Borel Games with Lower-Semi-Continuous Payoffs*

Janos Flesch[†], Jeroen Kuipers[‡], Ayala Mashiah-Yaakovi[§], Gijs Schoenmakers[†], Eilon Solan[¶] and Koos Vrieze[†]

December 22, 2009

Abstract

We prove that every two-player non-zero-sum Borel game with lower-semi-continuous payoffs admits a subgame-perfect ε-equilibrium. This result complements Example 3 in Solan and Vieille (2003), which shows that a subgame-perfect ε-equilibrium need not exists when the payoffs are not lower-semi-continuous.

1 Introduction

Borel games are sequential games where two players alternately choose actions. The payoff of each player is a function of the infinite sequence of actions that the players chose. Borel games were introduced by Gale and Stuart (1953), who studied 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 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 multi-player non-zero-sum Borel games to prove that for every $\varepsilon > 0$, every multi-player non-zero-sum Borel game has an ε-equilibrium, provided the payoff

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

[†]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

[§]School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel. e-mail: ayalam@post.tau.ac.il.

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

functions are bounded and measurable. 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 may lower 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 when the payoff function of each player is the indicator of some set. His proof is based on the following recursive construction. First, one identifies all finite histories which are a winning position to either player 1 or player 2; that is, if this finite history occurs, one of the players can ensure that his payoff is 1. After each such finite history one instruct the winning player to play his winning strategy, and the other player is instructed to play his best response. One then identifies winning positions to the two players in the new game, assuming the behavior in the first set of winning positions is set. The process repeats itself, until it reaches a stable state. Ummels (2005) proves that the pair of strategies thus defined is a subgame-perfect equilibrium.

In the present paper we show that every Borel game with bounded and lower-semi-continuous payoffs admits a subgame-perfect ε-equilibrium, for every $\varepsilon > 0$. This result complements Example 3 in Solan and Vieille (2003) that shows that in games with general payoff functions a subgame-perfect ε-equilibrium need not exist.

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 where 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) or Flesch et al. (2008, 2009). In particular, our result generalizes results in Flesch et al. (2008, 2009).

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 contains some comments.

2 The Model and the Main Result

Definition 1. *A (two-player non-zero-sum)* Borel game *is a triplet* (A, u^1, u^2) *where* A *is a set of actions, and* $u^1, u^2 : A^{\mathbf{N}} \to \mathbf{R}$ *are payoff functions.*

The game is played as follows. At every odd stage t player 1 chooses an action $a_t \in A$, and at every even stage t player 2 chooses an action $a_t \in A$. While making his choice at stage t, the player knows the sequence $(a_j)_{j<t}$ of actions that was chosen by the players in previous stages. The payoff to each player i is $u^i(a_1, a_2, \ldots)$. The description of the game is common knowledge among the players.

Comment 2. *The assumptions that (1) there are only two players, (2) both players have the same action set, and (3) the set of action is the same in all stages, are made for simplicity of notations only. Nothing that is said below would be affected if there are more than two players, if the players have different action sets, or if the sets of actions depend on past choices of actions.*

Denote by $\emptyset$ the empty history at the beginning of the game. The set of possible histories[1] at stage t is $H_t = A^{t-1}$. Denote by $H^1 = \bigcup_{k\in\mathbf{N}} H_{2k-1}$ and $H^2 = \bigcup_{k\in\mathbf{N}} H_{2k}$ the sets of possible histories at decision points of player 1 and player 2 respectively.

Definition 3. *A* strategy *for player* i *is a function* $\sigma^i : H^i \to A$.

We denote by Σ^i the strategy space of player i. Every pair of strategies $(\sigma^1, \sigma^2) \in \Sigma^1 \times \Sigma^2$ determines a unique play $p = p(\sigma^1, \sigma^2) = (a_t)_{t\in\mathbf{N}} \in A^{\mathbf{N}}$ as follows:

$$
\begin{aligned}
a_1 &= \sigma^1(\emptyset), && (1)\\
a_{2k} &= \sigma^2(a_1, a_2, \ldots, a_{2k-1}), \quad \forall k \in \mathbf{N}, && (2)\\
a_{2k+1} &= \sigma^1(a_1, a_2, \ldots, a_{2k}), \quad \forall k \in \mathbf{N}. && (3)
\end{aligned}
$$

We denote by $u^i(\sigma^1, \sigma^2) = u^i(p(\sigma^1, \sigma^2))$ the payoff of player i when the two players follow (σ^1, σ^2).

Definition 4. *Let* $\varepsilon \geq 0$. *A pair of strategies* (σ^1_*, σ^2_*) *is an* ε-equilibrium *if*

$$
\begin{aligned}
u^1(\sigma^1_*, \sigma^2_*) &\geq u^1(\sigma^1, \sigma^2_*) - \varepsilon, \quad \forall \sigma^1 \in \Sigma^1, && (4)\\
u^2(\sigma^1_*, \sigma^2_*) &\geq u^2(\sigma^1_*, \sigma^2) - \varepsilon, \quad \forall \sigma^2 \in \Sigma^2. && (5)
\end{aligned}
$$

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

The 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 merely a measurability condition.

[1] By convention, $A^0 = \{\emptyset\}$ contains only the empty history.

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 two-player non-zero-sum game.

Theorem 6 (Mertens, 1987)**.** *If u^1 and u^2 are bounded and Borel measurable, then an ε-equilibrium exists for every $\varepsilon > 0$.*

A stronger notion of equilibrium is the notion of subgame-perfect equilibrium. Every finite history $h = (a_1, a_2, \ldots, a_l) \in H^1 \cup H^2$, together with a pair of strategies (σ^1, σ^2), determines an infinite play $p = p(\sigma^1, \sigma^2 \mid h) = (b_t)_{t \in \mathbf{N}} \in A^{\mathbf{N}}$ as follows:

$$\begin{aligned} b_j &= a_j, \quad 1 \leq j \leq l, && (6) \\ b_{2j} &= \sigma^2(b_1, b_2, \ldots, b_{2j-1}), \quad l < 2j, && (7) \\ b_{2j+1} &= \sigma^1(b_1, b_2, \ldots, b_{2j}), \quad l < 2j+1. && (8) \end{aligned}$$

We denote by $u^i(\sigma^1, \sigma^2 \mid h) = u^i(p(\sigma^1, \sigma^2 \mid h))$ the payoff of player i at this play.

Definition 7. *Let $\varepsilon \geq 0$. A pair of strategies (σ^1_*, σ^2_*) is a* subgame-perfect ε-equilibrium *if for every finite history $h \in H^1 \cup H^2$ one has:*

$$\begin{aligned} u^1(\sigma^1_*, \sigma^2_* \mid h) &\geq u^1(\sigma^1, \sigma^2_* \mid h) - \varepsilon, \quad \forall \sigma^1 \in \Sigma^1, && (9) \\ u^2(\sigma^1_*, \sigma^2_* \mid h) &\geq u^2(\sigma^1_*, \sigma^2 \mid h) - \varepsilon, \quad \forall \sigma^2 \in \Sigma^2. && (10) \end{aligned}$$

Thus, every finite history h defines a subgame — the subgame that starts once the finite history h occurs. A strategy pair 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 $h' = (b_t)_{t=1}^m \in A^{\mathbf{N}}$, 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^i is* lower-semi-continuous *if for every sequence $(p^k)_{k \in \mathbf{N}}$ of infinite plays in $H^1 \cup H^2$ that converges to a limit p one has*

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

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

Theorem 9. *If u^1 and u^2 are lower-semi-continuous and bounded, then the game admits a subgame-perfect ε-equilibrium for every $\varepsilon > 0$.*

3 Proof

We first note that if one changes the payoffs u^1 and u^2 by at most ε (in the supremum norm), then every subgame-perfect ε-equilibrium in the original game is a subgame-perfect 3ε-equilibrium in the new game. Because the range of the payoff functions is bounded, we can assume w.l.o.g. that the range of u^1 and u^2 is in fact finite. We will show that if the payoff functions are bounded and have finite range, there is a subgame-perfect 0-equilibrium.

For $i \in \{1,2\}$ we denote by $-i$ the player who is not i. We denote by $i(h)$ the player who has to choose an action after the history h; $i(h) = 1$ if h has even length, and $i(h) = 2$ otherwise. We denote by $-i(h)$ the player who is not $i(h)$.

3.1 Subgame-perfect optimal strategies

Definition 10. *Let $h \in H^1 \cup H^2$ be a finite history. The real number $v(h)$ is called the* value at h *if*

$$v(h) = \max_{\sigma^{i(h)} \in \Sigma^{i(h)}} \min_{\sigma^{-i(h)} \in \Sigma^{-i(h)}} u^{i(h)}(\sigma^1, \sigma^2 \mid h) = \min_{\sigma^{-i(h)} \in \Sigma^{-i(h)}} \max_{\sigma^{i(h)} \in \Sigma^{i(h)}} u^{i(h)}(\sigma^1, \sigma^2 \mid h). \tag{11}$$

Because the range of the functions u^1 and u^2 is finite, the maximum and minimum in (11) are well defined.

Definition 11. *A strategy $\sigma^i_* \in \Sigma^i$ of player i is called* subgame-perfect optimal *if for every $h \in H^i$ one has*

$$u^i(\sigma^i_*, \sigma^{-i} \mid h) \geq v(h), \quad \forall \sigma^{-i} \in \Sigma^{-i}.$$

The next result follows from Martin (1975).

Theorem 12. *In every zero-sum game with bounded and discrete payoffs, if u^1 is Borel measurable then both players have subgame-perfect optimal strategies.*

3.2 Constructing some sequences

In this subsection we define for every finite history $h \in H^1 \cup H^2$ and every ordinal ξ, (a) a real number $\alpha_\xi(h)$, and (b) a set $H_\xi(h)$ of plays. The sequence $(\alpha_\xi(h))_\xi$ will be a non-decreasing sequence of lower bounds to subgame-perfect 0-equilibrium payoff for the player who makes the decision at h, in the game that start at h. The sequence $(H_\xi(h))_\xi$ will be a non-increasing (by inclusion) sequence of sets of histories, such that all plays generated by a subgame-perfect 0-equilibrium at the game that starts at h are included in every set in this sequence.

Every play $p = (a_t)_{t \in \mathbf{N}} \in A^{\mathbf{N}}$ defines a sequence $(h_n)_{n \in \mathbf{N}}$ of finite histories, where $h_n = (a_t)_{t=1}^n$ is the prefix of length n of p.

Suppose that one is given a real-valued function α that is defined over the set $H^1 \cup H^2$ of finite histories.

Definition 13. *Let h be a finite history, and p a play that extends h. The play p is called α-*monotonic *at h if the two sequences $(\alpha(h_{2k}))_{\{k\colon h \preceq h_{2k}\}}$ and $(\alpha(h_{2k+1}))_{\{k\colon h \preceq h_{2k+1}\}}$ are non-decreasing.*

A play p is called α-*viable* at a given finite history h if for each player i, the payoff $u^i(p)$ that p yields to player i is higher than $\alpha(h')$, for every prefix h' of p that extends h and that is a decision history for player i. Formally,[2]

Definition 14. *Let $\alpha : H^1 \cup H^2 \to \mathbf{R}$ be a real-valued function, and let h be a finite history. We say that a play p is α-*viable *at h if*

- *h is a prefix of p,*
- *$u^{i(h')}(p) \geq \alpha(h')$ for every history $h' \in H^1 \cup H^2$ that satisfies $h \preceq h' \prec p$.*

The following lemma lists two simple properties of the play that is realized when both players follow subgame-perfect optimal strategies. It follows from the definitions, and its proof is omitted.

Lemma 15. *Fix a finite history $h \in H^1 \cup H^2$. Let σ^1 and σ^2 be subgame-perfect optimal strategies of the two players in the subgame that starts at h. Then the realized play $p(\sigma^1, \sigma^2)$ is both α_1-monotonic and α_1-viable at h.*

Set

$$\widetilde{\alpha}_1(h) = v^{i(h)}(h), \tag{12}$$

$$H_1(h) = \{p \in A^{\mathbf{N}} \colon h \prec p, p \text{ is } \widetilde{\alpha}_1\text{-viable at } h\} \tag{13}$$

$$\alpha_1(h) = \min_{p \in H_1(h)} u^{i(h)}(p). \tag{14}$$

Thus, $\widetilde{\alpha}_1(h)$ is the value of the zero-sum subgame that "starts at h" with payoffs that are the payoffs of player $i(h)$. This is a lower bound on the equilibrium payoff of player $i(h)$ in this subgame. $H_1(h)$ is the set of all plays that yield both players at least their value in all subgames that may occur after h. Since, $\widetilde{\alpha}_1(h)$ is not necessarily the higher lower bound on the payoff of player $i(h)$ that correspond to the plays in $H_1(h)$, we set $\alpha_1(h)$ to be the higher lower bound on these payoffs. Note, every play $p \in H_1(h)$ is also α_1-viable at h. In addition, if a play $p \in H_1(h)$ is $\widetilde{\alpha}_1$-monotonic then it is also α_1-monotonic.

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 history of length $l + 1$ that starts with h and ends with a.

For every successor ordinal $\xi + 1$, define

$$\alpha_{\xi+1}(h) = \max_{a \in A} \min_{p \in H_\xi(h,a)} u^{i(h)}(p), \tag{15}$$

$$H_{\xi+1}(h) = \left\{ p \in \cup_{a \in A} H_\xi(h, a) \colon u^{i(h)}(p) \geq \alpha_{\xi+1}(h) \right\}. \tag{16}$$

[2]This definition is adapted from Flesch et al. (2008).

If $H_\xi(h,a)$ contains all plays that can be generated by subgame-perfect 0-equilibria in the subgame that starts at (h,a), then in every subgame-perfect 0-equilibrium of the subgame that starts at h player $i(h)$ will receive at least $\alpha_{\xi+1}(h)$. Moreover, after the history h, player $i(h)$ will not play an action that does not maximize the right-hand side of (15), and therefore $H_{\xi+1}(h)$ contains all plays that can be generated by subgame-perfect 0-equilibria in the subgame that starts at h.

For a limit ordinal ξ define

$$\widetilde{\alpha}_\xi(h) = \max_{\lambda<\xi} \alpha_\lambda(h), \tag{17}$$

$$H_\xi(h) = \{p \in A^{\mathbf{N}} \colon h \preceq p, p \text{ is } \widetilde{\alpha}_\xi\text{-viable at } h\}, \tag{18}$$

$$\alpha_\xi(h) = \min_{p \in H_\xi(h)} u^{i(h)}(p). \tag{19}$$

Every play $p \in H_\xi(h)$ is also α_ξ-viable at h, for every limit ordinal ξ. In addition, for every limit ordinal ξ, if a play $p \in H_\xi(h)$ is $\widetilde{\alpha}_\xi$-monotonic then it is also α_ξ-monotonic.

The following observation, which follows from the definitions, will be used later.

Lemma 16. *Suppose that $p \in H_\xi(h,a)$, where $a \in A$ achieves the maximum in the right-hand side of (15). Then $p \in H_{\xi+1}(h)$.*

3.3 Properties of the sequences $(\alpha_\xi(h))_\xi$ and $(H_\xi(h))_\xi$

The following theorem states some properties of the sequences $(\alpha_\xi(h))_\xi$ and $(H_\xi(h))_\xi$, which play a crucial role in the proof of the main result.

Theorem 17. *The following holds for every $h \in H^1 \cup H^2$:*

1. *The set $H_\xi(h)$ is not empty for every ordinal ξ.*
2. *The sequence $(H_\xi(h))_\xi$ is monotonic non-increasing (by inclusion).*
3. *The sequence $(\alpha_\xi(h))_\xi$ is monotonic non-decreasing.*
4. *For $\xi = 1$ and for every limit ordinal ξ, there is a play $p \in H_\xi(h)$ that is α_ξ-monotonic at h.*

Proof. The proof is by transfinite induction.
Part 1: $H_1(h) \neq \emptyset$ for every $h \in H^1 \cup H^2$. Moreover, there is $p \in H_1(h)$ that is α_1-monotonic at h.

This fact follows from Theorem 12, Lemma 15 and Eq. (14): once both players follows a subgame-perfect optimal strategy, the realized play is α_1-viable and α_1-monotonic at h.
Part 2: $\alpha_2(h) \geq \alpha_1(h)$ for every $h \in H^1 \cup H^2$.

Every play in $H_1(h,a)$ is α_1-viable at (h,a). Therefore $\alpha_2(h)$ is the value of the game where at the second stage player $i(h)$ receives the minimum payoff generated by an α_1-viable plays. But $\alpha_1(h)$ is the value of the game where at the second stage player $i(h)$ receives the minimum payoff generated by a play which satisfies the α_1-viability condition only for histories that are controlled by player $i(h)$. This implies that indeed $\alpha_2(h) \geq \alpha_1(h)$.

Part 3: $H_2(h) \subseteq H_1(h)$ for every $h \in H^1 \cup H^2$.

This follows from Part 2 and Eq. (16).

Part 4: If $\alpha_{\xi+1}(h) \geq \alpha_\xi(h)$ and $H_{\xi+1}(h) \subseteq H_\xi(h)$ for every $h \in H^1 \cup H^2$, then $\alpha_{\xi+2}(h) \geq \alpha_{\xi+1}(h)$ and $H_{\xi+2}(h) \subseteq H_{\xi+1}(h)$ for every $h \in H^1 \cup H^2$.

This follows from the definitions (15) and (16).

Part 5: For every limit ordinal ξ, every ordinal $\lambda < \xi$, and every $h \in H^1 \cup H^2$, one has $\alpha_\xi(h) \geq \alpha_\lambda(h)$ and $H_\xi(h) \subseteq H_\lambda(h)$.

This follows from the definitions (17), (18) and (19).

Part 6: For every ordinal ξ and every $h \in H^1 \cup H^2$, If $H_\xi(h) \neq \emptyset$ then $H_{\xi+1}(h) \neq \emptyset$.

This follows from the definitions (15) and (16).

Part 7: For every limit ordinal ξ and every $h \in H^1 \cup H^2$ one has $H_\xi(h) \neq \emptyset$. Moreover, there is a play $p \in H_\xi(h)$ that is α_ξ-monotonic at h.

This is the difficult part of the proof. Fix a limit ordinal ξ and a finite history h. We are going to generate a play that extends h, and we will show that it is in $H_\xi(h)$ and it is α_ξ-monotonic at h. The play will be generated in iterations, where 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 prove that it is in $H_\xi(h)$ and that it is α_ξ-monotonic at h.

Odd iterations:

Let h_1 be the history at the beginning of the iterations. For the first iteration, $h_1 = h$. For all other odd iterations, it is the 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 every set of ordinals has a minimal element.

2. Let a_1 be an action of player $i(h_1)$ that achieves the maximum in (15) for h_1 and ξ_1.

3. Set $h_2 = (h_1, a_1)$.

4. 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, $\xi_2 < \xi$.

5. Let a_2 be an action of player $i(h_2)$ that achieves the maximum in (15) for h_2 and ξ_2.

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

7. Continue this way to create a sequence $(h_1, \xi_1, a_1, h_2, \xi_2, a_2, \ldots)$ that extends h_1. The iteration ends when either $\xi_m = 1$ or ξ_m is a limit ordinal. If $\xi_m > 1$ is a successor ordinal for every $m \in \mathbf{N}$, the iteration never ends.

Note that because ξ is a limit ordinal we necessarily have $\xi_m < \xi$, for every $m \in \mathbf{N}$.

The next lemma states that odd iterations are finite.

Lemma 18. *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

$$H_{\xi_m - 1}(h_{m+1}) \supseteq H_{\xi_{m+2}-1}(h_{m+3}). \tag{20}$$

Note that because ξ_m and ξ_{m+2} are successor ordinals, the two ordinals $\xi_m - 1$ and $\xi_{m+2} - 1$ are well defined. Now, let p be any play in $H_{\xi_{m+2}-1}(h_{m+3})$. By Lemma 16, $p \in H_{\xi_{m+2}}(h_{m+2})$. Because $\xi_{m+1} - 1 \leq \xi_{m+2}$, by the induction hypothesis (Part 4), $p \in H_{\xi_{m+1}-1}(h_{m+2})$. By Lemma 16, $p \in H_{\xi_{m+1}}(h_{m+1})$. Because $\xi_m - 1 \leq \xi_{m+1}$, by the induction hypothesis (Part 4), $p \in H_{\xi_m - 1}(h_{m+1})$, as desired.

Because both ξ_m and ξ_{m+2} are successor ordinals, Eq. (15) implies that for every $m \in \mathbf{N}$

$$\alpha_{\xi_m}(h_m) \leq \alpha_{\xi_{m+2}}(h_{m+2}). \tag{21}$$

Because $\widetilde{\alpha}_\xi(h_m) = \alpha_{\xi_m}(h_m)$ and $\widetilde{\alpha}_\xi(h_{m+2}) = \alpha_{\xi_{m+2}}(h_{m+2})$ we deduce that for every $m \in \mathbf{N}$

$$\widetilde{\alpha}_\xi(h_m) \leq \widetilde{\alpha}_\xi(h_{m+2}). \tag{22}$$

Because the payoffs are discrete, the inequality $\widetilde{\alpha}_\xi(h_m) \leq \widetilde{\alpha}_\xi(h_{m+2})$ can be strict only finitely many times. That is, there is $M \in \mathbf{N}$ sufficiently large such that $\widetilde{\alpha}_\xi(h_m) = \widetilde{\alpha}_\xi(h_{m+2})$ for every $m \geq M$. Following the argument that Eq. (22) holds we deduce that for each such m we have

$$\widetilde{\alpha}_\xi(h_m) = \alpha_{\xi_m}(h_m) = \alpha_{\xi_{m+1}-1}(h_{m+2}) = \alpha_{\xi_{m+2}}(h_{m+2}) = \widetilde{\alpha}_\xi(h_{m+2}), \tag{23}$$

so that $\xi_{m+2} = \xi_{m+1} - 1$. Because this equality holds for every m sufficiently large, and because there is no infinite decreasing sequence of ordinals, there is m such that either $\xi_m = 1$ or ξ_m is a limit ordinal, as desired. □

Even iterations:

Let h_1 be the history that was generated by the previous iterations. Then it is the output of the previous odd iteration. Denote by λ the last ordinal ξ_m generated in the previous odd iteration. Then 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 hypothesis (Part 1 or Part 7), there is a play $p \in H_\lambda(h_1)$ that is α_λ-monotonic at h_1.

By the induction hypothesis (Parts 2, 3 and 4), $\widetilde{\alpha}_\xi(h') \geq \alpha_\lambda(h')$ 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. Otherwise, the output of the current even iteration is the shortest prefix h' of p that extends h_1 for which $\widetilde{\alpha}_\xi(h') > \alpha_\lambda(h')$.

Because $p \in H_\lambda(h_1)$, the play p is α_λ-viable at h_1. Because $\widetilde{\alpha}_\xi(h') = \alpha_\lambda(h')$ for prefixes h' of p that extend h_1, p satisfies the condition of $\widetilde{\alpha}_\xi$-viability for all such prefixes.

Lemma 19. *Let p_* be the play that we just generated. Then p_* is $\widetilde{\alpha}_\xi$-monotonic at h.*

Proof. For the part of the play added in odd iterations the monotonicity was proved in (22). For the part of the play added in even iterations it follows from the construction. □

We are now ready to prove Part 7.

Lemma 20. *$p_* \in H_\xi(h)$.*

Proof. If the number of iterations if finite, so that the last even iteration is infinite, then the play p_* is $\widetilde{\alpha}_\xi$-viable at h. Indeed, by construction it is $\widetilde{\alpha}_\xi$-viable at h_0, the history at the beginning of the last even iteration. The claim now follows from Lemma 19.

We now show that it cannot be that the number of iterations is infinite. Denote by $(h_n)_{n\in\mathbf{N}}$ all finite prefixes of p_*. Because p_* is $\widetilde{\alpha}_\xi$-monotonic at h, the two sequences $(\widetilde{\alpha}_\xi(h_{2k}))_{h\preceq h_{2k}}$ and $(\widetilde{\alpha}_\xi(h_{2k+1}))_{h\preceq h_{2k+1}}$ are non-decreasing. Because the range of the payoffs is finite, these two sequences are eventually constant. At the beginning of an even iteration we have $\widetilde{\alpha}_\xi(h_1) = \alpha_\lambda(h_1)$, and the sequence $(\alpha_\lambda(h'))_{h'}$ is non-decreasing in the part of p^* that is added in an even iteration (for the definition of λ, see the construction for even iterations). But $\widetilde{\alpha}_\xi(h') \geq \alpha_\lambda(h')$ for every prefix h' of p_* that is added in the even iteration, so that $\widetilde{\alpha}_\xi(h') = \alpha_\lambda(h')$ for every such h'. In particular, the even iteration does not end. □

Finally, since the play p_* that we generated is in $H_\xi(h)$ and is $\widetilde{\alpha}_\xi$-monotonic at h, it is also α_ξ-monotonic at h. □

By Theorem 17(2) and 17(3) it follows that

Theorem 21. *There is an ordinal ξ_* such that $\alpha_{\xi_*}(h) = \alpha_{\xi_*+1}(h)$ and $H_{\xi_*}(h) = H_{\xi_*+1}(h)$ for every $h \in H^1 \cup H^2$.*

3.4 Proof of Theorem 9

We now construct a pair (σ^1_*, σ^2_*) of strategies, and show that they form a subgame-perfect 0-equilibrium.

For every finite history h choose an α_{ξ_*}-viable play $p(h)$ that extends h and that satisfies

$$u^{-i(h)}(p(h)) = \min_{p \in H_{\xi_*}(h)} u^{-i(h)}(p). \tag{24}$$

If player i deviates, and h is the finite history right after the deviation (so that $i = -i(h)$), then $p(h)$ is an α_{ξ_*}-viable play at h that minimizes player's i's payoff.

Let σ^1_* be the following strategy: Follow the play $p(\emptyset)$ as long as player 2 follows $p(\emptyset)$. Suppose that at stage $2k_1$ player 2 deviates from $p(\emptyset)$. From stage $2k_1 + 1$ and on follow the play $p(h_{2k_1+1})$ as long as player 2 follows this play. Suppose that at stage $2k_2$ player 2 deviates from $p(h_{2k_1+1})$. From stage $2k_2 + 1$ and on follow the play $p(h_{2k_2+1})$ as long as player 2 follows this play. Continue this way.

The strategy σ^2_* of player 2 is defined symmetrically.

We now show that (σ^1_*, σ^2_*) is a subgame-perfect 0-equilibrium. To this end we fix a finite history $h \in H^1$ and we show that

$$u^2(\sigma^1_*, \sigma^2 \mid h) \leq u^2(\sigma^1_*, \sigma^2_* \mid h), \quad \forall \sigma^2 \in \Sigma^2.$$

One can use similar argument to show that the analog inequality for player 1 holds as well, so that (σ^1_*, σ^2_*) is indeed a subgame-perfect 0-equilibrium.

Let $\sigma^2 \in \Sigma^2$ be any strategy of player 2. Let $p_* = p(\sigma^1_*, \sigma^2_* \mid h)$ be the play induced by (σ^1_*, σ^2_*) given h. This is the play that is generated if player 2 does not deviate. Let $p = p(\sigma^1_*, \sigma^2 \mid h)$ be the play when player 2 deviates to σ^2.

Denote by $2k_1, 2k_2, \ldots$ the stages where σ^2 and σ^2_* differ along p; in those stages player 1 observes the deviations of player 2. The sequence $(2k_j)_j$ may be finite or infinite. Denote by $p_j = p(\sigma^1_*, \sigma^2)$ the play that player 1 start to follow at stage $2k_j$, for each j.

We complete the proof by showing that

$$u^2(p) \leq u^2(p_*). \tag{25}$$

It is sufficient to show that

$$u^2(p_j) \leq u^2(p_*), \quad \forall j. \tag{26}$$

If σ^2 and σ^2_* differ only finitely many times along p, Eq. (25) follows from Eq. (26). If σ^2 and σ^2_* differ infinitely many times along p, then the sequence $(p_j)_{j \in \mathbf{N}}$ converges to p, so that Eq. (25) follows from Eq. (26) and the lower-semi-continuity of u^2. This is the only place in the proof where the lower-semi-continuity of the payoff functions is used.

The proof of (26) is by induction on j. Because p_* is α_{ξ_*}-viable at h, $u^2(p_1) = \alpha_{\xi_*}(h_{2k_1}) \leq u^2(p_*)$. For every $j \geq 1$, because p_j is α_{ξ_*}-viable at h_{2k_j}, $u^2(p_{j+1}) = \alpha_{\xi_*}(h_{2k_j}) \leq u^2(p_j)$, which is at most $u^2(p_*)$ by the induction hypothesis. The proof is now complete.

3.5 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 22. *A play p is induced by some subgame-perfect 0-equilibrium if and only if p is α_{ξ_*}-viable.*

Proof. If p is α_{ξ_*}-viable, then the construction in Section 3.4 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^1_*, σ^2_*) is a subgame-perfect 0-equilibrium, then $p(\sigma^1_*, \sigma^2_* \mid h)$ is α_ξ-viable at h, for every ordinal ξ and every finite history h.

Because $\alpha_1(h)$ is the value of the subgame that starts at h, and because every subgame-perfect 0-equilibrium is an equilibrium in the subgame that starts at h, the claim follows for $\xi = 1$.

Suppose now that the claim holds for an ordinal ξ. Let h be any finite history. Because the claim holds for ξ, the play $p(\sigma^1_*, \sigma^2_* \mid h)$ is α_ξ-viable at (h, a) for every a, and therefore in any subgame-perfect 0-equilibrium of the subgame that starts at h, the payoff to player $i(h)$ is at least $\alpha_{\xi+1}(h)$. This implies that $p(\sigma^1_*, \sigma^2_* \mid h)$ is $\alpha_{\xi+1}$-viable at h, for every h.

Finally, the definition of α_ξ for limit ordinals ξ implies that if $p(\sigma^1_*, \sigma^2_* \mid h)$ is α_λ-viable at h for every ordinal $\lambda < \xi$, then it is also α_ξ-viable at h. □

4 Comments

4.1 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 was the first to play b, 2 if he played a and some player $j > t$ played b, and 0 otherwise. We claim that there is no subgame-perfect 0-equilibrium in this game. Suppose to the contrary that a subgame-perfect 0-equilibrium exists, and for each $t \in \mathbf{N}$ let $h^t = (a, a, \dots, a)$ be the finite history composed of t times the action a. If no player $j > t$ plays b after the finite history h^{j-1}, then after the history h^{t-1} player t plays b; this implies that infinitely many players t play b after h^{t-1}. However, if under this equilibrium player t plays b after h^{t-1}, then all players $j < t$ play a after h^{j-1}; this implies that no player t plays b after h^{t-1} in this equilibrium. Thus a subgame-perfect 0-equilibrium does not exist.

4.2 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 a third 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.

4.3 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. Many 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 (2000), 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. (2008) 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 depends only on the vertex where termination occurred, and not on the whole past play.

Flesch et al. (2008) prove that every such game admits a subgame-perfect 0-equilibrium.[3] 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. (2008) in two respects. First, when the number of vertices is infinite, there need not by $k_* \in \mathbf{N}$ 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, one also needs to take into account the possible plays, and introduce the sets $(H_\xi(h))_\xi$.

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

[3] When transitions are random, Flesch et al. (2008) prove the existence of a subgame-perfect ε-equilibrium, for every $\varepsilon > 0$.

there is a α_ξ-viable path. For positive recursive games one can prove that the same holds for every ordinal ξ. We do not know whether this property holds for Borel games with lower-semi-continuous payoffs as well.

4.4 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 (2008).

In our proof the lower-semi-continuity of the payoff functions was used only in the last part, to show that any deviation σ^2 that differs from σ^2_* infinitely many times cannot be profitable, as soon as any deviation σ^2 that differs from σ^2_* 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 H_ξ is not appropriate for general Borel games. Consider a 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

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 $\alpha_\xi(h) = 1$ for every finite history h and every ordinal ξ, so that the folk theorem, Theorem 22, does not hold, and our construction of the subgame-perfect 0-equilibrium in the proof of Theorem 12 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. (2008) Subgame Perfection in Positive Recursuve Games with Perfect Information, preprint.

[3] Gale D. and Stuart F. (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*, Volume 2996, 4-18.

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

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

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

[8] Mashiah-Yaakovi A. (2008) Subgame-Perfect Equilibrium in Stopping Games with Perfect Information. Preprint.

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

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

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

[12] 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.

[13] Shmaya E. (2008) Many inspections are manipulable. *Theoretical Economics*, **3**, 367382.

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

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

[16] 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.

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