

Subgame-perfect ϵ -equilibria in perfect information games with sigma-discrete discontinuities

János Flesch · Arkadi Predtetchinski

Received: 3 November 2014 / Accepted: 24 February 2015 / Published online: 17 March 2015
© The Author(s) 2015. This article is published with open access at Springerlink.com

Abstract Multi-player perfect information games are known to admit a subgame-perfect ϵ -equilibrium, for every $\epsilon > 0$, under the condition that every player's payoff function is bounded and continuous on the whole set of plays. In this paper, we address the question on which subsets of plays the condition of payoff continuity can be dropped without losing existence. Our main result is that if payoff continuity only fails on a sigma-discrete set (a countable union of discrete sets) of plays, then a subgame-perfect ϵ -equilibrium, for every $\epsilon > 0$, still exists. For a partial converse, given any subset of plays that is not sigma-discrete, we construct a game in which the payoff functions are continuous outside this set but the game admits no subgame-perfect ϵ -equilibrium for small $\epsilon > 0$.

Keywords Perfect information games · Subgame perfect equilibrium · Discontinuous games

JEL Classification C73

We would like to thank George Koumoullis, Ayala Mashiah-Yaakovi, Alex Ravsky, Eilon Solan, the Editor and the Reviewer for their helpful comments.

J. Flesch
Department of Quantitative Economics, Maastricht University, P.O.Box 616, 6200 MD Maastricht,
The Netherlands
e-mail: j.flesch@maastrichtuniversity.nl

A. Predtetchinski (✉)
Department of Economics, Maastricht University, P.O.Box 616, 6200 MD Maastricht, The Netherlands
e-mail: a.predtetchinski@maastrichtuniversity.nl

1 Introduction

We consider multi-player perfect information games, without chance moves, which have arbitrary action spaces and infinite duration. As a consequence of Flesch et al. (2010) and Purves and Sudderth (2011), these games admit a subgame-perfect ϵ -equilibrium, ϵ -SPE for brevity, for every $\epsilon > 0$, provided that every player's payoff function is bounded and continuous on the whole set of plays. Here, continuity is meant with respect to the product topology on the set of plays, with the set of actions given its discrete topology.

The question naturally arises what happens in games in which continuity of the payoffs is not satisfied everywhere, but only on a subset of plays. More precisely, on which subsets of plays discontinuity can be allowed without losing that an ϵ -SPE exists for every $\epsilon > 0$. We show that discontinuities can be allowed on a sigma-discrete set of plays, i.e., on a countable union of discrete sets. In the special case, when the action space is countable, a set of plays is sigma-discrete if and only if it is countable.

Our main findings are as follows:

- [1] If the set of discontinuities is sigma-discrete, then an ϵ -SPE exists for every $\epsilon > 0$.
- [2] Given a set of plays that is not sigma-discrete, there exists a game in which the payoff functions are continuous outside of the given set, and the game admits no subgame-perfect ϵ -equilibrium for small $\epsilon > 0$. This is achieved by embedding a game given in Flesch et al. (2014), which has no ϵ -SPE for small $\epsilon > 0$.

The structure of the paper is follows. In Sect. 2, we define the model. Then, in Sect. 3, we present the main results and discuss the main ideas of the proof. Sections 4 and 5 contain the proofs, and Sect. 6 provides some concluding remarks.

2 The model and preliminaries

The game Let $N = \{1, \dots, n\}$ denote the set of players and let A be an arbitrary non-empty set. Let $\mathbb{N} = \{0, 1, 2, \dots\}$. We denote by H the set of all finite sequences of elements of A , including the empty sequence \emptyset , and we denote by $P = A^{\mathbb{N}}$ the set of all infinite sequences of elements of A . The elements of A are called actions, the elements of H are called histories, and the elements of P are called plays. There is a function $\iota : H \rightarrow N$ which assigns an active player to each history. Further, each player $i \in N$ is given a so-called payoff function $u_i : P \rightarrow \mathbb{R}$.

The game is played as follows: At period 0, player $\iota(\emptyset)$ chooses an action a_0 . In general, suppose that up to period $t \in \mathbb{N}$ of the game, the sequence $h = (a_0, \dots, a_t)$ of actions has been chosen. Then, at period $t + 1$, player $\iota(h)$ chooses an action a_{t+1} . The chosen action is observed by all players. Continuing this way, the players generate a play $p = (a_0, a_1, \dots)$, and finally, each player $i \in N$ receives payoff $u_i(p)$.

We remark that this setup encompasses all games of finite duration. Another important special case is the situation when the players receive instantaneous payoffs at every period of the game and then aggregate them into one payoff, for example, by taking the discounted sum.

The topological structure We endow the set A with the discrete topology and $P = A^{\mathbb{N}}$ with the product topology \mathcal{T} . A basis for (P, \mathcal{T}) is formed by the cylinder sets $P(h) = \{p \in P : h \prec p\}$ for $h \in H$, where for a history $h \in H$ and a play $p \in P$ we write $h \prec p$ if h is the initial segment of p . In this topology, a sequence of plays $(p_n)_{n \in \mathbb{N}}$ converges to a play p precisely when for every $k \in \mathbb{N}$ there exists an $N_k \in \mathbb{N}$ such that for every $n \geq N_k$, the first k coordinates of p_n coincide with those of p . Note that (P, \mathcal{T}) is completely metrizable¹, and moreover, it is separable (and hence Polish) if and only if A is countable².

A function $f : P \rightarrow \mathbb{R}$ is said to be *continuous* at a play $p \in P$ if, for every sequence of plays $(p_n)_{n \in \mathbb{N}}$ converging to p , we have $\lim_{n \rightarrow \infty} f(p_n) = f(p)$. Thus, f is continuous at p precisely when for every $\delta > 0$ there is a history $h \prec p$ such that $|f(p) - f(q)| \leq \delta$ for every $q \in P(h)$. Intuitively, continuity at p means that, after following p for a large number of periods, further actions have little effect on the value of f . The function f is said to be *discontinuous* at a play p if it is not continuous at p . Further, f is said to be *continuous* if it is continuous at each play in P .

Subgames For the concatenation of histories and actions, we use the following notations. For a history $h = (a_0, \dots, a_t) \in H$ and finite sequence of actions (b_0, \dots, b_m) in A , let $(h, b_1, \dots, b_m) = (a_0, \dots, a_t, b_1, \dots, b_m)$, whereas for an infinite sequence of actions (b_0, b_1, \dots) in A , let $(h, b_0, b_1, \dots) = (a_0, \dots, a_t, b_0, b_1, \dots)$.

Consider a history $h = (a_0, \dots, a_t)$ for some $t \in \mathbb{N}$. The subgame $G(h)$ corresponding to h is played as follows: First, player $\iota(h)$ chooses an action a_{t+1} . In general, suppose that the sequence of actions $(a_{t+1}, \dots, a_{t+m})$ is chosen. Then, player $\iota(h, a_{t+1}, \dots, a_{t+m})$ chooses an action a_{t+m+1} . Continuing this way, the players generate a play $p = (h, a_{t+1}, a_{t+2}, \dots)$ in $P(h)$, and finally, each player $i \in N$ receives payoff $u_i(p)$.

Strategies A strategy for player i is a function $\sigma_i : \iota^{-1}(i) \rightarrow \Delta_c(A)$, where $\iota^{-1}(i)$ is the set of histories where player i moves and where $\Delta_c(A)$ is the set of probability measures on A with a countable support. The interpretation is that if a history $h \in \iota^{-1}(i)$ arises, then σ_i prescribes player i to choose an action according to $\sigma_i(h)$. A strategy is called *pure*, if it always places probability one on one action. A strategy profile is a tuple $(\sigma_1, \dots, \sigma_n)$ where σ_i is a strategy for player i . Given a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ and a strategy η_i for player i , we write σ/η_i to denote the strategy profile obtained from σ by replacing σ_i with η_i .

Every strategy profile σ induces a probability measure μ_σ on the Borel sets of P . Similarly, by considering the subgame for a history $h \in H$, every strategy profile σ

¹ One can take, for example, the metric $d : P \times P \rightarrow \mathbb{R}$ which is defined for each $p, q \in P$ as follows: if $p = q$, then $d(p, q) = 0$, and otherwise $d(p, q) = 2^{-m(p,q)}$ where $m(p, q) \in \mathbb{N}$ is the least $m \in \mathbb{N}$ such that $p_m \neq q_m$. For a detailed account of the properties of the metric d , see Kechris (1995, p.7).

² Suppose first that A is uncountable. Consider the cylinder sets $P(a)$, for $a \in A$. They are non-empty, mutually disjoint, and there are uncountably many of them, so (P, \mathcal{T}) is not separable. Now suppose that A is countable, which implies that H is countable too. For every history $h \in H$, take an arbitrary play $p_h \succ h$. Then, the set $\{p_h \mid h \in H\}$ is countable and dense in (P, \mathcal{T}) .

induces a probability measure $\mu_{\sigma,h}$ on the Borel sets of P , such that $\mu_{\sigma,h}(P(h)) = 1$. We denote the expected payoff for player $i \in N$ in the beginning of the game by $u_i(\sigma)$ and in the subgame starting at h by $u_i(\sigma, h)$.

Subgame-perfect ϵ -equilibrium Let $\epsilon \geq 0$ be an error term. A strategy profile σ is called an ϵ -equilibrium if no player can gain more than ϵ by a unilateral deviation, i.e., if for each player $i \in N$ and for each strategy σ'_i of player i , it holds that

$$u_i(\sigma) \geq u_i(\sigma/\sigma'_i) - \epsilon.$$

A stronger concept arises if we require that the strategy profile induces an ϵ -equilibrium in every subgame. A strategy profile σ is called a *subgame-perfect ϵ -equilibrium*, ϵ -SPE for brevity, if for each history $h \in H$, each player $i \in N$, and each strategy σ'_i of player i , it holds that

$$u_i(\sigma, h) \geq u_i(\sigma/\sigma'_i, h) - \epsilon.$$

A 0-equilibrium is simply called an equilibrium, and a 0-SPE is simply called an SPE.

Existence of ϵ -SPE An ϵ -SPE exists for every $\epsilon > 0$ provided that each player's payoff function is bounded and continuous. This follows from more general results in Flesch et al. (2010) and Purves and Sudderth (2011). Carmona (2005) shows that for every $\epsilon > 0$, an ϵ -SPE exists under the assumption that the payoff functions are bounded and continuous at infinity. Continuity at infinity implies continuity but not vice versa.³ Flesch et al. (2014) describe a game in which every player's payoff function is bounded and Borel measurable, yet the game admits no ϵ -SPE for small $\epsilon > 0$.

Sigma-discrete and perfect sets Let X be a topological space. Consider a subset $D \subseteq X$. A point $x \in D$ is called an isolated point of D if x has an open neighborhood that contains no point of $D \setminus \{x\}$. The set D is discrete if every point of D is isolated, and it is sigma-discrete if it is a countable union of discrete sets. If X is completely metrizable and separable (i.e., Polish), then D is sigma-discrete if and only if it is countable.

A non-empty subset of X is said to be perfect if it is closed and has no isolated points. By Koumoullis (1984), if X is completely metrizable, then a non-empty Borel subset D of X is sigma-discrete if and only if D does not contain a perfect subset of X . Since the set of plays P is completely metrizable, these results apply to Borel subsets of P .

³ For instance, consider a game with the action set $A = \mathbb{N}$. Fix an action $a^* \in A$ and let the function u_i be given by $u_i(a_0, a_1, \dots) = 1$ if $a_{a_0} = a^*$ and 0 otherwise. Then, u_i is continuous but is not continuous at infinity.

3 Main results

In this section, we present our main results.

Theorem 1 *Consider a perfect information game with bounded payoff functions. Suppose that for every player i the payoff function u_i is continuous outside a sigma-discrete subset of P . Then, for every $\epsilon > 0$, the game admits an ϵ -SPE.*


In Appendix 7.3, we prove that the payoff functions satisfying the condition of theorem 1 are Borel measurable. Let D_i denote the set of plays p such that u_i is not continuous at p , and let $D = \cup_{i \in N} D_i$. Under the hypothesis of theorem 1, the set D_i is sigma-discrete for each $i \in N$. Since the set of players N is finite, this is obviously equivalent to the requirement that the set D be sigma-discrete.

If D is a countable set, it is sigma-discrete. Moreover, if the set of actions A is countable, then D is sigma-discrete if and only if it is countable. The set of eventually constant plays is an example of a sigma-discrete set that is not discrete. Another such example is the set of eventually periodic plays.

The proof of theorem 1 is carried out in three steps. As a first step, we discretize the payoffs in the following sense: We show that there exist payoff functions $\bar{u}_1, \dots, \bar{u}_n$ such that each \bar{u}_i has finite range and is ϵ -close to u_i , and the set of discontinuity points \bar{D}_i of \bar{u}_i is contained in D_i . This approach has the advantage that for each play p that is a continuity point of each function \bar{u}_i there exists a history $h < p$ such that each \bar{u}_i is constant on the set $P(h)$. Obviously, the subgame starting at history h has a δ -SPE for every $\delta > 0$. In contrast, the existence of such a history in the game with the original payoff functions is not trivial to establish. It is for that reason that working with discretized payoffs is crucial for our method.

As a second step, we prove theorem 1 for the so-called stopping games with finite payoff range. Finally, we show theorem 1 along the following lines. We define H^* as the set of histories $h \in H$ such that the subgame $G(h)$ has an ϵ -SPE for each $\epsilon > 0$. We prove that $H^* = H$ by contradiction: assuming that $H^* \neq H$ we show that there is a history $h \in H \setminus H^*$ such that $G(h)$ can be reduced to a stopping game and thus shown to have an ϵ -SPE for each $\epsilon > 0$.

Theorem 1 cannot be strengthened to conclude that there is a pure ϵ -SPE, a counterexample is readily provided by Solan and Vieille (2003). The game is a two-player stopping game in which the two players move alternately, and each player can either stop the game or continue:



The game admits no pure ϵ -SPE for $\epsilon < 1$, but does admit an ϵ -SPE in randomized strategies: Player 1 always stops with probability 1, while player 2 always stops with probability ϵ . The main idea is that randomization by player 2 makes it impossible for player 1 to ever reach the only point of discontinuity in the game, the play (c, c, c, \dots) . This conveys much of the intuition for how randomized strategies will be used in the proof of theorem 1.

We also prove a partial converse to theorem 1. Recall that the game is specified by the set of players, the set of actions, the assignment of the active players to histories, and the payoff functions.

Theorem 2 *Let A be any set of actions, and let D be a subset of $P = A^{\mathbb{N}}$ that is not sigma-discrete. Then, there exists a two-player game with action set A and with payoff functions that are continuous outside of D , such that the game admits no ϵ -SPE for small $\epsilon > 0$.*

The proof makes use of a game in Flesch et al. (2014) that admits no ϵ -SPE for small $\epsilon > 0$. We embed this game into D and define the payoffs elsewhere in such a way that the large game admits no ϵ -SPE for small $\epsilon > 0$ either.

The question that motivates the paper is for what kind of subsets of plays the condition of payoff continuity can be dropped without losing the existence of ϵ -SPE. As shown above, the answer is sigma-discrete subsets of plays. Surprisingly, the answer is not strongly related to the topological notion of denseness, as illustrated by the following two examples.

Take D to be the set of eventually constant plays. Then, D is dense. However, D is sigma-discrete, and hence, theorem 1 implies that any game with payoffs continuous outside of D has an ϵ -SPE for each $\epsilon > 0$. Conversely, let the action set be $A = \{0, 1, 2\}$, and D be the set $\{1, 2\}^{\mathbb{N}}$. Then, D is nowhere dense. It is also perfect and thus is not sigma-discrete. Hence, theorem 2 implies that there is a game with payoff functions that are continuous outside of D which has no ϵ -SPE for small $\epsilon > 0$.

4 The proof of theorem 1

4.1 A reduction to a finite range of payoffs

In this section, we argue that it suffices to prove theorem 1 for games with payoff functions having finite range.

Lemma 3 *Let $f : P \rightarrow [-r, r]$ be a Borel measurable function and let $\epsilon > 0$. Then, there exists a Borel measurable function $\bar{f} : P \rightarrow [-r, r]$ such that [1] \bar{f} has finite range, [2] for every $p \in P$, $|\bar{f}(p) - f(p)| < \epsilon$, and [3] if f is continuous at a play $p \in P$, then so is \bar{f} .*

The proof of lemma 3 can be found in the Appendix.

We say that \bar{f} is an ϵ -discretization of the function f .


Now let G be a game with the payoff functions u_1, \dots, u_n satisfying the assumptions of theorem 1. Fix an $\epsilon > 0$, and let $\bar{u}_1, \dots, \bar{u}_n$ be ϵ -discretizations of the respective payoff functions. Then, the set of discontinuity points of the function \bar{u}_i is contained in the set of the discontinuity points of the function u_i and is therefore sigma-discrete. Moreover, any ϵ -SPE of the game with the payoff functions $\bar{u}_1, \dots, \bar{u}_n$ is a 2ϵ -SPE of the game with the payoff functions u_1, \dots, u_n .

This shows that it is sufficient to prove theorem 1 for games with payoff functions having finite range. In view of this result, we henceforth restrict our attention to games with a finite range of payoffs.

4.2 Stopping games with a finite range of payoffs

A game G is said to be a *stopping game* if the action space is $A = \{s, c\}$, where s stands for “stop”, c stands for “continue”, and if for each $t \in \mathbb{N}$ and each player i , the payoff function u_i is constant on the set $P(c^t, s)$. Here, we write (c^t, s) to denote the history where action c has been played successively t times followed by the action s . So, in a stopping game, the payoffs are fixed once the active player decides to play action “stop”.

For each $t \in \mathbb{N}$ and each player i , let $v^t = v(c^t)$, and let r_i^t denote player i 's constant payoff on $P(c^t, s)$ and let r_i^∞ denote $u_i(c^\infty)$. A stopping game can be represented as follows:



Lemma 4 *A stopping game with a finite range of payoffs admits an ϵ -SPE for each $\epsilon > 0$.*

This result follows from a more general theorem in [Mashiah-Yaakovi \(2014\)](#). For the sake of completeness, we give a direct proof of the lemma, see the Appendix.

4.3 Games with a finite range of payoffs

In this section, we prove theorem 1 for games with a finite range of payoffs. Let G be a game satisfying the conditions of theorem 1 such that moreover the payoff functions u_1, \dots, u_n all have finite range. Let D_i denote the set of plays p such that u_i is not continuous at p , and let $D = \cup_{i \in N} D_i$. Then, for each $p \in P \setminus D$, there exists a history $h < p$ such that each u_i is constant on the set $P(h)$. Consider

$$H^* = \{h \in H : \text{for each } \epsilon > 0 \text{ the game } G(h) \text{ has an } \epsilon\text{-SPE}\}.$$

The following lemma is straightforward.

Lemma 5 *Let $h \in H$. Then, $h \in H^*$ if and only if for each $a \in A$, $(h, a) \in H^*$.*

Define the set

$$Q = \{q \in P : \text{for each } h < q \text{ it holds that } h \notin H^*\}.$$

Lemma 6 *The set Q has the following properties:*

- [1] Q is closed,
- [2] $Q \subseteq D$,
- [3] $H^* = \{h \in H : P(h) \cap Q = \emptyset\}$,
- [4] Q is empty if and only if $H^* = H$.

Proof [1] To prove that Q is closed we show that the complement of Q is open. Thus, take $p \in P \setminus Q$. Then, there is an $h < p$ such that $h \in H^*$, and hence, $P(h) \subseteq P \setminus Q$.

[2] To prove that $Q \subseteq D$ take a $p \in P \setminus D$. Since this is a continuity point of u_1, \dots, u_n , and since each payoff function has finite range, there is an $h < p$ such that for every $i \in N$ the function u_i is constant on $P(h)$. But then $h \in H^*$ as any strategy profile is an SPE of $G(h)$. Hence, $p \in P \setminus Q$.

[3] The inclusion $H^* \subseteq \{h \in H : P(h) \cap Q = \emptyset\}$ is trivial. We prove the converse. Thus, take $h \in H$ such that $P(h) \cap Q = \emptyset$. Suppose that $h \in H \setminus H^*$. We recursively define a sequence $h_0 < h_1 < \dots$ of elements of $H \setminus H^*$, as follows: Let $h_0 = h$. Suppose we have defined $h_k \in H \setminus H^*$. By lemma 5, there is an action $a_k \in A$ such that $(h_k, a_k) \in H \setminus H^*$. Let $h_{k+1} = (h_k, a_k)$. Define $p = (h, a_0, a_1, \dots)$. Then, $p \in P(h) \cap Q$, which is a contradiction.

[4] It follows immediately from [3]. □

To prove theorem 1, we need to show that $H = H^*$. Suppose not. Then, Q is non-empty. If Q has no isolated points, by the above lemma, it is a perfect set contained in D , contradicting the assumption of theorem 1. Therefore, Q has an isolated point, say $p \in Q$. Choose an $h \in H$ such that $P(h) \cap Q = \{p\}$. We argue that $h \in H^*$, thus obtaining a contradiction to the item [3] of the lemma above.

Thus, take an $\epsilon > 0$. We show that $G(h)$ has an ϵ -SPE.

For this purpose, we argue that $G(h)$ can be reduced to a stopping game. First, we rename the actions so that $p = (c, c, c, \dots)$. Hence, $h = c^k$ for some $k \in \mathbb{N}$. Take an $m \geq k$, and consider the history c^m . If A is a singleton, there is nothing to prove. So we assume that A contains at least two actions. For each $a \neq c$, $P(c^m, a)$ does not contain p and is contained in $P(h)$. As $P(h) \cap Q = \{p\}$, we have $P(c^m, a) \cap Q = \emptyset$. Therefore, by item [3] of the previous lemma, $(c^m, a) \in H^*$. This allows us to replace the subgame $G(c^m, a)$ by a payoff on some $\frac{1}{2}\epsilon$ -SPE of $G(c^m, a)$, denoted $(w_1^m(a), \dots, w_n^m(a))$. Choose an action $a^m \in A \setminus \{c\}$ so as to maximize $w_i^m(a)$ for the player i active at the history c^m . We can rename the actions so that $a^m = s$ for each $m \geq k$.

We thus obtain a game $G'(h)$ where there are two distinguished actions, s and c . Playing any action other than c terminates the game (in the sense that subsequent actions do not affect the payoffs). Should a player choose to terminate the game, action s yields the highest payoff. It is clear that all actions other than c and s can be safely ignored, so that $G'(h)$ becomes a stopping game.

By lemma 4, the game $G'(h)$ has an $\frac{1}{2}\epsilon$ -SPE. It is clear that the combination of the $\frac{1}{2}\epsilon$ -SPE in $G'(h)$ with the chosen $\frac{1}{2}\epsilon$ -SPE in the subgames $G(c^m, a)$, for $m \geq k$ and $a \in A \setminus \{c\}$, constitutes an ϵ -SPE of $G(h)$.

5 Proof of theorem 2

In this section, we prove theorem 2. The proof makes use of the following game in Flesch et al. (2014) that does not admit an ϵ -SPE for small $\epsilon > 0$: There are two players, and the set of actions is $A = \{1, 2\}$. Player 1 starts the game. The active player decides who the next active player is by choosing the corresponding action. The payoffs are $(-1, 2)$ if player 2 is active only finitely many times, $(-2, 1)$ if player 1 is active only finitely many times, and $(0, 0)$ if both players are active infinitely many times. In this game, the payoff functions are Borel, but discontinuous at every play. Flesch et al. (2014) prove that this game admits no ϵ -SPE for $\epsilon \in (0, 0.1]$.

Proof of theorem 2 Let A be any set of actions, and let D be a subset of $P = A^{\mathbb{N}}$ that is not sigma-discrete. We construct a two-player game G , with action set A and with payoff functions that are continuous outside of D , such that the game admits no ϵ -SPE for small $\epsilon > 0$.

Step 1: Construction of the game G . Since D is not sigma-discrete, it contains a perfect subset P^* of the set of plays P . We first construct a subset D^* of P^* , on which we can define a game that is strategically equivalent with the game in Flesch et al. (2014).

Let S denote the set of finite sequences of elements of $\{1, 2\}$. Define a function $\varphi : S \rightarrow H$ recursively on the length of the sequence s . Let $\varphi(\emptyset) = \emptyset$. Suppose $\varphi(s)$ has been defined so that $P(\varphi(s)) \cap P^* \neq \emptyset$. The set $P(\varphi(s)) \cap P^*$ contains at least two distinct plays, say p and p' . Let h be the longest common prefix of p and p' , and let a and a' be the unique actions such that $(h, a) < p$ and $(h, a') < p'$. Define $\varphi(s, 1) = (h, a)$ and $\varphi(s, 2) = (h, a')$.

Now define the function $f : \{1, 2\}^{\mathbb{N}} \rightarrow P^*$ by letting $f(x)$ be the unique play that extends the history $\varphi(x_0, \dots, x_m)$ for each $m \in \mathbb{N}$. Let D^* be the image of f .⁴ Thus, $D^* \subseteq P^* \subseteq D \subset P$.

Let H^* denote the set of histories that are consistent with D^* , i.e., there is a play $p \in D^*$ such that $h < p$. At every $h \in H^*$, let $A^*(h)$ denote the set of actions that keep play consistent with D^* , i.e., $A^*(h) = \{a \in A : (h, a) \in H^*\}$.

By construction, for every $h \in H^*$, it holds that

- $A^*(h)$ is either a singleton or it contains exactly two actions,
- there exists a history $\bar{h} \in H^*$ such that $\bar{h} \succeq h$ and $A^*(\bar{h})$ contains exactly two actions.

Define the game G as follows: The set of players is $\{1, 2\}$. The function $\iota : H \rightarrow \{1, 2\}$ is defined recursively as follows. Let $\iota(\emptyset) = 1$. Suppose that $\iota(h)$ has been defined for $h \in H$. If $h \in H^*$ and $A^*(h)$ is the singleton $\{a\}$, let $\iota(h, a) = \iota(h)$. In this case, we rename the action a into $\iota(h)$. If $h \in H^*$ and $A^*(h)$ consist of two actions a and a' , take the unique $s \in S$ such that $\varphi(s, 1) = (h, a)$ and $\varphi(s, 2) = (h, a')$. Let $\iota(h, a) = 1$ and $\iota(h, a') = 2$. In this case, rename a into 1 and a' into 2. If $h \in H \setminus H^*$, then $\iota(h, a)$ can be arbitrary, say equal to 1, for each $a \in A$.

⁴ It is well known that f is a homeomorphism of the Cantor set $\{1, 2\}^{\mathbb{N}}$ with D^* .

The payoffs on D^* are defined exactly as in the game of Flesch et al. (2014). It remains to define the payoffs on the complement of D^* . Take a play $p \in P \setminus D^*$ and let h be the shortest prefix of p at which the active player, say player i , chooses an action a such that $(h, a) \notin H^*$. In that case, regardless of future play, we define the payoff for player i to be -10 , and for the other player to be 5 .⁵ Notice that the payoffs are continuous outside of D^* , and hence, outside of D as well.

Thus, the game G proceeds just as the game in Flesch et al. (2014), unless an active player chooses an action outside of $A^*(h)$, in which that player is punished and the other player is rewarded.

Step 2: Proving that G has no ϵ -SPE for small $\epsilon > 0$. We let H_i^* be the set of histories in H^* where player i is active. Let S_1 denote the set of plays with the payoff $(-1, 2)$; S_2 the set with the payoff $(-2, 1)$; Q_1 the set with the payoff $(-10, 5)$; Q_2 denote the set of plays with the payoff $(5, -10)$, and R the set with the payoff $(0, 0)$.

Let $\sigma = (\sigma_1, \sigma_2)$ be an ϵ -SPE where $\epsilon \in (0, \frac{1}{7})$. We first argue that

$$\mu_{\sigma,h}(Q_1) \leq \frac{\epsilon}{9}, \quad \mu_{\sigma,h}(Q_2) \leq \frac{\epsilon}{9} \quad \text{for each } h \in H^*. \tag{5.1}$$

To prove the first inequality, consider the following strategy τ_1 for player 1: Follow σ_1 unless the outcome of the lottery $\sigma_1(h)$ at a history $h \in H_1^*$ is an action outside of $A^*(h)$. If that happens, play action 1 for the rest of the game. Recall playing an action outside $A^*(h)$ gives player 1 a payoff of -10 , whereas playing 1 forever gives -1 . Then, for each $h \in H^*$

$$\epsilon \geq u_1((\tau_1, \sigma_2), h) - u_1((\sigma_1, \sigma_2), h) \geq 9 \cdot \mu_{\sigma,h}(Q_1),$$

from which the first inequality follows. The proof of the second inequality is similar.

The rest of the proof follows very closely the analysis in Flesch et al. (2014).

We argue that

$$\mu_{\sigma,h}(R) = 0 \quad \text{for each } h \in H^*. \tag{5.2}$$

Otherwise by Lévy’s zero-one law, there is a history $h \in H^*$ where $\mu_{\sigma,h}(R) \geq \frac{19}{20}$, which implies that $u_2(\sigma, h) \leq 10 \cdot \frac{1}{20} < 1 - \epsilon$. This is impossible since at h player 2 can guarantee a payoff of 1 by always playing action 2 whenever it is his turn.

At any history $h \in H_1^*$, player 1 can guarantee a payoff of -1 by always playing action 1, hence

$$u_1(\sigma, h) \geq -1 - \epsilon \quad \text{for each } h \in H_1^*. \tag{5.3}$$

Now

$$u_1(\sigma, h) = \mu_{\sigma,h}(S_1) \cdot (-1) + \mu_{\sigma,h}(S_2) \cdot (-2) + \mu_{\sigma,h}(Q_1) \cdot (-10) + \mu_{\sigma,h}(Q_2) \cdot 5.$$

Combining this with (5.1), (5.2), and (5.3), we obtain $\mu_{\sigma,h}(S_1) \geq 1 - 2\epsilon$ for every $h \in H_1^*$. This last inequality together with (5.1) give $u_2(\sigma, h) \geq 2 - 6\epsilon$ for each

⁵ Here, one has to be careful, because two low payoffs such as $(-10, -10)$ does not work. This would namely allow player i to punish his opponent, which in some games could create ϵ -SPE.

$h \in H_1^*$. It is then easy to conclude that $u_2(\sigma, h) \geq 2 - 7\epsilon$ for each $h \in H^*$. Since $u_2(\sigma, h) \leq \mu_{\sigma,h}(S_2) + 10(1 - \mu_{\sigma,h}(S_2))$, we obtain

$$\mu_{\sigma,h}(S_2) \leq \frac{8 + 7\epsilon}{9} < 1 \quad \text{for each } h \in H^*.$$

Therefore, applying Lévy’s zero-one law yields $\mu_{\sigma,h}(S_2) = 0$ for each $h \in H^*$. Then, for each $h \in H^*$, we have $u_1(\sigma, h) \leq -\mu_{\sigma,h}(S_1) + 10\mu_{\sigma,h}(Q_2) \leq -1 + 4\epsilon < -\epsilon$. On the other hand, player 1 can guarantee a payoff of 0 at any history in H^* by playing action 2 whenever it is available. We arrive at a contradiction. \square

6 Concluding remarks

An interesting direction for future research is extending the results of this paper to the case where the set of players is infinite. Consider a game with infinitely many players.

First suppose that the payoff functions are all continuous. If the set of actions is finite, the existence of a pure SPE follows by the truncation approach of [Fudenberg and Levine \(1983\)](#). When the set of actions is infinite, an SPE need not exist, but one could follow the approach in [Flesch and Predtetchinski \(2015\)](#) to obtain existence of a pure ϵ -SPE, for all $\epsilon > 0$.

Much more challenging is the case where the continuity of payoff functions is only assumed outside of a sigma-discrete set. In this case, the question remains open whether an ϵ -SPE exists for all $\epsilon > 0$. We conjecture that the existence can be established at least in the special case when, on every play, infinitely many players are active. The intuition is that we can require a player who appears for the first time to randomize and place probability at least δ on at least two actions, where δ is a small positive number. This way the set of discontinuities can be avoided almost surely. A similar idea appeared in [Cingiz et al. \(2015\)](#).

Open Access This article is distributed under the terms of the Creative Commons Attribution License which permits any use, distribution, and reproduction in any medium, provided the original author(s) and the source are credited.

7 Appendix

7.1 Proof of lemma 3

Let E denote the set of histories $h \in H$ such that

$$\sup_{p,q \in P(h)} |f(p) - f(q)| \leq \frac{\epsilon}{2}.$$

We also let

$$P_E = \bigcup_{h \in E} P(h).$$

Notice that P_E is an open set. For each $p \in P_E$, let $h(p)$ be the minimal history such that $h \prec p$ and $h \in E$. Now fix an element $a \in A$ arbitrarily and define the function $\ell : P \rightarrow P$ as follows:

$$\ell(p) = \begin{cases} (h(p), a, a, \dots) & \text{if } p \in P_E \\ p & \text{otherwise.} \end{cases}$$

We argue that ℓ is a continuous function. Let p_k be a sequence of plays converging to a play $p \in P$. We must show that $\ell(p_k)$ converges to $\ell(p)$.

Suppose first that $p \in P_E$. Then, for k large enough, we have $h(p) \prec p_k$, hence $p_k \in P_E$ with $h(p_k) = h(p)$, and so $\ell(p_k) = \ell(p)$.

Suppose now that $p \notin P_E$. Then, $\ell(p) = p$. We have to show that for each $h \prec p$, there is a $K \in \mathbb{N}$ such that $h \prec \ell(p_k)$ for all $k \geq K$. Thus, take any $h \prec p$. There exists a $K \in \mathbb{N}$ such that $h \prec p_k$ for all $k \geq K$. Take any $k \geq K$. If $p_k \notin P_E$ then $\ell(p_k) = p_k$, and hence $h \prec \ell(p_k)$, as desired. Suppose $p_k \in P_E$. Since p_k extends both histories h and $h(p_k)$, we have either $h \preceq h(p_k)$ or $h(p_k) \prec h$. But if $h(p_k) \prec h$ then $h(p_k) \prec p$, implying that $p \in P_E$, contrary to our supposition. Hence $h \preceq h(p_k)$, which yields $h \prec \ell(p_k)$, as desired. Thus, we have shown that $\ell(p_k)$ indeed converges to $\ell(p)$.

Without loss of generality assume that r is a multiple of $\frac{\epsilon}{2}$. Let U be the set consisting of points $x \in [-r, r]$ that are multiples of $\frac{\epsilon}{2}$. Let $g : [-r, r] \rightarrow U$ be given by $g(x) = \min\{y \in U : x \leq y\}$. Notice that g is Borel and that $|g(x) - x| < \frac{\epsilon}{2}$ for every $x \in [-r, r]$.

Define the function \bar{f} by letting $\bar{f}(p) = g(f(\ell(p)))$. It is Borel since g and f are Borel and ℓ is continuous. Moreover, \bar{f} has finite range since g does, so property [1] of the lemma is satisfied.

To check property [2], take a $p \in P$. If $p \notin P_E$, then

$$|\bar{f}(p) - f(p)| = |g(f(p)) - f(p)| < \frac{\epsilon}{2}.$$

If $p \in P_E$, then

$$|\bar{f}(p) - f(p)| \leq |g(f(\ell(p))) - f(\ell(p))| + |f(\ell(p)) - f(p)| < \frac{\epsilon}{2} + \frac{\epsilon}{2} = \epsilon.$$

To check property [3], suppose that f is continuous at a play $p \in P$. Then, $p \in P_E$ and so $p \in P(h(p))$. Since ℓ is constant on the whole set $P(h(p))$, so is \bar{f} . Thus, \bar{f} is continuous on $P(h(p))$, and in particular at p . □

7.2 Proof of lemma 4

The proof is by the induction on the number of players. Obviously, every 1-player stopping game with finite range of payoffs has an SPE, which is an ϵ -SPE for each $\epsilon > 0$. So suppose that we have proven theorem 1 for all stopping games with finite range of payoffs with at most $n - 1$ players. Consider a stopping game G with n players with finite range of payoffs.

Given a $k \in \mathbb{N}$, we consider a “truncated” game G_k where at time k player i^k is forced to stop.⁶ By using backward induction, we can find a pure SPE σ_k in the game G_k with an additional property: for each $t < k$, whenever player i^t at history c^t is indifferent between taking action c and action s , the strategy profile σ_k assigns action s . That is, $\sigma_k(c^t) = s$ if $u_{i^t}(\sigma_k, c^{t+1}) = r_{i^t}^t$.

We identify the strategy profile σ_k with a point of $\{c, s\}^{\mathbb{N}}$. By taking a subsequence if necessary, we can assume that the sequence σ_k converges to a strategy profile σ . We consider the following three cases:

- [1] There is no player $i \in N$ for which the set $\{t \in \mathbb{N} : \sigma_i(c^t) = s\}$ is infinite.
- [2] There exist at least two players $i \in N$ for which the set $\{t \in \mathbb{N} : \sigma_i(c^t) = s\}$ is infinite.
- [3] There exists exactly one player $i \in N$ for which the set $\{t \in \mathbb{N} : \sigma_i(c^t) = s\}$ is infinite.

We prove that case 1 is impossible, and that in case 2, the strategy profile σ is an SPE. In case 3, we modify σ to construct an ϵ -SPE.

Case 1: We show that this case cannot occur. Let $T \in \mathbb{N}$ be so large that $\sigma(c^t) = s$ for every $t \geq T$. By replacing the game G with the subgame starting with history c^T , we may assume that in fact $\sigma(c^t) = c$ for all $t \in \mathbb{N}$.

Let ℓ_k denote the first time that, under the strategy profile σ_k , action s is taken. That is $\ell_k = \min\{t \in \mathbb{N} : \sigma_k(c^t) = s\}$. Notice that $\ell_k \leq k$, since at history c^k , player i^k plays s by the definition of G_k . We argue that the sequence $\{\ell_k\}$ is bounded.

Given a player $i \in N$, consider the set $R_i = \{r_i^t : t \in \mathbb{N}, i^t = i\}$ and for each $k \in \mathbb{N}$ let $R_i^k = \{r_i^t : t \in \mathbb{N}, i^t = i, t \leq k\}$. The sets R_i and R_i^k are finite. Clearly $R_i^0 \subseteq R_i^1 \subseteq \dots$ is an increasing sequence of sets with the limit R_i , so there exists a k_i such that $R_i^{k_i} = R_i$. Now let $K = \max\{k_1, \dots, k_n\}$. We argue that $\ell_k \leq K$ for each $k \geq K$. Suppose on the contrary that $\ell_k > K$ for some $k \geq K$. Let $i = i^{\ell_k}$. By the choice of K , there exists a $t \leq K$ such that $r_i^t = r_i^{\ell_k}$. Since, under the strategy profile σ_k , no player stops before period ℓ_k , at history c^t player i is indifferent between taking action s and action c . Hence, our requirement on the strategy profile σ_k implies that $\sigma_k(c^t) = s$. But this contradicts the definition of ℓ_k .

Thus, we have shown that the sequence $\{\ell_k\}$ is bounded. By taking a subsequence if necessary, we may then assume that it converges, say to an $l \in \mathbb{N}$. However, then $\sigma(c^l) = s$, contradicting our assumption that $\sigma(c^t) = c$ for all $t \in \mathbb{N}$.

Case 2: We prove that σ is an SPE.

Take a $t \in \mathbb{N}$. Let $m > t$ be the minimal number such that $\sigma(c^m) = s$ and $i^m \neq i^t$. Notice that such an m exists by the condition defining case 2. Let $k \in \mathbb{N}$ be so large that σ_k coincides with σ on histories c^0, \dots, c^m . Then, for any strategy η_{i^t} of player i^t , we have $u_{i^t}(\sigma_k/\eta_{i^t}, c^t) = u_{i^t}(\sigma/\eta_{i^t}, c^t)$. Since η_{i^t} is not a profitable deviation from σ_k , it is not a profitable deviation from σ .

⁶ Formally, one could define the payoffs in the game G_k such that any play extending the history c^{k+1} gives player i^k a payoff smaller than $r_{i^k}^k$.

Case 3: Let i be the only player such that the set $S = \{t \in \mathbb{N} : \sigma_i(c^t) = s\}$ is infinite. One can show exactly as in case 2 that no player $j \neq i$ has a profitable deviation from the strategy profile σ .

Let $t_0 < t_1 < t_2 \dots$ be the increasing enumeration of the elements of S . It is easy to see that for each m

$$r_i^{t_m} = \max\{r_i^\ell : \ell \geq t_m, \ell \neq i\}. \tag{7.1}$$

In particular, $r_i^{t_0} \geq r_i^{t_1} \geq \dots$ is a non-increasing sequence. Define α to be the limit of this sequence. Define also

$$\beta = \max_{T \in \mathbb{N}} \min\{r_i^t : t \geq T, t \neq i\}.$$

Intuitively, β is the minimal payoff in far subgames to player i that he receives when other players decide to play action s . This payoff β will serve as a threat payoff against player i . We distinguish the following three cases:

- [3.1] $\alpha \geq r_i^\infty$,
- [3.2] $\alpha < r_i^\infty$ and $\alpha < \beta$,
- [3.3] $\alpha < r_i^\infty$ and $\beta \leq \alpha$.

Case 3.1: In this case, the strategy profile σ is an SPE of the game G . Indeed, the fact that player i has no profitable deviations from σ follows since $r_i^{t_0} \geq r_i^{t_1} \geq \dots \geq r_i^\infty$. Since, as noticed above, other players have no profitable deviations, σ is an SPE.

Case 3.2: Choose an $m, \ell \in \mathbb{N}$ so large that $r_i^{t_m} = \alpha$ and $\min\{r_i^t : t \geq \ell, t \neq i\} = \beta$. Let $k = \max\{t_m, \ell\}$. We argue that the subgame $G(c^k)$, that starts at history c^k , has an ϵ -SPE. Then, due to backward induction, G admits an ϵ -SPE as well.

Notice that, in $G(c^k)$, player i has no incentive to ever stop. Indeed, stopping the game yields player i a payoff of at most α . If another player stops the game, player i 's payoff is at least β . And if no one stops, player i 's payoff is r_i^∞ .

Let η_i be the strategy for player i in $G(c^k)$ that requires player i to always continue. Fixing player i 's strategy to η_i yields a new stopping game where the player set is $N \setminus \{i\}$. In view of our inductive hypothesis, this game has an ϵ -SPE η_{-i} . It is now not difficult to see that $\eta = (\eta_i, \eta_{-i})$ is an ϵ -SPE in the game $G(c^k)$.

Case 3.3: We construct an ϵ -SPE of G .

Since $\beta \leq \alpha$, there are infinitely many periods $t \in \mathbb{N}$ such that $t \neq i$ and $r_i^t \leq \alpha$. Consequently, we can construct an infinite increasing sequence t'_0, t'_1, \dots such that

- [A] for each $\ell \in \mathbb{N}$, we have $t'_\ell \neq i$ and $r_i^{t'_\ell} \leq \alpha$, and
- [B] for each $m \in \mathbb{N}$, there is at most one $\ell \in \mathbb{N}$ such that $t_m < t'_\ell < t_{m+1}$.

Let M the maximum of $|r_j^t|$ over all $j \in N$ and $t \in \mathbb{N} \cup \{\infty\}$. If $M < \frac{\epsilon}{2}$, then any strategy is an ϵ -SPE of G , so there is nothing to prove. Suppose that $M \geq \frac{\epsilon}{2}$ and let $\delta = \frac{\epsilon}{2M}$. Define the strategy profile η by letting

$$\eta(t) = \begin{cases} \delta s + (1 - \delta)c & \text{if } t \in \{t'_0, t'_1, \dots\} \\ \sigma(t) & \text{otherwise.} \end{cases}$$

In particular, $\eta_i = \sigma_i$. We argue that η is an ϵ -SPE.

Consider a player $j \neq i$. Take a $t \in \mathbb{N}$ with $t^t = j$. We show that for each strategy η'_j of player j , it holds that

$$|u_j(\eta/\eta'_j, t) - u_j(\sigma/\eta'_j, t)| \leq \epsilon. \tag{7.2}$$

It suffices to verify (7.2) for all pure strategies η'_j of player j . Thus, let m be the minimal natural number such that $t < t_m$. By property [B], there is at most one ℓ for which $t \leq t'_\ell < t_m$. Consider now the play of the game starting at period t :

- If under the strategy profile σ/η'_j player j stops the game at or before period t'_ℓ , the same occurs under η/η'_j .
- If under the strategy profile σ/η'_j player j stops the game at period $s > t'_\ell$ (in which case $s < t_m$), then under η/η'_j the game stops at period t'_ℓ with probability δ and at period s with probability $1 - \delta$.
- Suppose that under the strategy profile σ/η'_j player i stops the game at time t_m . If $t^{t'} = j$, then under η/η'_j the game stops at t_m with probability 1. If $t^{t'} \neq j$, then the game stops at with t'_ℓ with probability δ and stops at t_m with probability $1 - \delta$.

In all three cases, the difference in the payoff to player j is at most ϵ .

Since for player j no deviation from the strategy profile σ is profitable, no deviation from the strategy profile η increases player j 's payoff by more than ϵ .

Consider player i . Take a $t \in \mathbb{N}$ with $t^t = i$. Let m be the least natural number with $t \leq t_m$. It is not difficult to see that the payoff to player i on the strategy profile η is at least $r_i^{t_m} - \epsilon$.

Now let η'_i be a pure strategy for player i . Under the strategy profile η/η'_i with probability 1 one of the following outcomes prevails: Either the game stops at some period in the set $\{t'_0, t'_1, \dots\}$, or the game is stopped by player i at a period $t'' \geq t$. In the first case, the payoff to player i is bounded above by α by property [A] above. Notice that $\alpha \leq r_i^{t_m}$. In the second case, the payoff is $r_i^{t''}$. Notice that $r_i^{t''} \leq r_i^{t_m}$ by (7.1). We conclude that under the strategy profile η/η'_i the payoff to player i is at most $r_i^{t_m}$. Hence, the deviation to η'_i improves player i 's payoff by at most ϵ .

This concludes the proof of theorem 1 for stopping games with finite payoff range. We remark that the method to solve case 3.3 is exactly the one that is needed to find an ϵ -SPE in the above-mentioned game by Solan and Vieille.

7.3 sigma-discrete subsets of P are Borel

Lemma 7 *A sigma-discrete subset of P is Borel. Furthermore, every function $f : P \rightarrow \mathbb{R}$ that is continuous outside a sigma-discrete subset of P is Borel measurable.*

Proof

Step 1: Let $Q \subset P$ be a discrete set of plays. We argue that Q is a G_δ -subset of P , i.e., a Π_2^0 -set of P . For every $q \in Q$, let

$$m(q) = \inf_{q' \in Q \setminus \{q\}} d(q, q'),$$

where d is the metric on P . Because Q is discrete, $m(q) > 0$. Now for every $n \in \mathbb{N}$, consider the set

$$U_n = \bigcup_{q \in Q} B\left(q, \frac{1}{2^n}m(q)\right),$$

where $B(q, \varepsilon)$ denotes the open ball around a play q of radius ε . Notice that the sequence U_n is non-increasing.

We show that U_1 is a disjoint union. Take a $p \in P$ and suppose that there are two different $q_1, q_2 \in Q$ such that $p \in B(q_1, \frac{1}{2}m(q_1))$ and $p \in B(q_2, \frac{1}{2}m(q_2))$. Assume that $m(q_1) \geq m(q_2)$. Then,

$$d(q_1, q_2) \leq d(q_1, p) + d(p, q_2) < \frac{1}{2}m(q_1) + \frac{1}{2}m(q_2) \leq m(q_1),$$

which is in contradiction with the choice of $m(q_1)$.

It follows that

$$Q = \bigcap_{n \in \mathbb{N}} U_n.$$

Indeed, the inclusion \subseteq is trivial. To see the other inclusion \supseteq , take a $p \in \bigcap_{n \in \mathbb{N}} U_n$. Suppose by way of contradiction that $p \notin Q$. Then, $p \in U_1 \setminus Q$. Since U_1 is a disjoint union, there is a unique $q \in Q$ such that $p \in B(q, \frac{1}{2}m(q))$. Because $d(p, q) > 0$, we conclude that $p \notin B(q, \frac{1}{2^n}m(q))$ for large n , and hence, $p \notin U_n$ for large n . This is a contradiction with the choice of p .

Since each U_n is open, Q is a G_δ -set of P , as desired.

Step 2: Let $Q \subset P$ be a sigma-discrete set of plays. Then, Q is a countable union of G_δ -sets of P , and hence, Q is a Borel subset of P , in particular it is a Σ_3^0 -set of P .

Step 3: Let $f : P \rightarrow \mathbb{R}$ be continuous outside a sigma-discrete subset $D \subset P$. We argue that f is Borel. Take a Borel set X of \mathbb{R} . Let f_C and f_D denote the restriction of f to the domain $P \setminus D$ and, respectively, to the domain D .

Since f_C is continuous with respect to the relative topology on $P \setminus D$, the set $f_C^{-1}(X)$ is a Borel subset of $P \setminus D$. As $P \setminus D$ is a Borel subset of P by step 2, $f_C^{-1}(X)$ is a Borel subset of P . Further, since any subset of D is sigma-discrete, by step 2 $f_D^{-1}(X)$ is a Borel subset of P . Hence,

$$f^{-1}(X) = f_C^{-1}(X) \cup f_D^{-1}(X)$$

is a Borel subset of P , as desired. □

References

Carmona, G.: On games of perfect information: equilibria, ϵ -equilibria and approximation by simple games. *Int. Game Theor. Rev.* 7, 491–499 (2005)
 Cingiz, K., Flesch, J., Herings, P.J.J., Predtetchinski, A.: Perfect information games with infinitely many players. Working paper (2015)

- Flesch, J., Kuipers, J., Mashiah-Yaakovi, A., Schoenmakers, G., Solan, E., Vrieze, K.: Perfect-information games with lower semicontinuous payoffs. *Math. Oper. Res.* **35**, 742–755 (2010)
- Flesch, J., Kuipers, J., Mashiah-Yaakovi, A., Shmaya, E., Schoenmakers, G., Solan, E., Vrieze, K.: Equilibrium refinements in perfect information games with infinite horizon. *Int. J. Game Theor.* **43**, 945–951 (2014)
- Flesch, J., Predtetchinski, A.: On refinements of subgame perfect epsilon-equilibrium. *Int. J. Game Theor.* (2015, accepted)
- Fudenberg, D., Levine, D.: Subgame-perfect equilibria of finite and infinite-horizon games. *J. Econ. Theor.* **31**, 251–268 (1983)
- Kechris, A.S.: *Classical Descriptive Set Theory*. Springer, Berlin (1995)
- Koumoullis, G.: Cantor sets in Prohorov spaces. *Fundam. Math.* **124**, 155–161 (1984)
- Mashiah-Yaakovi, A.: Subgame perfect equilibria in stopping games. *Int. J. Game Theor.* **43**, 89–135 (2014)
- Purves, R.A., Sudderth, W.D.: Perfect information games with upper semicontinuous payoffs. *Math. Oper. Res.* **36**, 468–473 (2011)
- Solan, E., Vieille, N.: Deterministic multi-player Dynkin games. *J. Math. Econ.* **39**, 911–929 (2003)