Nothing that is said below would be affected if the action sets would depend on the player, on the stage, or in fact on the whole past play.*

The set of finite histories where player j is the active player is:

$$\mathcal{H}^j = i^{-1}(j) = \left\{h \in \cup_{t\in\mathbf{N}} A^{t-1} : i(h) = j\right\}.$$

The set of all finite histories is then $\mathcal{H} = \cup_{j\in I}\mathcal{H}^j$.

Definition 3. *A* (pure) strategy *for player* j *is a function* $\sigma^j : \mathcal{H}^j \to A$. *A* (pure) strategy profile *is a vector of strategies* $\sigma = (\sigma^j)_{j\in I}$.

In the present paper we discuss only pure strategies, and by a strategy or by a strategy profile we will always mean a pure one. We denote by Σ^j the strategy space of player j, and by $\Sigma = \times_{j\in I}\Sigma^j$ the set of all strategy profiles.

[1]The set of actions A may be finite or infinite.

[2]By convention, the initial history is the empty history $h_1 = \emptyset$, and $A^0 = \{\emptyset\}$.

An infinite sequence of actions $p \in A^{\mathbf{N}}$ is called a play. Every strategy profile $\sigma \in \Sigma$ determines a unique play $p(\sigma) = (a_t)_{t\in\mathbf{N}} \in A^{\mathbf{N}}$ recursively as follows:

$$\begin{aligned} a_1 &= \sigma^{i(\emptyset)}(\emptyset), \\ a_t &= \sigma^{i(h_t)}(h_t), \quad \text{where } h_t = (a_1, a_2, \ldots, a_{t-1}),\ \forall t \in \mathbf{N}. \end{aligned}$$

We denote by $u^j(\sigma) = u^j(p(\sigma))$ the payoff of player j when the players follow σ.

For $j \in I$ we denote by $-j = I \setminus \{j\}$ the set of all players excluding j. If σ is a strategy profile and j is a player, then $\sigma^{-j} = (\sigma^k)_{k\in I\setminus\{j\}}$.

Definition 4. *Let $\varepsilon \geq 0$. A strategy profile $\sigma_* = (\sigma_*^j)_{j\in I}$ is an* ε-equilibrium *if $u^j(\sigma_*) \geq u^j(\sigma_*^{-j}, \sigma^j) - \varepsilon$ for every player $j \in I$ and every strategy $\sigma^j \in \Sigma^j$.*

Throughout the paper we endow A with the discrete topology, and $A^{\mathbf{N}}$ with the product topology.

The two-player game is called *zero-sum* if $u^1(p)+u^2(p) = 0$ for every $p \in A^{\mathbf{N}}$. The result of Martin (1975) implies that in zero-sum games, an ε-equilibrium exists for every $\varepsilon > 0$ under quite general conditions.

Theorem 5. *If the game is zero-sum, and if u^1 is bounded and Borel measurable, then an ε-equilibrium exists for every $\varepsilon > 0$.*

This results implies the existence of an ε-equilibrium in every multi-player Borel game.

Theorem 6 (Mertens and Neyman, see Mertens, 1987)**.** *If u^j is bounded and Borel measurable for every player $j \in I$, then an ε-equilibrium exists for every $\varepsilon > 0$.*

A stronger notion of equilibrium is the notion of subgame-perfect equilibrium. Every finite history $h = (a_1, a_2, \ldots, a_l) \in \mathcal{H}$, together with a strategy profile σ, determines an infinite play $p(\sigma \mid h) = (b_t)_{t\in\mathbf{N}} \in A^{\mathbf{N}}$ recursively as follows:

$$\begin{aligned} b_t &= a_t, \quad 1 \leq t \leq l, \\ b_t &= \sigma^{i(h_t)}(h_t), \quad \text{where } h_t = (b_1, b_2, \ldots, b_{t-1}),\ l < t. \end{aligned}$$

This is the play that σ generates given that the history h occurred. We denote by $u^j(\sigma \mid h) = u^j(p(\sigma \mid h))$ the payoff of player j at this play.

Definition 7. *Let $\varepsilon \geq 0$. A strategy profile $\sigma_* = (\sigma_*^j)_{j\in I}$ is a* subgame-perfect ε-equilibrium *if for every finite history $h \in \mathcal{H}$, one has*

$$u^j(\sigma_* \mid h) \geq u^j((\sigma_*^{-j}, \sigma^j) \mid h) - \varepsilon \quad \forall j \in I \quad \forall \sigma^j \in \Sigma^j.$$

In other words, every finite history h defines the subgame that starts at h — namely, the game with payoff function $u^j(\cdot \mid h)$ for player $j \in I$. A strategy profile is a subgame perfect ε-equilibrium if it induces an ε-equilibrium in all subgames.

We say that a finite history $h = (a_t)_{t=1}^l$ is a *prefix* of the play $p = (b_t)_{t\in\mathbf{N}} \in A^{\mathbf{N}}$, or that p is an *extension* of h, if $a_t = b_t$ for every $t \in \{1, 2, \ldots, l\}$, and we denote it by $h \prec p$. We say that a finite history $h = (a_t)_{t=1}^l$ is a *prefix* of the finite history $h' = (b_t)_{t=1}^m$, or that h' is an *extension* of h, if $l \leq m$ and $a_t = b_t$ for every $t \in \{1, 2, \ldots, l\}$, and we denote it by $h \preceq h'$.

When A is endowed with the discrete topology, and $A^{\mathbf{N}}$ is endowed with the product topology, then a sequence $(p^k)_{k\in\mathbf{N}}$ of plays converges to a limit p if and only if every prefix h of p is a prefix of all the plays $(p^k)_{k\in\mathbf{N}}$ except possibly of finitely many of them.

Definition 8. *The payoff function u^j is* lower-semi-continuous *if for every sequence $(p^k)_{k\in\mathbf{N}}$ of plays in $A^{\mathbf{N}}$ that converges to a limit p one has*

$$\liminf_{k\to\infty} u^j(p^k) \geq u^j(p).$$

Note that every lower-semi-continuous function is Borel measurable. Our main result is the following.

Theorem 9. *If the payoff function u^j is bounded and lower-semi-continuous for every player $j \in I$, then the game admits a subgame-perfect ε-equilibrium for every $\varepsilon > 0$.*

This result is tight, in the sense that if the payoff function of one of the players is not lower-semi-continuous, then the game need not admit a subgame-perfect ε-equilibrium for every $\varepsilon > 0$ (see Example 3 in Solan and Vieille (2003)).

3 Proof of Theorem 9 and a Folk Theorem

We first argue that we can assume w.l.o.g. that the range of the payoff functions $(u^j)_{j\in I}$ is finite. Indeed,[3] let $\widehat{u}^j(p)$ be the highest multiple of ε that is *strictly* smaller than $u^j(p)$:

$$\widehat{u}^j(p) = \varepsilon \lfloor \frac{u^j(p)}{\varepsilon} \rfloor.$$

Note that if u^j is bounded then $\widehat{u}^j$ has finite range, and if u^j is lower-semi-continuous then so is $\widehat{u}^j$. Moreover, every subgame-perfect ε-equilibrium in the game with payoff functions $(\widehat{u}^j)_{j\in I}$ is a subgame-perfect 2ε-equilibrium in the game with payoff functions $(u^j)_{j\in I}$. Therefore, for the proof of Theorem 9, we may assume w.l.o.g. that the payoff functions have finite range.

From now on, we assume that the payoff functions $(u^j)_{j\in I}$ have finite range and are lower-semi-continuous. We will prove under these assumptions that there exists a subgame-perfect 0-equilibrium. In the proof, we use the finiteness of the range of the payoffs to have a maximal payoff and a minimal payoff in every non-empty subset of payoffs. The lower-semi-continuity of the payoff

[3] Below, $\lfloor x \rfloor$ is the largest integer that is *strictly* smaller than x, for every real number x.

functions, on the other hand, will be used to obtain the following property: when the players are supposed to play according to a strategy profile $\sigma = (\sigma^\ell)_{\ell \in I}$, if some player j cannot deviate profitably by not playing the action prescribed by σ^j finitely many times, then he cannot deviate profitably by disobeying σ^j infinitely many times.

3.1 Constructing some sequences

In this subsection we define for every finite history $h \in \mathcal{H}$ and every ordinal ξ, (a) a real number $\alpha_\xi(h)$, and (b) a set $P_\xi(h)$ of plays. The sequence $(\alpha_\xi(h))_\xi$ will be a non-decreasing sequence of lower bounds to the set of subgame-perfect 0-equilibrium payoffs for player $i(h)$ in the subgame that starts at h. The sequence $(P_\xi(h))_\xi$ will be a non-increasing (by inclusion) sequence of sets of plays; a play that is not in $P_\xi(h)$ cannot be induced by a subgame-perfect 0-equilibrium in the subgame that starts at h.

We will in fact prove a Folk theorem: $\max_\xi \alpha_\xi(h)$ will be the minimal subgame-perfect 0-equilibrium payoff of player $i(h)$ in the subgame that starts at h, and a play will be in all the sets $(P_\xi(h))_\xi$ if and only if it is induced by some subgame-perfect 0-equilibrium in the subgame that starts at h.

For every finite history $h \in \mathcal{H}$ set:

$$P_1(h) := \{p \in A^{\mathbf{N}} \colon h \prec p\}, \tag{1}$$

$$\alpha_1(h) := \min_{p \in P_1(h)} u^{i(h)}(p). \tag{2}$$

The set $P_1(h)$ consists of all plays that extend h, and the quantity $\alpha_1(h)$ is a naive lower bound to the set of subgame-perfect 0-equilibrium payoffs in the subgame that starts at h.

If $h = (a_t)_{t=1}^l$ is a finite history with length l, and $a \in A$, we denote by $(h, a) = (a_1, a_2, \ldots, a_l, a)$ the finite history of length $l+1$ that starts with h and ends with a.

For every successor ordinal $\xi + 1$ and every finite history $h \in \mathcal{H}$ define

$$\alpha_{\xi+1}(h) := \max_{a \in A} \min_{p \in P_\xi(h,a)} u^{i(h)}(p), \tag{3}$$

$$P_{\xi+1}(h) := \left\{ p \in \cup_{a \in A} P_\xi(h,a) \colon u^{i(h)}(p) \geq \alpha_{\xi+1}(h) \right\}. \tag{4}$$

As we will show, a play that is not in $P_\xi(h,a)$ cannot be induced by a subgame-perfect 0-equilibrium in the subgame that starts at (h,a). Therefore, when player $i(h)$ considers the subgame that starts at h, he can ignore the plays that are not in $\cup_{a \in A} P_\xi(h,a)$. In particular, when player $i(h)$ plays optimally at h, the quantity $\alpha_{\xi+1}(h)$ is a lower bound to his payoff in subgame-perfect 0-equilibria in the subgame that starts at h, and a play that is not in $P_{\xi+1}(h)$ cannot be induced by a subgame-perfect 0-equilibrium in this subgame.

For every limit ordinal ξ and every finite history $h \in \mathcal{H}$ define

$$P_\xi(h) := \cap_{\lambda < \xi} P_\lambda(h), \tag{5}$$

$$\alpha_\xi(h) := \min_{p \in P_\xi(h)} u^{i(h)}(p). \tag{6}$$

As we will show, a play that is not in $P_\lambda(h)$ for some $\lambda < \xi$ cannot be induced by a subgame-perfect 0-equilibrium in the subgame that starts at h. Therefore, the same holds for any play that is not in $P_\xi(h)$. Moreover, the quantity $\alpha_\xi(h)$ is then a lower bound to the payoff of player $i(h)$ in subgame-perfect 0-equilibria in the subgame that starts at h.

3.2 Properties of the sequences $(\alpha_\xi(h))_\xi$ and $(P_\xi(h))_\xi$

We first list a few simple properties of the sequences $(\alpha_\xi)_\xi$ and $(P_\xi)_\xi$ that easily follow from the definitions and that we will use later.

Lemma 10. *Let ξ be an ordinal and let h be a finite history.*

1. $\alpha_\xi(h) = \min_{p \in P_\xi(h)} u^{i(h)}(p)$. *In particular, there is a play $p \in P_\xi(h)$ such that $u^{i(h)}(p) = \alpha_\xi(h)$.*
2. *Let $a \in A$ be an action that achieves the maximum in the right-hand side of (3). If $p \in P_\xi(h, a)$ then $p \in P_{\xi+1}(h)$.*
3. *Suppose that $\xi = 1$ or ξ is a limit ordinal. Let $p \in P_\xi(h)$ and let h' be a finite history that satisfies $h \preceq h' \prec p$. Then $p \in P_\xi(h')$.*

Proof. Part 1 follows from the definitions (2) and (6) for $\xi = 1$ and for limit ordinals, and from the definitions (3) and (4) for successor ordinals.

Part 2 follows from the definitions (3) and (4).

We now prove Part 3. For $\xi = 1$, the claim follows from definition (1). So, assume that ξ is a limit ordinal. We will show that $p \in P_\xi(h, a)$, where $a \in A$ is the action in p right after h. Let λ be an ordinal such that $\lambda < \xi$. As $\lambda + 1 < \xi$, it follows from definition (5) that $p \in P_{\lambda+1}(h)$. Hence, by definition (4), we have $p \in P_\lambda(h, a)$. Since $\lambda < \xi$ was arbitrary, it follows that $p \in P_\xi(h, a)$. The proof for any finite history h' that satisfies $h \prec h' \prec p$ follows by induction. □

The following theorem states additional properties of the sequences $(\alpha_\xi(h))_\xi$ and $(P_\xi(h))_\xi$, which play a crucial role in the proof of Theorem 9.

Theorem 11. *The following holds for every $h \in \mathcal{H}$:*

1. *The set $P_\xi(h)$ is not empty for every ordinal ξ.*
2. *The sequence $(P_\xi(h))_\xi$ is monotonic non-increasing (by inclusion).*
3. *The sequence $(\alpha_\xi(h))_\xi$ is monotonic non-decreasing.*

Before proving the theorem we define another property of plays, ξ-monotonicity, that will be used in the proof of Theorem 11.

Definition 12. *Let h be a finite history, and p a play that extends h. The play p is called* ξ-monotonic *at h if the sequence $(P_\xi(h'))_{h \preceq h' \prec p}$ is non-increasing:*

$$P_\xi(h') \supseteq P_\xi(h''), \quad \forall h', h'' \text{ such that } h \preceq h' \preceq h'' \prec p.$$

Proof of Theorem 11. The proof is by transfinite induction, and it will follow once we prove the following claims:

Claim 1: $P_1(h) \neq \emptyset$ for every $h \in \mathcal{H}$. Moreover, every $p \in P_1(h)$ is 1-monotonic at h.

Claim 2: For every ordinal ξ, if $P_\xi(h) \neq \emptyset$ for every $h \in \mathcal{H}$, then $P_{\xi+1}(h) \neq \emptyset$ for every $h \in \mathcal{H}$.

Claim 3: For $\xi = 1$ and for every limit ordinal ξ we have $\alpha_{\xi+1}(h) \geq \alpha_\xi(h)$ and $P_{\xi+1}(h) \subseteq P_\xi(h)$ for every $h \in \mathcal{H}$.

Claim 4: If $\alpha_{\xi+1}(h) \geq \alpha_\xi(h)$ and $P_{\xi+1}(h) \subseteq P_\xi(h)$ for every $h \in \mathcal{H}$, then $\alpha_{\xi+2}(h) \geq \alpha_{\xi+1}(h)$ and $P_{\xi+2}(h) \subseteq P_{\xi+1}(h)$ for every $h \in \mathcal{H}$.

Claim 5: For every limit ordinal ξ, every ordinal $\lambda < \xi$, and every $h \in \mathcal{H}$, one has $\alpha_\xi(h) \geq \alpha_\lambda(h)$ and $P_\xi(h) \subseteq P_\lambda(h)$.

Claim 6: For every limit ordinal ξ and every $h \in \mathcal{H}$ one has $P_\xi(h) \neq \emptyset$. Moreover, there is a play $p \in P_\xi(h)$ that is ξ-monotonic at h.

As we will see, Claims 1-5 easily follow from the definitions. The proof of claim 6 is the challenging part of the whole proof. The existence of ξ-monotonic plays in Claims 1 and 6 is needed for the inductive step of the proof of Claims 3 and 6.
<u>Proof of Claim 1:</u> This claim follows from the definitions (1) and (2).
<u>Proof of Claim 2:</u> This claim follows from the definitions (3) and (4).
<u>Proof of Claim 3:</u> By the induction hypothesis (Claims 1 and 6) there is a play $\widehat{p} \in P_\xi(h)$ that is ξ-monotonic at h. Let a_0 be the first action in $\widehat{p}$ after h. Then $P_\xi(h) \supseteq P_\xi(h, a_0)$. Therefore, by (3) and (6),

$$\begin{aligned} \alpha_{\xi+1}(h) &= \max_{a \in A} \min_{p \in P_\xi(h,a)} u^{i(h)}(p) \geq \min_{p \in P_\xi(h,a_0)} u^{i(h)}(p) \\ &\geq \min_{p \in P_\xi(h)} u^{i(h)}(p) = \alpha_\xi(h), \end{aligned}$$

and the first part of the claim holds.

We now prove that $P_{\xi+1}(h) \subseteq P_\xi(h)$. For $\xi = 1$ this follows trivially from (1). We therefore assume that ξ is a limit ordinal. Let $p \in P_{\xi+1}(h)$. By (4) we have $p \in P_\xi(h, a_0)$, where a_0 is the first action in p after h. By (5), $p \in P_\lambda(h, a_0)$ for every ordinal $\lambda < \xi$. By the first part of Claim 3 and by Claim 5 applied to ξ,

$$u^{i(h)}(p) \geq \alpha_{\xi+1}(h) \geq \alpha_\xi(h) \geq \alpha_{\lambda+1}(h),$$

so that by (4), $p \in P_{\lambda+1}(h)$ for every $\lambda < \xi$. It follows by (5) that $p \in P_\xi(h)$, as desired.
<u>Proof of Claim 4:</u> This claim follows from the definitions (3) and (4).
<u>Proof of Claim 5:</u> The claim $P_\xi(h) \subseteq P_\lambda(h)$ follows from definition (5). Hence, part 1 of Lemma 10 implies $\alpha_\xi(h) \geq \alpha_\lambda(h)$.

Proof of Claim 6: Fix a limit ordinal ξ. We will prove that if Claims 1-6 hold for every ordinal λ smaller than ξ, and Claim 5 holds for ξ as well, then Claim 6 also holds for the ordinal ξ. The following lemma follows from Claims 3-5.

Lemma 13. *For every finite history h, the sequence $(\alpha_\lambda(h))_{\lambda\leq\xi}$ is monotonic non-decreasing, and the sequence $(P_\lambda(h))_{\lambda\leq\xi}$ is monotonic non-increasing (by inclusion).*

For every finite history h, we define

$$\widetilde{\alpha}_\xi(h) := \max_{\lambda<\xi} \alpha_\lambda(h). \tag{7}$$

By Lemma 13, $\widetilde{\alpha}_\xi(h) \leq \alpha_\xi(h)$.

Fix a finite history h. We are going to generate a play that extends h, and we will show that it is in $P_\xi(h)$ and that it is ξ-monotonic at h. The play will be generated in iterations; the output of the first iteration is an extension of h, and the output of each subsequent iteration extends the output of the previous iteration. The construction in odd iterations differs from the construction in even iterations. We will then prove that an infinite play is generated after an even number of iterations. Finally we will show that this play is in $P_\xi(h)$ and that it is ξ-monotonic at h.

Odd iterations:

Let h_1 be the finite history at the beginning of the odd iteration. For the first iteration, $h_1 = h$. For all other odd iterations, it is the finite history generated by the previous even iteration.

Consider the following algorithm that generates a finite history or a play that extends h_1.

1. Let $\xi_1 < \xi$ be a successor ordinal that satisfies $\widetilde{\alpha}_\xi(h_1) = \alpha_{\xi_1}(h_1)$. Such an ordinal exists because (a) the range of payoffs is finite, and (b) every nonempty set of ordinals has a minimal element.

2. Let a_1 be an action of player $i(h_1)$ that achieves the maximum in (3) for h_1 and ξ_1, that is,

$$\alpha_{\xi_1}(h_1) = \min_{p\in P_{\xi_1-1}(h_1,a_1)} u^{i(h_1)}(p).$$

   Set $h_2 = (h_1, a_1)$.

3. If $\xi_1 > 1$, let $\xi_2 \geq \xi_1 - 1$ be the minimal ordinal that satisfies $\widetilde{\alpha}_\xi(h_2) = \alpha_{\xi_2}(h_2)$. Note that because ξ is a limit ordinal, $\xi_2 < \xi$.

4. If ξ_2 is a successor ordinal, let a_2 be an action of player $i(h_2)$ that achieves the maximum in (3) for h_2 and ξ_2, that is,

$$\alpha_{\xi_2}(h_2) = \min_{p\in P_{\xi_2-1}(h_2,a_2)} u^{i(h_2)}(p).$$

   Set $h_3 = (h_1, a_1, a_2)$.

5. Continue this way to create a sequence $(h_1, \xi_1, a_1, h_2, \xi_2, a_2, \ldots)$. The iteration ends when either $\xi_m = 1$ or ξ_m is a limit ordinal; in this case, the output of the odd iteration is the finite history $h_m = (h_1, a_1, a_2, \ldots, a_{m-1})$. If $\xi_m > 1$ is a successor ordinal for every $m \in \mathbf{N}$, the iteration never ends.

Note that every ordinal ξ_k generated along an odd iteration satisfies $\xi_k < \xi$. As the next lemma states, odd iterations are finite.

Lemma 14. *There is $m \in \mathbf{N}$ such that either $\xi_m = 1$ or ξ_m is a limit ordinal.*

Proof. Assume that the algorithm never terminates: $\xi_m > 1$ is a successor ordinal for every $m \in \mathbf{N}$, so that the algorithm generates an infinite sequence $(h_1, \xi_1, a_1, h_2, \xi_2, a_2, \ldots)$.

We first argue that for every $m \in \mathbf{N}$ one has[4]

$$P_{\xi_m}(h_m) \supseteq P_{\xi_m - 1}(h_{m+1}) \supseteq P_{\xi_{m+1}}(h_{m+1}). \tag{8}$$

Indeed, the left-hand side inclusion holds by Lemma 10(2), whereas the right-hand side inclusion holds by Lemma 13 and since $\xi_m - 1 \le \xi_{m+1}$.

By (8), for every player j

$$\min_{p \in P_{\xi_m}(h_m)} u^j(p) \le \min_{p \in P_{\xi_m - 1}(h_{m+1})} u^j(p) \le \min_{p \in P_{\xi_{m+1}}(h_{m+1})} u^j(p). \tag{9}$$

Because the payoffs are discrete, the inequalities in (9) can be strict only finitely many times, for every player j. That is, there is $M \in \mathbf{N}$ sufficiently large such that for every player $j \in I$ and every $m \ge M$,

$$\min_{p \in P_{\xi_m}(h_m)} u^j(p) = \min_{p \in P_{\xi_m - 1}(h_{m+1})} u^j(p) = \min_{p \in P_{\xi_{m+1}}(h_{m+1})} u^j(p). \tag{10}$$

Let m, m' be two integers satisfying (a) $M \le m < m'$, and (b) $i(h_m) = i(h_{m'})$. By repeated use of Eq. (10),

$$\begin{aligned} \widetilde{\alpha}_\xi(h_m) &= \alpha_{\xi_m}(h_m) = \min_{p \in P_{\xi_m}(h_m)} u^{i(h_m)}(p) = \min_{p \in P_{\xi_{m'-1}-1}(h_{m'})} u^{i(h_m)}(p) \\ &= \min_{p \in P_{\xi_{m'}}(h_{m'})} u^{i(h_m)}(p) = \alpha_{\xi_{m'}}(h_{m'}) = \widetilde{\alpha}_\xi(h_{m'}). \end{aligned} \tag{11}$$

Because

$$\alpha_{\xi_{m'-1}-1}(h_{m'}) = \min_{p \in P_{\xi_{m'-1}-1}(h_{m'})} u^{i(h_m)}(p) = \alpha_{\xi_{m'}}(h_{m'}),$$

it follows that $\xi_{m'} = \xi_{m'-1} - 1$. Because this equality holds for every m' sufficiently large, there is m such that either $\xi_m = 1$ or ξ_m is a limit ordinal, as desired. □

[4]Because ξ_m is a successor ordinal, the ordinal $\xi_m - 1$ is well defined.

Observe that inclusion (8) holds as long as ξ_m is a successor ordinal, and therefore it holds all along the odd iteration.

Even iterations:

Let h_1 be the finite history that is the output of the previous odd iteration, and denote by λ the last ordinal ξ_m generated by the previous odd iteration. In particular, either $\lambda = 1$ or λ is a limit ordinal, and $\lambda < \xi$. Moreover, $\widetilde{\alpha}_\xi(h_1) = \alpha_\lambda(h_1)$.

By the induction hypotheses of Claim 1 (if $\lambda = 1$) or Claim 6 applied to λ (if $1 < \lambda < \xi$), there is a play $p \in P_\lambda(h_1)$ that is λ-monotonic at h_1.

By the definition of $\widetilde{\alpha}_\xi(h')$, we have $\widetilde{\alpha}_\xi(h') \geq \alpha_\lambda(h')$ for every prefix h' of p that extends h_1. If $\widetilde{\alpha}_\xi(h') = \alpha_\lambda(h')$ for every prefix h' of p that extends h_1, the even iteration is infinite and its output is p. Otherwise, the output of the even iteration is the shortest prefix h' of p that extends h_1 for which $\widetilde{\alpha}_\xi(h') > \alpha_\lambda(h')$, and in this case we proceed with the next odd iteration.

Denote by p_* the play that extends h, which is generated by (a possibly infinite) use of odd and even iterations. We will now show that p_* is ξ-monotonic at h and that it is in $P_\xi(h)$.

Denote by $(h_m)_{m\in\mathbf{N}}$ all finite prefixes of p_* that extend h, so that $h_1 = h$. Denote by ξ_m the ordinal that is attached to h_m in the construction of p_*; it is a successor ordinal along odd iterations, and a limit ordinal or 1 along even iterations.

Lemma 15. *The play p_* is ξ-monotonic at h.*

Proof. Let $m \geq 1$, and let $p \in P_\xi(h_{m+1})$. We will prove that $p \in P_\xi(h_m)$. By (5) it follows that $p \in P_\tau(h_{m+1})$, for every $\tau < \xi$, and in particular $p \in P_{\xi_{m+1}}(h_{m+1})$. We deduce that $p \in P_{\xi_m}(h_m)$; if h_m is in an odd iteration (and it is not the last history of that odd iteration) this follows from (8), while if h_{m+1} is in an even iteration (and it is not the last history of that even iteration) this follows because then $\xi_m = \xi_{m+1}$ (both are equal to the ordinal λ of that even iteration) and the part of the play added in an even iteration is ξ_m-monotonic.

Because $p \in P_{\xi_m}(h_m)$ we have

$$u^{i(h_m)}(p) \geq \alpha_{\xi_m}(h_m) = \widetilde{\alpha}_\xi(h_m) \geq \alpha_{\tau+1}(h_m),$$

for every ordinal $\tau < \xi$. By (4) this implies that $p \in P_{\tau+1}(h_m)$ for every $\tau < \xi$, so that by (5) we have $p \in P_\xi(h_m)$. ☐

We are now ready to prove Claim 6.

Lemma 16. $p_* \in P_\xi(h)$.

Proof. Suppose first that the number of iterations is finite, so that the last even iteration is infinite. Denote by h_m the history at the beginning of the last even block. Then $\xi_m = \xi_{m+1} = \cdots =: \lambda$. We will show that $p_* \in P_\xi(h_m)$, so that by the ξ-monotonicity of p_* (Lemma 15) it will follow that $p_* \in P_\xi(h)$, as desired. Note that by the definition of even blocks, $\widetilde{\alpha}_\xi(h_{m'}) = \alpha_\lambda(h_{m'})$ for every $m' \geq m$.

Assume to the contrary that $p_* \notin P_\xi(h_m)$. Let τ be the smallest ordinal such that $p_* \notin P_\tau(h_{m'})$ for some $m' \geq m$. Note that $\tau > \lambda$: because $p_* \in P_\lambda(h_m)$, by Lemma 10(3) we have $p_* \in P_\lambda(h_{m'})$ for every $m' \geq m$. By definition (5), τ cannot be a limit ordinal, so that τ is a successor ordinal. It follows that $p_* \in P_{\tau-1}(h_{m'})$ for every $m' \geq m$. To derive a contradiction we argue that $p_* \in P_\tau(h_{m'})$ for every $m' \geq m$. Indeed, for every $m' \geq m$, because $p_* \in P_\lambda(h_{m'})$, $\widetilde{\alpha}_\xi(h_{m'}) = \alpha_\lambda(h_{m'})$, and $\xi > \tau$, it follows that

$$u^{i(h_{m'})}(p_*) \geq \alpha_\lambda(h_{m'}) = \widetilde{\alpha}_\xi(h_{m'}) \geq \alpha_\tau(h_{m'}),$$

so that by definition (4) we have $p_* \in P_\tau(h_{m'})$, as claimed.

We now show that the number of iterations cannot be infinite. Note that $P_{\xi_m}(h_m) \supseteq P_{\xi_{m+1}}(h_{m+1})$ for every $m \in \mathbf{N}$. Indeed, if h_m is in an odd iteration then the inclusion follows from (8); if h_m and h_{m+1} are both in the same even iteration then this inclusion holds because the part of the play added in this even iteration is ξ_m-monotonic and $\xi_m = \xi_{m+1}$; in case h_m is the last history of an even iteration and h_{m+1} is the first history of the subsequent odd iteration, then the construction implies that $\xi_{m+1} > \xi_m$, so that the inclusion follows by $P_{\xi_m}(h_m) \supseteq P_{\xi_m}(h_{m+1})$ and Lemma 13.

It follows that for every $m \in \mathbf{N}$ and every player j,

$$\min_{p \in P_{\xi_m}(h_m)} u^j(p) \leq \min_{p \in P_{\xi_{m+1}}(h_{m+1})} u^j(p). \tag{12}$$

Because the range of the payoffs is finite, the inequality (12) can be strict only finitely many times, for every player j.

Let h_{m+1} be the last finite history of an even iteration (so it is also the first history of the next odd iteration). Then the finite history h_m is generated by the same even iteration as h_{m+1}, and

$$\min_{p \in P_{\xi_m}(h_m)} u^{i(h_{m+1})}(p) \leq \min_{p \in P_{\xi_m}(h_{m+1})} u^{i(h_{m+1})}(p) < \min_{p \in P_{\xi_{m+1}}(h_{m+1})} u^{i(h_{m+1})}(p) \tag{13}$$

where the weak inequality holds because the part of p_* generated by the even iteration is ξ_m-monotonic, and the strict inequality holds because h_{m+1} is the last finite history of an even iteration. In particular, by Eq. (12), since the range of the payoffs is finite and since there are finitely many players, Eq. (13) can hold only finitely many times, so there can be only finitely many even iterations. □

The proof of Theorem 11 is now complete. □

The next theorem states that the process of defining the sequences $(\alpha_\xi(h))_\xi$ and $(P_\xi(h))_\xi$ reaches a fixed point.

Theorem 17. *There is an ordinal ξ_* such that for every finite history $h \in \mathcal{H}$ we have $\alpha_{\xi_*}(h) = \alpha_{\xi_*+1}(h)$, and $P_{\xi_*}(h) = P_{\xi_*+1}(h) \neq \emptyset$.*

In section 4.1 we provide an example which shows that ξ_* can be any ordinal.

Proof. Denote by ρ the cardinality of the set of functions that assign to each finite history $h \in \mathcal{H}$ a subset of $A^{\mathbf{N}}$. For every finite history $h \in \mathcal{H}$, the sequence $(P_\xi(h))_\xi$ is non-increasing. Moreover, if $P_\xi(h) = P_{\xi+1}(h)$ for every $h \in \mathcal{H}$ then $P_{\xi+1}(h) = P_{\xi+2}(h)$ for every $h \in \mathcal{H}$, which implies that $P_\xi(h) = P_\tau(h)$ for every $h \in \mathcal{H}$ and every ordinal $\tau > \xi$. It follows that for every ordinal ξ whose cardinality is larger than ρ, $P_\xi(h) = P_{\xi+1}(h)$ for every $h \in \mathcal{H}$.

By Lemma 10(1) it follows that for each such ordinal ξ, $\alpha_\xi(h) = \alpha_{\xi+1}(h)$ for every $h \in \mathcal{H}$, and the result follows. □

3.3 Proof of Theorem 9

We now construct a strategy profile $\sigma_* = (\sigma_*^j)_{j \in I}$, and show that it is a subgame-perfect 0-equilibrium.

For the initial history $\emptyset$ choose an arbitrary play $p(\emptyset) \in P_{\xi_*}(\emptyset)$. For every other finite history $h = (a_l)_{l<t}$, denote by $h^- = (a_l)_{l<t-1}$ the prefix of h excluding the last action, and by $i(h^-)$ the active player at h^-. Choose a play $p(h) \in P_{\xi_*}(h)$ that extends h and that satisfies

$$u^{i(h^-)}(p(h)) = \min_{p \in P_{\xi_*}(h)} u^{i(h^-)}(p). \tag{14}$$

If player j deviates, and h is the finite history right after the deviation (so that $j = i(h^-)$), then $p(h)$ is a play at h that minimizes player's j's payoff within $P_{\xi_*}(h)$.

Let σ_*^j be the following strategy: Follow the play $p(\emptyset)$ as long as all other players follow $p(\emptyset)$. Suppose that at stage t_1 player j_1 deviates from $p(\emptyset)$. From stage $t_1 + 1$ and on follow the play $p(h_{t_1+1})$ as long as all other players follow $p(h_{t_1+1})$. Suppose that at stage $t_2 > t_1$ player j_2 deviates from $p(h_{t_1+1})$. From stage $t_2 + 1$ and on follow the play $p(h_{t_2+1})$ as long as all other players follow $p(h_{t_2+1})$. Continue this way.

We now show that σ_* is a subgame-perfect 0-equilibrium. To this end, we fix a finite history $h \in \mathcal{H}$, and we show for an arbitrary player j that

$$u^j(\sigma_*^{-j}, \sigma^j \mid h) \le u^j(\sigma_* \mid h), \quad \forall \sigma^j \in \Sigma^j.$$

Let $\sigma^j \in \Sigma^j$ be any strategy of player j. Let $p_* = p(\sigma_* \mid h)$ be the play induced by σ_* given h. This is the play that is generated given h if player j does not deviate. Let $p = p(\sigma_*^{-j}, \sigma^j \mid h)$ be the play given h when player j deviates to σ^j.

Denote by $t_1, t_2, \ldots$ the stages where σ^j and σ_*^j differ along p; in those stages all the players observe the deviations of player j. The sequence $(t_k)_k$ may be finite or infinite. Denote by $p_k = p(h_{t_k+1})$ the play that the players start to follow from stage $t_k + 1$ on, for each k.

We complete the proof by showing that

$$u^j(p) \le u^j(p_*). \tag{15}$$

It is sufficient to show that

$$u^j(p_k) \leq u^j(p_*), \quad \forall k. \tag{16}$$

Indeed, if σ^j and σ^j_* differ only finitely many times along p, Eq. (15) follows from Eq. (16) applied to the last k; if σ^j and σ^j_* differ infinitely many times along p, then the sequence $(p_k)_{k\in\mathbf{N}}$ converges to p, so that Eq. (15) follows from Eq. (16) and the lower-semi-continuity of u^j. This is the only place in the proof where the lower-semi-continuity of the payoff functions is used.

The proof of (16) is by induction on k. Due to the construction and to Lemma 10(3), we have $p_* \in P_{\xi_*}(h)$. Hence, Lemma 10(3) and Eq. (14) imply $u^j(p_1) \leq u^j(p_*)$. For every $k \geq 1$, because $p_k \in P_{\xi_*}(h_{t_k})$, and by Lemma 10(3) and Eq. (14), $u^j(p_{k+1}) \leq u^j(p_k)$, which is at most $u^j(p_*)$ by the induction hypothesis. The proof is now complete.

3.4 A Folk Theorem

Our construction enables us to characterize the set of plays that can arise in a subgame-perfect 0-equilibrium in the game with discrete payoffs.

Theorem 18. *A play p is induced by some subgame-perfect 0-equilibrium if and only if $p \in P_{\xi_*}(\emptyset)$.*

It follows from this result that for every $h \in \mathcal{H}$, $\alpha_{\xi_*}(h)$ is the lowest subgame-perfect 0-equilibrium payoff in the subgame that starts at h, and that p is the play induced by some subgame-perfect 0-equilibrium if and only if $u^{i(h)}(p) \geq \alpha_{\xi_*}(h)$ for every prefix h of p.

Proof. If p is in $P_{\xi_*}(\emptyset)$, then the construction in Section 3.3 shows that it is the play that is induced by some subgame-perfect 0-equilibrium.

To see that the converse is true, we show that if σ_* is a subgame-perfect 0-equilibrium, then $p(\sigma_* \mid h)$ is in $P_\xi(h)$, for every ordinal ξ and every finite history h. The proof is by transfinite induction on ξ.

Because every play that extends h is in $P_1(h)$, the claim holds for $\xi = 1$.

Suppose now that the claim holds for a given ordinal ξ. Let h be any finite history. Because the claim holds for ξ, the play $p(\sigma_* \mid (h,a))$ is in $P_\xi(h,a)$ for every action a. Therefore, with respect to σ_*, the payoff to player $i(h)$ is at least $\alpha_{\xi+1}(h)$ in the subgame that starts at h. This implies that $p(\sigma_* \mid h)$ is in $P_{\xi+1}(h)$, for every h.

Finally, the definition of $P_\xi(h)$ for limit ordinals ξ implies that if $p(\sigma_* \mid h)$ is in $P_\lambda(h)$ for every ordinal $\lambda < \xi$, then it is also in $P_\xi(h)$. □

4 Concluding Remarks

4.1 How large can ξ_* be?

A natural question is whether the ordinal ξ_* of Theorem 17 is at most the first infinite ordinal ω, or whether it can be larger than ω. As the following example shows, ξ_* can be any ordinal. For simplicity of exposition, in this example the set of actions is history dependent.

Let τ be any ordinal. Roughly speaking, we consider a game in which two players I and II choose a non-increasing sequence $(a_t)_t$ of ordinals, with $a_0 = \tau$, according to the following rules. If the current ordinal a_{t-1} is a successor ordinal, then player I chooses a_t from $\{a_{t-1}, a_{t-1} - 1\}$. If the current ordinal a_{t-1} is a limit ordinal, then player II chooses any ordinal a_t smaller than a_{t-1}. Finally, if $a_{t-1} = 1$ then all further choices $(a_m)_{m \geq t}$ will be 1 as well. The payoff for player I equals 1 if the sequence $(a_t)_t$ eventually reaches 1 and equals 0 otherwise. The payoff for player II equals 0 for every play.

The idea of this game is that player I can make sure that the sequence $(a_t)_t$ eventually reaches 1 and thereby obtain the best payoff 1, but the number of stages needed to reach 1 depends on player II. More precisely, if player I, whenever he is the active player, always lowers the current ordinal by 1, then $a_t < a_{t-1}$ holds as long as $a_{t-1} > 1$. Since there is no infinite sequence of decreasing ordinals, this strategy of player I guarantees that $a_t = 1$ for some $t \in \mathbf{N}$, regardless the actions that player II chooses. Still, if λ_1 and λ_2 are two limit ordinals such that $\lambda_2 < \lambda_1 \leq \tau$, there is no bound on the number of stages needed to descend from λ_1 to λ_2, as player II can choose the ordinal $\lambda_2 + k$ for any $k \in \mathbf{N}$ when the current ordinal is λ_1. As we will show, one needs τ steps in our iterative method to realize that the sequence $(a_t)_t$ eventually reaches 1 whatever ordinals player II chooses.

Formally we consider the following two-player Borel game. For any history[5] $h = (a_0, a_1, \ldots, a_{t-1})$, the active player $i(h)$ and his action set $A(h)$ are defined as follows:

- If a_{t-1} is a successor ordinal: $i(h) = \text{I}$ and $A(h) = \{a_{t-1}, a_{t-1} - 1\}$.
- If a_{t-1} is a limit ordinal: $i(h) = \text{II}$ and $A(h) = \{$all ordinals smaller than $a_{t-1}\}$.
- If $a_{t-1} = 1$: $i(h) = \text{I}$ and $A(h) = \{1\}$.[6]

Let W denote the set of all plays $p = (a_t)_{t \geq 0}$ such that $a_t = 1$ for some t. For an arbitrary play p, the payoff to player I is as follows: $u^{\text{I}}(p) = 1$ if $p \in W$, and $u^{\text{I}}(p) = 0$ otherwise. The payoff to player II is $u^{\text{II}}(p) = 0$ for every play p. Because there is no infinite strictly decreasing sequence of ordinals, the payoff functions are lower-semi-continuous.

We claim that for every finite history $h = (a_0, a_1, \ldots, a_{t-1})$:

(a) If $\xi < a_{t-1}$ and a_{t-1} is a successor ordinal, then $(h, a_{t-1}, a_{t-1}, \ldots) \in P_\xi(h) \setminus W$. If $\xi < a_{t-1}$ and a_{t-1} is a limit ordinal, then $(h, \rho, \rho, \ldots) \in P_\xi(h) \setminus W$ for every successor ordinal ρ satisfying $\xi + 1 \leq \rho < a_{t-1}$.

(b) If $\xi \geq a_{t-1}$ then $P_\xi(h) \subseteq W$.

[5] To simplify notations, we denote the initial history by $h_1 = (a_0)$.

[6] In this case, it makes no difference which player is the active player.

In particular, this will imply that $\xi_* = a_0 = \tau$. The proof of the claim is by transfinite induction on ξ.

For $\xi = 1$, the claim is obvious.

Assume that the claim holds for some ordinal ξ. We will now prove the claim for $\xi + 1$.

Suppose that $\xi + 1 < a_{t-1}$. If a_{t-1} is a successor ordinal, then whichever action $a_t \in \{a_{t-1}, a_{t-1} - 1\}$ player I chooses, we have $\xi < a_t$, and therefore the induction hypothesis implies that $P_\xi(h, a_t) \setminus W$ is non-empty. Hence, $\alpha_{\xi+1}(h) = 0$ and $(h, a_{t-1}, a_{t-1}, \ldots) \in P_{\xi+1}(h) \setminus W$. If a_{t-1} is a limit ordinal, since player II can choose any successor ordinal ρ satisfying $\xi + 1 \leq \rho < a_{t-1}$, we obtain $(h, \rho, \rho, \ldots) \in P_{\xi+1}(h) \setminus W$.

Suppose that $\xi + 1 \geq a_{t-1}$. If $\xi \geq a_{t-1}$ then, by the induction hypothesis and because the sequence $(P_\rho(h))_\rho$ is monotonic non-increasing (by inclusion), we obtain $P_{\xi+1}(h) \subseteq P_\xi(h) \subseteq W$. Assume then that $\xi + 1 = a_{t-1}$. Since player I can choose the action $a_t = \xi$, and since $P_\xi(h, \xi) \subseteq W$ by the induction hypothesis, it follows that $P_{\xi+1}(h) \subseteq W$.

Finally, let ξ be a limit ordinal, and assume that the claim holds for all ordinals $\lambda < \xi$. If either $\xi < a_{t-1}$ or $\xi > a_{t-1}$ then the respective parts of the claim for ξ follow by (5). Suppose then that $\xi = a_{t-1}$ and take any play $p \in P_\xi(h)$. We will show that $p \in W$. Let a_t denote the action in p right after h. Since $p \in P_\xi(h)$ and $a_t < \xi$, we have $p \in P_{a_t+1}(h)$, and hence $p \in P_{a_t}(h, a_t)$. By the induction hypothesis, $p \in W$ as desired.

4.2 Other applications of the technique

The driving force behind the proof is the following property, that holds in games with perfect information. Denote by h the current finite history. Suppose that for every possible action a, $v(h, a)$ is the minimal continuation payoff possible for the decision maker at h if he chooses a, and suppose that if the decision maker chooses the action a_0, he is supposed to get a payoff x which is at least $\max_{a \in A} v(h, a)$. Then even if the decision maker at h will eventually receive a payoff higher than x after playing a_0 at h, one can construct a strategy profile that ensures that he plays a_0, and is punished by $v(h, a)$ otherwise.

This property does not hold, e.g., for mixed equilibria in sequential games with simultaneous moves, because in such games, if the continuation payoffs change, then the set of mixed equilibria may change as well, and a deviation from the original mixed equilibrium may not be detected.

The property does hold for extensive-form correlated equilibria in games with simultaneous moves. In this type of equilibrium, a mediator sends a private signal to each player at every stage. If the signal contains a recommended action for the current stage, as well as the recommendations made to all players in the previous stage, then a deviation from the recommendation is detected immediately and can be punished. We hope that our approach can be used to prove the existence of an extensive-form correlated equilibrium in multi-player Borel games with simultaneous moves.

4.3 Tightness of the result

It is well known that a 0-equilibrium, and therefore also a subgame-perfect 0-equilibrium, may fail to exist when the range of the payoff functions is not finite.

As the following example shows, when there are infinitely many players, a subgame-perfect 0-equilibrium need not exist even when the range of the payoff functions is finite. Suppose that the set of players is the set $\mathbf{N}$ of natural numbers, and the set of actions is $A = \{a, b\}$. Each player $t \in \mathbf{N}$ plays only once, at stage t. The payoff of player t is 1 if he played b, 2 if he played a and some player $j > t$ played b, and 0 if he played a and every player $j > t$ also played a. We claim that there is no subgame-perfect 0-equilibrium in this game. Suppose to the contrary that σ is a subgame-perfect 0-equilibrium. Since every player can guarantee 1 by simply playing action b, it cannot happen in any subgame that σ prescribes for all further players to play action a. This means in particular that, with respect to σ, infinitely many players play action b, and receive 1. But then each of those players is better off by deviating to a and receiving 2.

4.4 Chance moves

Borel games are deterministic, and the sequence of actions chosen by the players uniquely determines the outcome. In many situations there are chance moves along the game, where actions are chosen according to a known probability distribution. This situation is equivalent to the case where there is an additional player who follows a specific non-deterministic strategy, whatever the other players play. There are indications that our proof can be adapted to this more general situation, and this will be done elsewhere.

4.5 Positive recursive Borel games

Recursive Borel games are games where some finite histories are terminating, in the sense that once they occur the payoff is determined (and the play that follows them does not affect the players' payoffs), and the payoff of every infinite (non-terminating) play is 0. Various positional games that are studied in the computer science literature have this form. The significance of this class of games to game theory was exhibited in the context of stochastic games by Vieille (2000a,b), who used it as a step to proving the existence of an equilibrium payoff in every two-player stochastic game. A recursive Borel game is called positive if the terminal payoffs are positive for both players.

Flesch et al. (2010) studied positive recursive Borel game with finitely many states; these are positional games that are played on a finite directed graph, where each vertex is controlled by some player, and when the game reaches some vertex, the controlling player can choose whether to terminate the game, or whether to continue the game by choosing one of the edges that leaves the vertex. The terminal payoff, which is positive for all players, depends only on the vertex where termination occurred, and not on the whole past play.

Flesch et al. (2010) prove that every such game admits a subgame-perfect 0-equilibrium.[7] In their proof, they define for every vertex s a sequence $(\alpha_k(s))_{k\in\mathbf{N}}$ that is similar to our sequence $(\alpha_\xi(h))_\xi$, they prove that this sequence is non-decreasing, and, because there are finitely many vertices, they deduce that there is $k_* \in \mathbf{N}$ such that $\alpha_{k_*+1}(s) = \alpha_{k_*}(s)$ for every vertex s. They then use a similar construction of the subgame-perfect 0-equilibrium as the one that we used.

In Borel games every history is a different vertex. Therefore one needs to employ a much more delicate construction, that differs from the one in Flesch et al. (2010) in two respects. First, when the number of vertices is infinite, there need not be $k_* \in \mathbf{N}$ such that $\alpha_{k_*+1}(s) = \alpha_{k_*}(s)$ for every vertex s, and therefore $(\alpha_\xi(h))_\xi$ should be defined for every ordinal. Second, since play never terminates, one has to deal with plays of infinite length and introduce the sets $(P_\xi(h))_\xi$.

It turns out that for positive recursive Borel games our construction can be simplified, and a single odd iteration is sufficient to show that $P_\xi(h)$ is not empty for limit ordinals ξ.

4.6 Borel games with general payoffs

Example 3 in Solan and Vieille (2003) shows that without the condition that payoffs are lower-semi-continuous, a subgame-perfect ε-equilibrium need not exist. However, Solan and Vieille (2003) show that a subgame-perfect ε-equilibrium does exist if one allows behavior strategies. The existence of a subgame-perfect ε-equilibrium in behavior strategies was proved in other setups where the payoff functions are not lower-semi-continuous, see, e.g., Solan (2005) and Mashiah-Yaakovi (2009).

In our proof, the lower-semi-continuity of the payoff functions was used only in the last part, to show that any deviation σ^j that differs from σ^j_* infinitely many times cannot be profitable, as soon as any deviation σ^j that differs from σ^j_* finitely many times is not profitable. We do not know how the proof should be adapted to handle general payoff functions.

In fact, the following example shows that our definition of α_ξ and P_ξ is not appropriate for general Borel games. Consider a two-player Borel game with $A = \{a, b\}$. The payoff functions of the two players are as follows:

Condition	$u^1(h)$	$u^2(h)$
Both players played b finitely many times	2	2
Only player 1 played b finitely many times	2	1
Only player 2 played b finitely many times	1	2
No player played b finitely many times	0	0

[7] When transitions are random, Flesch et al. (2010) prove the existence of a subgame-perfect ε-equilibrium, for every $\varepsilon > 0$.

Note that u^1 and u^2 are not lower-semi-continuous. Playing b finitely many times is a dominant strategy for both players, so that the unique subgame-perfect 0-equilibrium payoff is $(2, 2)$. However, one can verify that for every finite history h and every ordinal ξ, $P_\xi(h)$ contains all plays in which at least one player plays b finitely many times, so that the Folk Theorem, Theorem 18, does not hold, and our construction of the subgame-perfect 0-equilibrium in the proof of Theorem 9 is invalid.

References

[1] Emerson E.A., Jutla C.S. and Sistla A.P. (2001) On Model Checking for the μ-Calculus and its Fragments. *Theoretical Computer Science*, **258**, 491-522.

[2] Flesch J., Kuipers J., Schoenmakers G. and Vrieze K. (2010) Subgame Perfection in Positive Recursive Games with Perfect Information, *Mathematics of Operations Research*, DOI:10.1287/moor.1090.0437.

[3] Gale D. and Stewart F.M. (1953) Infinite Games with Perfect Information. *Contributions to the Theory of Games*, Volume II, *Annals of Mathematical Studies*, **28**, 254-266.

[4] Grädel E. (2004) Positional Determinacy of Infinite Games. STACS 2004, *Lecture Notes in Computer Science*, **2996**, 4-18.

[5] Kechris A.S. (1995) Classical Descriptive Set Theory. Graduate Texts in Mathematics, **156**. Springer.

[6] Kuipers J., Flesch J., Schoenmakers G. and Vrieze K. (2008) Pure Subgame-Perfect Equilibria in Free Transition Games. *European Journal of Operational Research*, **199**, 442-447.

[7] Maitra A. and Sudderth W. (1993) Borel Stochastic Games with Limsup Payoffs. *The Annals of Probability*, **21**, 861-885.

[8] Martin D.A. (1975) Borel Determinacy. *Annals of Mathematics*, **102**, 363-371.

[9] Mashiah-Yaakovi A. (2009) Subgame-Perfect Equilibrium in Stopping Games with Perfect Information. Preprint.

[10] Mertens J.-F. (1987) Repeated Games. *Proceedings of the International Congress of Mathematicians*, **1**, (Berkeley, California, 1986), 1528-1577, American Mathical Society, Providence, RI.

[11] Schilling K. and Vaught R. (1983) Borel Games and the Baire Property. *Transactions of the American Mathematical Society*, **279**, 411-428.

[12] Selten R. (1965) Spieltheoretische Behandlung eines Oligopolm-odells mit Nachfrageträgheit. *Zeitschrift für die gesamte Staat-wissenschaft*, **121**, 301-324 and 667-689.

[13] Selten R. (1973) A Simple Model of Imperfect Competition, where Four are Few and Six are Many. *International Journal of Game Theorey*, **2**, 141-201.

[14] Shmaya E. (2008) Many inspections are manipulable. *Theoretical Economics*, **3**, 367-382.

[15] Solan E. (2005) Subgame-Perfection in Quitting Games with Perfect Information. *Mathematics of Operations Research*, **30**, 51-72.

[16] Solan E. and Nicolas N. (2003) Deterministic Multi-Player Dynkin Games. *Journal of Mathematical Economics*, **39**, 911-929.

[17] Thomas W. (2002) Infinite Games and Verification. In *Proceedings of the International Conference on Computer Aided Verification*, CAV'02, volume 2404, *Lecture Notes in Computer Science*, 58-64.

[18] Ummels M. (2005) Rational Behaviour and Strategy Construction in Infinite Multiplayer Games. Thesis, Rheinisch-Westf alischen Technischen Hochschule Aachen.

[19] Vieille N. (2000a) Equilibrium in 2-person stochastic games I: A reduction. *Israel Journal of Mathematics*, **119**, 55-91.

[20] Vieille N. (2000b) Equilibrium in 2-person stochastic games II: The case of recursive games. *Israel Journal of Mathematics*, **119**, 93-126.