

A Characterization of Subgame-Perfect Equilibrium Plays in Borel Games of Perfect Information

Citation for published version (APA):

Flesch, J., & Predtetchinski, A. (2017). A Characterization of Subgame-Perfect Equilibrium Plays in Borel Games of Perfect Information. *Mathematics of Operations Research*, 42(4), 1162-1179.
<https://doi.org/10.1287/moor.2016.0843>

Document status and date:

Published: 01/11/2017

DOI:

[10.1287/moor.2016.0843](https://doi.org/10.1287/moor.2016.0843)

Document Version:

Publisher's PDF, also known as Version of record

Document license:

Taverne

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

A Characterization of Subgame-Perfect Equilibrium Plays in Borel Games of Perfect Information

János Flesch,^a Arkadi Predtetchinski^b

^a Department of Quantitative Economics, Maastricht University, 6200 MD Maastricht, Netherlands; ^b Department of Economics, Maastricht University, 6200 MD Maastricht, Netherlands

Contact: j.flesch@maastrichtuniversity.nl (JF); a.predtetchinski@maastrichtuniversity.nl (AP)

Received: March 4, 2016

Revised: October 14, 2016

Accepted: November 10, 2016

Published Online in Articles in Advance:
May 25, 2017

MSC2010 Subject Classification: Primary:
91A25 dynamic games

OR/MS Subject Classification: Primary:
games/group decisions, noncooperative

<https://doi.org/10.1287/moor.2016.0843>

Copyright: © 2017 INFORMS

Abstract. We provide a characterization of subgame-perfect equilibrium plays in a class of perfect information games where each player's payoff function is Borel measurable and has finite range. The set of subgame-perfect equilibrium plays is obtained through a process of iterative elimination of plays. Extensions to games with bounded Borel measurable payoff functions are discussed. As an application of our results, we show that if every player's payoff function is bounded and upper semicontinuous, then, for every positive epsilon, the game admits a subgame-perfect epsilon-equilibrium. As we do not assume that the number of players is finite, this result generalizes the corresponding result of Purves and Sudderth [24] [Purves RA, Sudderth WD (2011) Perfect information games with upper semicontinuous payoffs. *Math. Oper. Res.* 36(3):468–473].

Keywords: perfect information games • subgame-perfect equilibrium • semicontinuity

1. Introduction

The setup. Considered are perfect information games played over a game tree of infinite depth. A single active player is assigned to every history in the tree. The play begins at the empty history, called the root of the tree. Once a given history is reached, the corresponding active player chooses an action leading to one of the successor histories. This generates an infinite sequence of actions, called a play, which determines the payoffs.

We allow for an arbitrary set of players. Without loss of generality we require that each history has at least one successor, and we impose no further restrictions on the game tree. We assume that each player's payoff function is bounded and Borel measurable with respect to the topology generated by the cylinder sets, as e.g., in Martin [20]. This is a very mild assumption. Games with semicontinuous payoffs (Purves and Sudderth [24], Flesch et al. [7]), perfect information stopping games (e.g., Solan [25], Mashiah–Yaakovi [21]), recursive stochastic games with perfect information (e.g., Flesch et al. [8], Kuipers et al. [16]) are all examples of games with Borel measurable payoffs.

Parts of the paper are written under an additional assumption that each player's payoff function only takes finitely many distinct values. The main reason to focus on payoff functions with finite range is that they serve as natural approximations of bounded payoff functions.

This paper concentrates on the concept of subgame-perfect (ϵ -)equilibrium in pure strategies. The main contribution of the paper is twofold. We discuss each contribution, and the additional results, in turn.

Contribution I. Our first contribution is a characterization of subgame-perfect equilibrium plays, in games where the payoff functions are Borel measurable and have finite range. The stated characterization is obtained using an algorithm of iterative elimination of plays.

The algorithm proceeds as follows. At each step of the algorithm and for each history of the game tree we determine a certain payoff level, called the security level, for the player who is active at the given history. This is the highest payoff level that the active player can be sure to attain when playing against any strategy profile that only prescribes plays that have survived the elimination thus far. We then eliminate all plays with a payoff lower than the active player's security level. The algorithm terminates when no more plays can be eliminated. The plays that remain are exactly those that are supported by some subgame-perfect equilibrium of the game. If no plays remain, the game admits no subgame-perfect equilibrium. In general, the algorithm takes a transfinite number of steps to terminate.

The main technical novelty of our work is in the definition of the security payoff level. The definition relies on an auxiliary zero-sum game played by the deviator (Player I) against the strategy builder (Player II). Player I effectively represents the active player at the given history in the original game, while Player II acts as his artificial adversary.

The game begins with Player I choosing an action. Player II then proposes a play that respects Player I's move. Player I then chooses to accept the proposed play or to deviate. In the former case, the game ends. In the latter case, Player I also announces a deviation, and Player II proposes a new play that accommodates Player I's deviation. This process continues ad infinitum. Player II is restricted to propose plays that have survived the elimination thus far. The value of this auxiliary zero-sum game is defined to be the active player's security level.

We believe that the characterization established here could yield useful insights into the nature of subgame perfection. Our second contribution on upper semicontinuous games, obtained as an application of this characterization, serves as an illustration.

Contribution II. As an application of the characterization result, we address an open question raised by Purves and Sudderth [24]. Purves and Sudderth work in a setup that is identical to ours, except for their assumption that the number of players is finite. They show that if each player's payoff function is bounded and upper semicontinuous, then for each positive ϵ , the game admits a subgame-perfect ϵ -equilibrium. The authors further pose the question of whether the assumption that the number of players be finite could be dispensed with.

The technique that Purves and Sudderth use is an induction on the number of distinct payoff vectors in the game. Because of this feature of their proof, extending it beyond the case of finitely many players seems to be difficult. We take a different approach. We show that in games where each player's payoff function has a finite range and is upper semicontinuous, our algorithm returns a nonempty set of plays at each step of the iteration. This allows us to deduce that any such game has a subgame-perfect equilibrium. The result for games with bounded upper semicontinuous payoff functions then follows by discretizing the payoff functions, exactly as is done by Purves and Sudderth.

Additional results. For games with bounded Borel measurable payoffs we obtain the following result: subgame-perfect equilibrium plays survive all stages of the elimination. Conversely, a play that survives all stages of elimination could be supported as a subgame-perfect ϵ -equilibrium play for each positive ϵ . We show by means of an example that this result cannot be strengthened: there might be plays that survive all stages of elimination and yet that are not supportable by any subgame-perfect equilibrium. We argue that such a discrepancy arises because Player II might not have an optimal strategy in the auxiliary zero-sum game.

Our final result is motivated by the frequent use of the technique of discretizing payoff functions: approximating a given bounded payoff function by a function with finite range, as is done, for instance, by Mertens and Neyman (see Mertens [22]), Flesch et al. [7], Purves and Sudderth [24], Flesch and Predtetchinski [10, 11]. The technique allows one to reduce the problem of existence of (subgame-perfect) ϵ -equilibrium for a game with bounded payoff functions to the problem of existence of (subgame-perfect) 0-equilibrium for a game where the payoff functions have finite range. We show that a game with bounded Borel measurable payoff functions admits a subgame-perfect ϵ -equilibrium for each $\epsilon > 0$ if and only if, for each $\epsilon > 0$, there exists a Borel measurable ϵ -discretization of its payoff functions that admits a subgame-perfect 0-equilibrium. Our result could be interpreted as saying that passing to discretizations is without loss of generality, at least if one is only interested in the question of existence of subgame-perfect ϵ -equilibrium for every positive ϵ .

Related literature. Closely related to the algorithm developed here are recursive algorithms in Harris [14] and in Flesch et al. [7]. Both algorithms terminate with the set of subgame-perfect equilibrium plays. Harris considers the class of continuous games, whereas Flesch et al. assume that the payoff functions are lower semicontinuous. For general payoff functions having finite range, the algorithm in Flesch et al. outputs the set of plays that could be supported by a one-deviation immune strategy profile, that is, a strategy profile such that no player could improve his payoff by deviating at a single history only.

Our algorithm shares a number of features with those in Harris and Flesch et al.: all three algorithms prescribe iterative elimination of plays. At each step of the algorithm one determines a certain payoff level that a player could, in some sense, guarantee for himself. What distinguishes the algorithm presented here is the way the security payoff level is determined.

Harris and Flesch et al. only take into account deviations at a single history. More precisely, each of the player's actions is evaluated, assuming that the ensuing play will be the worst possible among those plays that survived the elimination thus far. Equivalently, one could see this as a value of a "short" zero-sum game, where the deviator is only allowed to deviate once, and is forced to accept the first proposal of the strategy builder. In contrast, the "full" auxiliary zero-sum game employed here allows the deviator to deviate finitely or even infinitely many times. Unlike the algorithms of Harris and Flesch et al., the algorithm proposed here does not require the game to satisfy the one-shot deviation principle: indeed, the security payoffs capture the possibility of infinite as well as finite deviations.

Flesch et al. [12] consider the class of free transition games and proves the existence of subgame-perfect ϵ -equilibrium in mixed strategies using an algorithm of iterative elimination of plays. This algorithm operates

with pairs of plays, rather than individual plays. In the construction of subgame-perfect ϵ -equilibria, the primary play is to be followed with a high probability, and the secondary play, which acts as a threat, is to be followed with a small probability. Kuipers et al. [16] refine this algorithm to establish the result for perfect information stochastic games with the crucial feature that every player controls only one state. Alós-Ferrer and Ritzberger [1] consider an algorithm very different from ours to obtain an existence of subgame-perfect equilibrium in a class of continuous games.

Fudenberg and Levine [13] and Harris [15] develop a characterization of subgame-perfect equilibrium strategies (rather than plays) by means of convergent sequences of subgame-perfect ϵ -equilibria in the truncated games. Carmona [5] provides a characterization in terms of approximations of the game by payoff functions with finite range. Alós-Ferrer and Ritzberger [2] derive a condition on the topology of the set of plays that is both necessary and sufficient for the existence of a subgame-perfect equilibrium.

Brihaye et al. [4] obtain a characterization of one-deviation immune strategy profiles (called very weak subgame-perfect equilibria there) for quantitative games played on finite graphs, using an algorithm motivated by the one in Flesch et al. (2010a).

Le Roux [18] provides an intriguing analysis of games with Δ_2^0 -payoff functions. Essentially, he shows that, in the two-player case, the game by Solan and Vieille [26] is the only payoff pattern that could rule out the existence of subgame-perfect equilibrium.

Le Roux and Pauly [19] derive a number of existence results for subgame-perfect ϵ -equilibrium with infinitely many players. These results are transfer results, i.e., starting with the existence of some type of equilibria for a small class of games, they allow one to conclude the existence of some type of equilibria for a larger class. Our result on discretization could also be seen as such a transfer result.

T. Brihaye, V. Bruyère, N. Meunier, and J.-F. Raskin (reported in Meunier [23]) provide a characterization of games that admit a subgame-perfect equilibrium. With each perfect information game, they associate a two-player zero-sum game, called the prover game. They show that the original perfect information game has a subgame-perfect equilibrium if and only if player I has a winning strategy in the corresponding prover game. They also give a similar characterization for one-deviation immune strategy profiles (called very weak subgame-perfect equilibria there).

As we allow the player set to be infinite, our setup encompasses some classes of intergenerational games: perfect information games where a player represents a generation and thus acts exactly once in the course of the game. For a thorough review of such games, we refer to Balbus et al. [3]. In a related work, Cingiz et al. [6] consider infinite centipede games where each player acts exactly once. Using truncations of the game tree, the authors show the existence of subgame-perfect equilibrium under the assumption that the payoff functions are upper semicontinuous.

The organization of the paper. In Section 2 we describe the class of perfect information games. In Section 3 we provide an informal overview of the algorithm, describe a formal definition, and discuss two illustrating examples. Section 4 provides an overview of the main results. Proofs of the main results are carried out in Sections 5 and 6. Section 7 is dedicated to the discussion of interesting special classes of games. Section 8 contains additional results.

2. The Model and Preliminaries

Finite and infinite sequences. Let $\mathbb{N} = \{0, 1, 2, \dots\}$ and let A be a nonempty set. We denote by $A^{<\mathbb{N}}$ the set of all finite sequences of elements of A , including the empty sequence $\emptyset$. A finite sequence of length $t + 1$, where $t \in \mathbb{N}$, is denoted by $(a_0, \dots, a_t)$. However, with a slight abuse of notation, we write a to denote the sequence (a) of length 1. For $h \in H$ we write (h, a) to denote the concatenation of h by a . We let $A^{\mathbb{N}}$ denote the set of all infinite sequences of elements of A . A typical element of $A^{\mathbb{N}}$ is denoted $p = (a_0, a_1, \dots)$.

For a finite sequence $h = (a_0, \dots, a_t) \in A^{<\mathbb{N}}$, sequences $\emptyset, (a_0), (a_0, a_1), \dots, (a_0, \dots, a_t)$, are called prefixes of h . By definition, the only prefix of $\emptyset$ is $\emptyset$. We write $g \leq h$ to mean that g is a prefix of a finite sequence h , and we write $g < h$ if g is a prefix of h and $g \neq h$. For an infinite sequence $p = (a_0, a_1, \dots) \in A^{\mathbb{N}}$, finite sequences $\emptyset, (a_0), (a_0, a_1), (a_0, a_1, a_2), \dots$, are called prefixes of p . We write $h < p$ to mean that h is a prefix of p .

Given $h, e \in A^{<\mathbb{N}}$ and $p \in A^{\mathbb{N}}$ we let (h, e) and (h, p) denote concatenations defined in an obvious way.

The game Ω . Let A be an arbitrary nonempty set. Elements of A are called actions. Let H be a subset of $A^{<\mathbb{N}}$ such that (1) $\emptyset \in H$, (2) if $h \in H$, then every prefix of h is an element of H , and (3) if $h \in H$, then there is $a \in A$ such that $(h, a) \in H$. Elements of H are called histories. A play p is an element of $A^{\mathbb{N}}$ such that every prefix of p is an element of H . The set of all plays is denoted¹ by P .

Let I denote an arbitrary nonempty set of players. Let $\iota: H \rightarrow I$ be a function that assigns an active player to each history. Further, each player $i \in I$ is given a so-called payoff function $u_i: P \rightarrow \mathbb{R}$. Thus a game Ω is characterized by the sets I, A, H , and the functions ι and u_i for each $i \in I$.

The game Ω is played as follows: let $h_0 = \emptyset$. Suppose a history $h_t \in H$ has been determined for some $t \in \mathbb{N}$. Player $i(h_t)$ then chooses an action a_t such that $(h_t, a_t) \in H$. The players thus generate a play $p = (a_0, a_1, \dots)$, and each player $i \in I$ receives payoff $u_i(p)$.

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, the set $A^{\mathbb{N}}$ with the product topology, and the set P with the topology as a subspace of $A^{\mathbb{N}}$. Denote this topology by $\mathcal{T}$. A basis for $(P, \mathcal{T})$ is formed by the cylinder sets $P(h) = \{p \in P: h < p\}$ for $h \in H$.

The space $(P, \mathcal{T})$ is completely metrizable. One can take, for example, the metric $d: P \times P \rightarrow \mathbb{R}$ given by $d(p, q) = 2^{-k(p, q)}$ if $p \neq q$, where $k(p, q) \in \mathbb{N}$ is the length of the longest common prefix of p and q . 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 plays p_n and p have a common prefix of length k . The space $(P, \mathcal{T})$ is separable (and hence Polish) if and only if H is a countable set.² This is clearly equivalent to the condition that for every $h \in H$ there are only countably many $a \in A$ such that $(h, a) \in H$.

A function $f: P \rightarrow \mathbb{R}$ is said to be *upper semicontinuous* if for each $x \in \mathbb{R}$, the set $\{p \in P: f(p) \geq x\}$ is closed. It is easy to see that the following conditions are equivalent: (A) f is upper semicontinuous, (B) for every sequence of plays $(p_n)_{n \in \mathbb{N}}$ converging to a play p , we have $\limsup_{n \rightarrow \infty} f(p_n) \leq f(p)$, and (C) for every $p \in P$ and every $\delta > 0$ there is a prefix h of p such that $f(p) \geq f(q) - \delta$ for every $q \in P(h)$. Intuitively, condition (C) means that any play p has a prefix h such that p is maximal, or at least almost maximal, within $P(h)$.

A function $f: P \rightarrow \mathbb{R}$ is said to be *lower semicontinuous* if the function $-f$ is upper semicontinuous. A function f is continuous if and only if it is both upper and lower semicontinuous.

Strategies. In this paper, we only consider pure strategies, unless specified otherwise. A (pure) strategy for player i is a function σ_i assigning to each history $h \in H$ with $i(h) = i$ an element $\sigma_i(h)$ of A such that $(h, \sigma_i(h)) \in H$. The interpretation is that if a history h as above arises, then σ_i tells player i to choose the action $\sigma_i(h)$. A strategy profile is a tuple $\sigma = (\sigma_i)_{i \in I}$ where σ_i is a strategy for player i . Given a strategy profile σ 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 .

A strategy profile σ induces a unique play at a history $h \in H$, denoted $\pi(\sigma, h)$. It is an element of $P(h)$. In particular, the induced play at the root is $\pi(\sigma, \emptyset)$.

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 $i \in I$ and for each strategy σ'_i of player i it holds that

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

Mertens and Neyman (see Mertens [22]) showed that an ϵ -equilibrium exists for every $\epsilon > 0$, provided that the payoff functions are bounded and Borel measurable. Their proof is for finitely many players, but it can be extended to infinitely many players along the same lines (see also Le Roux and Pauly [19]).

Subgame-perfect ϵ -equilibrium is a refinement of ϵ -equilibrium. A strategy profile σ is called a *subgame-perfect ϵ -equilibrium*, or ϵ -SPE for brevity, if no player can gain more than ϵ by a unilateral deviation conditional on any history. Formally σ is an ϵ -SPE if for each history $h \in H$, each player $i \in I$, and each strategy σ'_i of player i it holds that

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

A 0-SPE is simply called an SPE. For every $h \in H$ define

$$P^*(h) = \{\pi(\sigma, h): \sigma \text{ is an SPE in } \Omega\}. \quad (1)$$

The set $P^*(h)$ consists of all plays that are induced by an SPE of the game Ω at history h . The set $P^*(h)$ is a subset of $P(h)$. We refer to elements of the set $P^*(\emptyset)$ as SPE plays.

Flesch et al. [9] provided a two-player game in which each player's payoff function is bounded and Borel measurable, yet the game admits no ϵ -SPE for small $\epsilon > 0$, not even in randomized strategies.

Zero-sum games. We make use of the following lemma due to Mertens and Neyman (see Mertens [22], p. 1567). The result ultimately rests on Borel determinacy by Martin [20].

Lemma 2.1. *Let Ω be a two-player zero-sum perfect information game having a bounded and Borel measurable payoff function. Then the game Ω has a value. Moreover, for each $\epsilon > 0$, each player has a strategy that is ϵ -optimal in every subgame (and, consequently, the game has an ϵ -SPE). If, in addition, the payoff function has finite range, then each player has a strategy that is optimal in every subgame (and, consequently, the game Ω has an SPE).*

3. The Algorithm

In this section we discuss the algorithm of iterative elimination of plays that outputs the collection of sets $P^*(h)$ for $h \in H$. We split this section into three sections. In Section 3.1, which is not intended to be rigorous, we give a nontechnical overview and a motivation for the algorithm. In Section 3.2 we give a formal definition of the algorithm. In Section 3.3 we illustrate the algorithm by means of two examples.

Throughout this section, we assume that the payoff function of each player has finite range and is Borel measurable.

3.1. An Overview

We recursively define, for each history h , the sequence $P_\xi(h)$ of sets, and sequences $\beta_\xi(h)$ and $\alpha_\xi(h)$ of payoffs, indexed by an ordinal number ξ . The set $P_\xi(h)$ contains the plays that have survived the elimination up to step ξ of the algorithm. Naturally, the sequence $P_\xi(h)$ is nested: as the algorithm proceeds, fewer and fewer plays remain. The number $\alpha_\xi(h)$ is the security level for the player active at the history h . This is the payoff level that the active player can be sure to guarantee for himself when playing against any strategy profile that only prescribes plays that have survived the elimination process up to stage ξ . The number $\beta_\xi(h)$, which is only interesting when ξ is a limit ordinal, is the maximum of the security payoff levels $\alpha_\lambda(h)$ over the ordinals λ smaller than ξ .

Initialization. We initialize the recursion by letting $P_0(h)$ be the set of all plays having h as a prefix, and letting $\beta_0(h) = \alpha_0(h)$ be the worst payoff for the player $\iota(h)$ in the subgame starting at the history h .

Step 1. At step 1 of the algorithm, we set $\beta_1(h) = \alpha_0(h)$ and $P_1(h) = P_0(h)$, and we let the security level $\alpha_1(h)$ be the value of the following zero-sum game, denoted $G_1(h)$.

Player I of the game $G_1(h)$ is called the deviator, and Player II is called the strategy builder. Player I effectively represents player $\iota(h)$ in the original game, while Player II acts as his artificial opponent. The game starts with Player I taking an action a_0 at the history $h_0 = h$. Player II responds by proposing a play p_0 extending the prefix (h_0, a_0) . Player I now has a choice of either to accept p_0 , or to deviate. If Player I accepts, the game ends. If Player I chooses to deviate, he announces a prefix h_1 of p_0 extending (h_0, a_0) such that $\iota(h_1) = \iota(h_0)$, and an action a_1 that is incompatible with p_0 . Player II responds by proposing a new play, p_1 , that extends the history (h_1, a_1) . This process continues ad infinitum. In this way, the players generate the outcome play q : if Player I accepts Player II's proposal p_k , then $q = p_k$, while if Player I never accepts any proposal, then q is defined to be the play extending the histories $h_0, h_1, \dots$. Player I's payoff is given by $u_{\iota(h)}(q)$.

Thus $\alpha_1(h)$ is the highest payoff player $\iota(h)$ can be sure to attain when playing against an arbitrary strategy profile of his opponents. It is clear that in an SPE, player $\iota(h)$'s payoff cannot be smaller than $\alpha_1(h)$. Indeed, if it were smaller than $\alpha_1(h)$, player $\iota(h)$ could “beat” the assumed equilibrium and gain at least $\alpha_1(h)$, by adopting Player I's optimal strategy in the game $G_1(h)$. This means that, at the next step of the algorithm, we can safely eliminate any play p having h as a prefix that gives player $\iota(h)$ a payoff of less than $\alpha_1(h)$.

Step 2. At the second step of the iteration, we let $\beta_2(h) = \alpha_1(h)$, and we let $P_2(h)$ consist of plays p having h as a prefix, such that for each prefix g of p extending h , player $\iota(g)$'s payoff on p is at least $\beta_2(g)$. In view of the above discussion, $P_2(h)$ includes $P^*(h)$. The security payoff level $\alpha_2(h)$ is defined to be the value of the game $G_2(h)$. The rules of the game $G_2(h)$ are similar to those of $G_1(h)$, but Player II is now restricted to propose a play that is an element of the set $P_2(h_k, a_k)$, after a move (h_k, a_k) of Player I.

Thus $\alpha_2(h)$ is the highest payoff level player $\iota(h)$ could be sure to get against any strategy profile that only prescribes plays that survived the first round of elimination. It is also clear that player $\iota(h)$'s payoff at history h in any SPE is not smaller than $\alpha_2(h)$.

General step. At a general step ξ of the iteration, we let $\beta_\xi(h)$ denote the highest security level $\alpha_\lambda(h)$ over all iterations $\lambda < \xi$, and we define $P_\xi(h)$ as the set consisting of all plays p extending h that give a payoff of at least $\beta_\xi(g)$ for the player active at a prefix g of p . By a reasoning similar to that above we can see that $P_\xi(h)$ includes $P^*(h)$. The security level $\alpha_\xi(h)$ is defined as the value of the game $G_\xi(h)$ where Player II is restricted to propose plays from the set $P_\xi(h_k, a_k)$.

Limit properties. For every history h the sequences thus defined enjoy the obvious monotonicity properties: the sequence $P_\xi(h)$ is nonincreasing by inclusion, and the sequences $\alpha_\xi(h)$ and $\beta_\xi(h)$ are nondecreasing. For a sufficiently large ordinal number ξ^* , we have $P_{\xi^*}(h) = P^*(h)$ for every history $h \in H$. Whenever Ω has an SPE, the sets $P^*(h)$ are nonempty. In this case $\alpha_{\xi^*}(h) = \beta_{\xi^*}(h)$ is the minimum of $\iota(h)$'s payoff function on the set $P^*(h)$.

This algorithm shares a number of features with that in Flesch et al. [7]. In particular, it produces a nested sequence of sets of plays $P_\xi(h)$, and a nondecreasing sequence of security payoff levels $\alpha_\xi(h)$. The novelty of our approach lies in the way the security payoff levels $\alpha_\xi(h)$ are determined.

The definition of the payoff levels $\alpha_\xi(h)$ in Flesch et al. could be interpreted as being the value of the “short” auxiliary zero-sum game in which Player I chooses action a_0 , Player II chooses a play p_0 extending the history (h_0, a_0) , and Player I has no choice but to accept p_0 . In contrast, in the “full” zero-sum game used here, Player I is allowed to deviate arbitrarily often from the proposed plays.

This explains why our algorithm enjoys very different limit properties than the algorithm in Flesch et al. The limit of the algorithm in Flesch et al. is the set of plays that could be supported by strategy profiles that are immune to one-shot deviations. It is typically a superset of plays that could be supported by an SPE. Our algorithm outputs the latter set. The two sets coincide if the game satisfies the one-shot deviation principle. For a more elaborate discussion of one-deviation immunity we refer to Section 7.

Similar remarks apply to the comparison between our algorithm and that in Harris [14]. The essential difference (apart from the obvious differences in the setup) is the definition of the security levels.

3.2. The Formal Definition

For the ordinal 0 we define

$$P_0(h) = P(h),$$

$$\alpha_0(h) = \beta_0(h) = \min_{p \in P_0(h)} u_{\iota(h)}(p).$$

Let $\Phi(0)$ denote the statement that $P_0(h)$ is nonempty for each $h \in H$. The statement $\Phi(0)$ is obviously true.

For each ordinal $\xi > 0$ define

$$\beta_\xi(h) = \max_{\lambda < \xi} \alpha_\lambda(h), \quad (2)$$

$$P_\xi(h) = \left\{ p \in P(h) : \begin{array}{l} \text{for each } g \in H \text{ with } h \leq g < p \\ u_{\iota(g)}(p) \geq \beta_\xi(g) \end{array} \right\}. \quad (3)$$

We also define $\Phi(\xi)$ to be the following statement:

$$“P_\lambda(h) \neq \emptyset \text{ for each } \lambda \leq \xi \text{ and } h \in H.” \quad (4)$$

If $\Phi(\xi)$ is false, the algorithm terminates. In that case, for the sake of completeness we adopt the convention that $P_\lambda(h) = \emptyset$ for each $\lambda > \xi$ and each history $h \in H$.

Assuming that the statement $\Phi(\xi)$ is true, consider an auxiliary perfect information game $G_\xi(h)$. This is a zero-sum game where Player³ I essentially represents player $\iota(h)$. Players I and II move alternately, starting with Player I, who makes an initial move of the form $(h_0, a_0) \in H$, where $h_0 = h$. Suppose that, at period k of the game, Player I chooses a move $(h_k, a_k) \in H$. Then, Player II chooses a play $p_k \in P_\xi(h_k, a_k)$. Subsequently, at period $k+1$, Player I chooses either to accept p_k , or to make a move $(h_{k+1}, a_{k+1}) \in H$, where (1) $(h_k, a_k) \leq h_{k+1} < p_k$ and $\iota(h_{k+1}) = \iota(h)$, and (2) the sequence (h_{k+1}, a_{k+1}) is not a prefix of p_k . In the latter case we say that Player I deviates from p_k at the prefix h_{k+1} with the action a_{k+1} . If there is no move (h_{k+1}, a_{k+1}) with the stated properties, Player I has no choice but to accept p_k . The outcome of the game is the unique play q extending the histories $h = h_0 < h_1 < \dots$ if Player I never accepts, or the play $q = p_k$ if Player I accepts p_k at period $k+1$. Player I's payoff is then $u_{\iota(h)}(q)$, which is to be paid by Player II.

Finally, we define the *security level* as

$$\alpha_\xi(h) = \text{val}(G_\xi(h)),$$

the value of the game $G_\xi(h)$. The existence of the value is shown below in Lemma 3.1. This completes the recursive definition.

Lemma 3.1. *Let ξ be an ordinal such that the statement $\Phi(\xi)$ is true. Let $h \in H$. The payoff function of the game $G_\xi(h)$ is Borel measurable and has finite range. Consequently, the game has a value, and each player has a strategy that is optimal in every subgame.*

Proof. Take such an ordinal ξ and a history $h \in H$. The game $G_\xi(h)$ is a game in the sense of Section 2, once we adopt the following convention: once Player I accepts a proposal by Player II (let us denote this action by a^*), the game continues, Player I being the active player henceforth, and the only available action being a^* . With this convention, a play in $G_\xi(h)$ has either the form

$$((h_0, a_0), p_0, (h_1, a_1), p_1, \dots) \quad (5)$$

or the form

$$((h_0, a_0), p_0, \dots, (h_k, a_k), p_k, a^*, a^*, \dots), \quad (6)$$

where the actions (h_m, a_m) and p_m of $G_\xi(h)$ obey the rules of the game as described above.

Let Z denote the set of plays in $G_\xi(h)$. For each $z \in Z$ let $f(z)$ denote the outcome of the game $G_\xi(h)$. Thus $f(z)$ is the unique play that extends the histories $h = h_0 < h_1 < \dots$ if z is of the form (5), or $f(z) = p_k$ if z is of the form (6) for some $k \in \mathbb{N}$. Then Player I's payoff function v in the game $G_\xi(h)$ is given by $v(z) = u_{i(h)}(f(z))$ for each $z \in Z$, that is, v is a composition of f and $u_{i(h)}$.

We endow Z with the topology as in Section 2. Then $f: Z \rightarrow P$ is a continuous function: indeed, it is easy to see that if a sequence $(z_n)_{n \in \mathbb{N}}$ converges to z then the sequence $(f(z_n))_{n \in \mathbb{N}}$ converges to $f(z)$. It follows that f is Borel measurable. Now recall that the function $u_{i(h)}$ is assumed to be Borel measurable. It follows that v is Borel measurable, being a composition of Borel measurable functions. Since $u_{i(h)}$ is further assumed to have finite range, the function v has finite range. Now the result follows by Lemma 2.1. $\square$

Notice that if the statement $\Phi(\xi)$ is true, then the game $G_\xi(h)$, the quantities $\alpha_\xi(h)$, $\beta_{\xi+1}(h)$, and the sets $P_{\xi+1}(h)$ are all well defined for each $h \in H$.

Lemma 3.2 (Monotonicity of the Sequences). *Let ξ be an ordinal such that the statement $\Phi(\xi)$ is true. Then, for each history $h \in H$ it holds that*

- (a) $P_\lambda(h) \supseteq P_\xi(h)$ for every ordinal $\lambda \leq \xi$,
- (b) $\alpha_\lambda(h) \leq \alpha_\xi(h)$ for every ordinal $\lambda \leq \xi$,
- (c) $\beta_\xi(h) \leq \alpha_\xi(h) = \beta_{\xi+1}(h)$.

Proof. We first prove claims (a) and (b). Both claims are obvious if $\lambda = 0$. Thus, suppose that $\lambda > 0$. As $\lambda \leq \xi$, we have

$$\beta_\lambda(h) = \max_{\eta < \lambda} \alpha_\eta(h) \leq \max_{\eta < \xi} \alpha_\eta(h) = \beta_\xi(h)$$

for each $h \in H$. Therefore, (a) follows. Now consider an optimal strategy for Player I in the game $G_\lambda(h)$. By (a), this same strategy can also be used in the game $G_\xi(h)$, while at the same time any strategy for Player II in the game $G_\xi(h)$ is also a legal strategy in the game $G_\lambda(h)$. It follows that Player I can guarantee a payoff of at least $\alpha_\lambda(h)$ in the game $G_\xi(h)$. Therefore, (b) holds. The claim (c) follows directly from (b). $\square$

We argue that there exists an ordinal ξ^* such that

$$P_{\xi^*}(h) = P_\lambda(h) \quad \text{for all } h \in H \text{ and all } \lambda > \xi^*. \quad (7)$$

Now, if $P_\xi(h) = \emptyset$ for some $h \in H$ and some ordinal ξ , then by our convention $P_\lambda(h) = \emptyset$ for all histories h and all $\lambda > \xi$. Suppose that $P_\xi(h) \neq \emptyset$ for all h and all ξ . Then in view of Lemma 3.2 there exists⁴ an ordinal ξ^* such that $P_{\xi^*}(h) = P_{\xi^*+1}(h)$ for every $h \in H$. Finally the definitions imply that ξ^* satisfies (7). Henceforth we let ξ^* be the least ordinal satisfying (7).

3.3. Examples

We illustrate the algorithm by means of two examples. The first example comes from Flesch et al. [9]. The game has no SPE,⁵ and the algorithm of Section 3.2 outputs a collection of empty sets. The second example, which admits a unique SPE, demonstrates that the algorithm might take more than ω steps to terminate.

Example 3.3 (Flesch et al. [9]). Consider the following two-player game. The action set is $A = \{1, 2\}$. Player 1 starts the game. The active player decides who the next active player is by choosing the corresponding action. Let S_0 denote the set of plays where both actions 1 and 2 are played infinitely many times, S_1 the set of plays of the form $h1^\infty = (h, 1, 1, 1, \dots)$ for $h \in H$, and S_2 the set of plays of the form $h2^\infty = (h, 2, 2, 2, \dots)$ for $h \in H$. The payoffs are $(0, 0)$ for any play in S_0 , $(-1, 2)$ for any play in S_1 , and $(-2, 1)$ for any play in S_2 .

The values $P_\xi(h)$ and $\alpha_\xi(h)$ are given by the Table 1, where $S_1(h) = S_1 \cap P(h)$. The algorithm terminates after four rounds with an empty set of plays at each history. Corollary 4.2 thus implies that the game has no SPE.

The entries of the table are obtained as follows. In view of item (3) of Lemma 3.2, the calculation of $\beta_\xi(h)$ is straightforward, so we focus on the computation of $P_\xi(h)$ and $\alpha_\xi(h)$.

Step 0. By definition.

Step 1. We have $P_1(h) = P(h)$ by definition. Consider player 1's history h . In the game $G_1(h)$ Player I, who represents player 1 of the original game, "realizes" that he can guarantee the payoff of -1 by always taking action 1. Player II makes sure that Player I's payoff is not greater than -1 by proposing, after each move (h_k, a_k) of Player I, the play $(h_k, a_k)2^\infty$. This implies $\alpha_1(h) = -1$.

Consider player 2's history h . In the game $G_1(h)$ Player I (who is essentially representing player 2 of the original game) can guarantee the payoff of 1 since he can induce the play $h2^\infty$. Player II makes sure that I's payoff is not greater than 1 by proposing the play $(h_k, a_k)2^\infty$ after each move of Player I. We obtain $\alpha_1(h) = 1$.

Table 1. The iterations in Example 3.3.

ξ	$\beta_\xi(h)$		$P_\xi(h)$		$\alpha_\xi(h)$	
	$\iota(h) = 1$	$\iota(h) = 2$	$\iota(h) = 1$	$\iota(h) = 2$	$\iota(h) = 1$	$\iota(h) = 2$
0	-2	0	$P(h)$	$P(h)$	-2	0
1	-2	0	$P(h)$	$P(h)$	-1	1
2	-1	1	$S_1(h)$	$\{h2^\infty\} \cup S_1(h)$	-1	2
3	-1	2	$S_1(h)$	$S_1(h)$	0	2
4	0	2	$\emptyset$	$\emptyset$	*	*

Step 2. For each player 1's history h we obtain $P_2(h) = S_1(h)$. Consider player 2's history h . The set $P_2(h)$ consists of the play $h2^\infty$ and all plays of the set $S_1(h)$. Notice that $h2^\infty$ is an element of $P_2(h)$ since only player 2 is active along this play after history h .

Consider player 1's history h . We have $\alpha_2(h) = -1$ as in Step 1.

Consider player 2's history h and the corresponding auxiliary game $G_2(h)$. Player I realizes that, if he plays action 1, he is proposed a play in the set $S_1(h, 1)$. Accepting such a proposal yields him the payoff of 2. This leads to $\alpha_2(h) = 2$.

Step 3. For a player 1's history h we obtain $P_3(h) = S_1(h)$. Consider player 2's history h . The play $h2^\infty$ is not in $P_3(h)$ since it violates the beta value for player 2, which has been updated to 2. Consequently we have $P_3(h) = S_1(h)$.

Consider player 1's history h and the corresponding auxiliary game $G_3(h)$. Suppose Player I's makes the first move $(h_0, a_0) = (h, 2)$. He is then offered a play p_0 , an element of the set $P_3(h, 2) = S_1(h, 2)$. The play p_0 could be written in the form h_11^∞ for some player 1's history h_1 extending (h_0, a_0) . Consequently Player I could deviate from p_0 by playing $(h_1, a_1) = (h_1, 2)$. He is then offered a play $p_1 \in S_1(h_1, 2)$. He could likewise deviate from p_1 . Continuing this way, Player I induces a play in which both actions 1 and 2 appear infinitely often, yielding him the payoff of 0. Thus $\alpha_3(h) = 0$.

For player 2's history h we have $\alpha_3(h) = 2$, since $\alpha_2(h) = 2$.

Step 4. It is clear that there are no plays satisfying the beta levels of 0 and 2 for players 1 and 2, respectively. $\triangleleft$

Example 3.4. Consider an infinite centipede game with infinitely many players in which player 0 plays once, and all other players move infinitely many times. All payoff functions are binary and are upper semicontinuous.

The set of players is $I = \mathbb{N}$. Let $\mathbb{N}^+ = \mathbb{N} \setminus \{0\}$. Take any infinite partition $\{S^i: i \in \mathbb{N}^+\}$ of $\mathbb{N}^+$ such that each S^i is infinite. Player 0 is active in period 0 and each player $i \in \mathbb{N}^+$ is active at the infinitely many periods belonging to S^i . The action set is $A = \{s, c\}$, where s stands for stopping the game, and c for continuing. Player 0's payoff is 0 if any player stops, and 1 if no one stops. Player $i \in \mathbb{N}^+$ gets a payoff of 0 if a player $j \leq i$ stops the game, and his payoff is 1 otherwise.

The algorithm is illustrated by the table below. The rows correspond to the iteration steps. The first column shows the complement of $P_\xi(\emptyset)$, and the numbers in the remaining columns show the security levels $\alpha_\xi(h)$ for the respective players, for nonterminal histories h , that is, histories where no one yet has stopped the game. We let B_i be the set of plays where player i is the first player to stop.

The calculation of the values in Table 2 is straightforward. Consider iteration step $\xi = 1$. Take a player 1's history h . Then Player I (who represents player 1 of the original game) can guarantee the payoff of 1 in the game $G_1(h)$ by never stopping. Thus $\alpha_1(h) = 1$. Take a player i 's history h where $i > 1$. Now Player I (representing

Table 2. The iterations in Example 3.4.

	$P \setminus P_\xi(\emptyset)$	$\iota(h) = 0$	$\iota(h) = 1$	$\iota(h) = 2$	$\iota(h) = 3$	...
$\xi = 0$	$\emptyset$	0	0	0	0	...
$\xi = 1$	$\emptyset$	0	1	0	0	...
$\xi = 2$	B_1	0	1	1	0	...
$\xi = 3$	$B_1 \cup B_2$	0	1	1	1	...
$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$	$\vdots$
$\xi = \omega$	$B_1 \cup B_2 \cup \dots$	1	1	1	1	...
$\xi = \omega + 1$	$B_0 \cup B_1 \cup B_2 \cup \dots$	1	1	1	1	...

player i of the original game) cannot guarantee 1, since after any period where player i is active, eventually there comes a period where player 1 is active, and if player 1 stops, player i gets payoff 0. Thus $\alpha_1(h) = 0$. Likewise, $\alpha_1(\emptyset) = 0$.

At iteration step $\xi = 2$, player 1 never stops. Starting with player 2, the analysis is similar to that of the previous paragraph. That is, for any history h of player 2 we have $\alpha_2(h) = 1$, and for any history h of player $i > 2$ we have $\alpha_2(h) = 0$. Likewise, $\alpha_2(\emptyset) = 0$.

A similar argument holds for each $\xi < \omega$.

At iteration step $\xi = \omega$, we obtain $\beta_\omega(\emptyset) = 0$ and $\beta_\omega(h) = 1$ for any history h other than $\emptyset$. As a consequence, no player $i \in \mathbb{N}^+$ stops. Player 0, on the other hand, is still allowed to stop since $\beta_\omega(\emptyset) = 0$.

At $\xi = \omega + 1$, player 0 can guarantee 1 by playing continue. The iteration reaches a fixed point at $\xi^* = \omega + 1$. Note that $P_{\xi^*}(\emptyset)$ is the singleton consisting of the play $(c, c, c, \dots)$. And indeed, in the unique SPE, each player plays continue at each of his histories. $\triangleleft$

It is not difficult to show, by means of examples like the one above, that the ordinal ξ_* can be arbitrarily large. Examples of this sort have been previously constructed in Flesch et al. [7] and Flesch and Predtetchinski [11].

4. Overview of the Main Results

In this section we provide an overview of the main findings of the paper. We first state two main results, and then turn to the discussion of various additional results.

Our first result is a characterization of SPE plays in terms of the limit of the recursive sequence as defined in the previous section. The proof is delegated to Section 5.

Theorem 4.1 (A Characterization of SPE Plays). *Consider a perfect information game Ω . Assume that the payoff function of each player has finite range and is Borel measurable. Then $P^*(h) = P_{\xi^*}(h)$ for every $h \in H$. In particular, for a play $p \in P$, there exists an SPE σ such that $p = \pi(\sigma, \emptyset)$ if and only if $p \in P_{\xi^*}(\emptyset)$.*

Corollary 4.2. *Consider a perfect information game Ω . Assume that the payoff function of each player has finite range and is Borel measurable. Then the game has an SPE if and only if $P_{\xi^*}(\emptyset) \neq \emptyset$, or equivalently, if and only if $P_\xi(h) \neq \emptyset$ for each history $h \in H$ and each ordinal ξ .*

We now turn to our second contribution, the existence theorems for games with upper semicontinuous payoffs, which is from a technical point of view more demanding than the first. We apply Corollary 4.2 to address an open problem raised in Purves and Sudderth [24]. The authors show that there is an SPE provided that each player's payoff function has finite range and is upper semicontinuous, and that the number of players is finite. They further raise the question of whether one could dispense with the second assumption, the finiteness of the player set. We answer the question affirmatively.

The technique employed by Purves and Sudderth is based on the induction on the number of distinct payoff vectors in the game. This feature of the proof makes it difficult to extend it beyond the case of finitely many players. Instead, we show that in games where each player's payoff function has a finite range and is upper semicontinuous, our algorithm returns a nonempty set of plays at each step of the iteration. This is Theorem 4.3. Corollary 4.4 then follows by discretizing the payoff functions, exactly as is done in Purves and Sudderth [24].

Theorem 4.3 (Existence of SPE in Upper Semicontinuous Games). *Consider a perfect information game. Assume that the payoff function of each player has finite range and is upper semicontinuous. Then, $P_\xi(h) \neq \emptyset$ for each history $h \in H$ and each ordinal ξ . Consequently, the game admits an SPE.*

Corollary 4.4. *Consider a perfect information game. Assume that the payoff function of each player is bounded and upper semicontinuous. Then, for every $\epsilon > 0$, the game admits an ϵ -SPE.*

The assumption of upper semicontinuity implies that the sets $P_\xi(h)$ are all closed. This observation plays a crucial role in the proof of Theorem 4.3. The proof is delegated to Section 6.

We presently turn to additional results. Our goal is to evaluate to what extent our results depend on the assumption that every player's payoff function only takes finitely many values. Therefore in the final section of the paper we consider a richer class of games, games where the payoff functions are only required to be bounded and Borel measurable.

In this larger class of games one can still show that $P_{\xi^*}(\emptyset)$ contains all SPE plays. The converse is no longer true. However, we are able to show that each play in $P_{\xi^*}(\emptyset)$ could be induced as an ϵ -SPE for every $\epsilon > 0$ (Theorem 8.1). This discrepancy arises because in the auxiliary zero-sum game Player II might not have an optimal strategy. We refer the reader to Section 8.1 for a more thorough discussion of this point.

Our final result concerns the technique of “discretizing” payoff functions, that is, approximating a bounded payoff function by a payoff function having finite range. As mentioned in the introduction, the technique has been used frequently to reduce the problem of existence of ϵ -SPE in games with bounded payoff functions to the problem of existence of SPE in a game where the payoff functions have finite range. Also in this paper, it is by means of a discretization of payoff functions that we pass from Theorem 4.3 to Corollary 4.4. Our result could be interpreted as saying that passing to discretizations is without loss of generality, at least if one is only interested in the question of existence of ϵ -SPE for every $\epsilon > 0$. We refer to Theorem 8.3 for a rigorous statement of the result.

5. The Proof of Theorem 4.1

The proof of the inclusion $P^*(h) \subseteq P_{\xi^*}(h)$ is straightforward: we show by induction on ξ that if σ is an SPE, then $\pi(\sigma, h) \in P_{\xi}(h)$ for every history $h \in H$.

Proof that $P^*(h) \subseteq P_{\xi^*}(h)$ for each $h \in H$. Let σ be an SPE of Ω . Let $\Lambda(\xi)$ be the following statement: $\pi(\sigma, h) \in P_{\xi}(h)$ for every history $h \in H$. We prove that $\Lambda(\xi)$ is true for each ξ by transfinite induction on ξ .

For $\xi = 0$, the statement $\Lambda(\xi)$ is true by the definition of $P_0(h)$. Now, for an ordinal $\xi > 0$, suppose that $\Lambda(\lambda)$ is true for all $\lambda < \xi$. We need to show that $\Lambda(\xi)$ is also true.

Consider a history $h \in H$. Take any history $g \in H$ such that $h \leq g < \pi(\sigma, h)$. We need to show that $u_{i(g)}(\pi(\sigma, h)) \geq \beta_{\xi}(g)$. Notice that $\pi(\sigma, h) = \pi(\sigma, g)$. Thus, in view of (2) it is sufficient to show that

$$u_{i(g)}(\pi(\sigma, g)) \geq \alpha_{\lambda}(g) \quad (8)$$

for every for every ordinal $\lambda < \xi$.

Take an ordinal $\lambda < \xi$ and consider the game $G_{\lambda}(g)$. The strategy profile σ induces a strategy Δ_{II} for Player II in the game $G_{\lambda}(g)$, as follows: after Player I's move $(h_k, a_k) \in H$, let Player II propose the play $p_k = \pi(\sigma, (h_k, a_k))$. This is a legal strategy for Payer II since $\pi(\sigma, (h_k, a_k)) \in P_{\lambda}(h_k, a_k)$ by the induction hypothesis. Let Δ_{I} be a strategy of Player I in the game $G_{\lambda}(g)$ that makes the initial move (h_0, a_0) , where $h_0 = g$ and $a_0 = \sigma_{i(g)}(g)$, and then accepts the first proposal of Player II. Clearly, using Δ_{I} against the strategy Δ_{II} results in the play $\pi(\sigma, g)$ being proposed and accepted. Since under the strategy Δ_{II} Player II essentially proposes to follow the SPE strategy profile σ , Player I can gain no more than $u_{i(g)}(\pi(\sigma, g))$ by playing against Δ_{II} . Hence, Δ_{I} is a best response to Δ_{II} . The inequality (8) follows.

This proves that the statement $\Lambda(\xi)$ is true. $\square$

The intuition behind the proof of the inclusions $P^*(h) \supseteq P_{\xi^*}(h)$ is as follows. We first prove the inclusion for $h = \emptyset$ and then generalize to an arbitrary history h .

If $p \in P_{\xi^*}(\emptyset)$, we can define an SPE σ thus: from the root of the game Ω , follow the play p , as long as no player deviates. If a player i deviates from p at a history h_0 by playing action a_0 , choose some strategy $\Sigma_{\text{II}}^{h_0}$ for Player II in the auxiliary game $G_{\xi^*}(h_0)$ that is optimal in every subgame (recall that such a strategy exists by Lemma 2.1). Then σ prescribes to follow the play p_0 proposed by Player II under the strategy $\Sigma_{\text{II}}^{h_0}$ after the history (h_0, a_0) . Suppose that some player deviates from p_0 at a history h_1 with an action a_1 . If it is player i again, then σ prescribes to follow the play p_1 proposed by Player II still under the strategy $\Sigma_{\text{II}}^{h_0}$, after the history $((h_0, a_0), p_0, (h_1, a_1))$. If the deviating player is other than i , choose some strategy $\Sigma_{\text{II}}^{h_1}$ for Player II in the auxiliary game $G_{\xi^*}(h_1)$ that is optimal in every subgame. Then σ prescribes to follow the play p_1 proposed by Player II under his strategy $\Sigma_{\text{II}}^{h_1}$ after the history (h_1, a_1) . The process thus continues ad infinitum.

Before we proceed to the proof, we recall that the set $P_{\xi^*}(h)$ is nonempty for some history h if and only if $P_{\xi^*}(h)$ is nonempty for every h . We can thus assume that $P_{\xi^*}(h) \neq \emptyset$ for every $h \in H$, for otherwise there is nothing to prove.

Proof that $P^*(\emptyset) \supseteq P_{\xi^*}(\emptyset)$. Take a $p \in P_{\xi^*}(\emptyset)$. We construct an SPE σ in Ω such that $p = \pi(\sigma, \emptyset)$.

For each history $g \in H$, fix a strategy Σ_{II}^g for Player II in the game $G_{\xi^*}(g)$ as in Lemma 3.1, i.e., a strategy that is optimal in every subgame of $G_{\xi^*}(g)$.

Define a strategy profile σ in Ω as follows. The players are to follow the play p . Suppose that the players follow p until history h_0 , with $h_0 < p$, where the first deviation occurs, by player $i(h_0)$ taking action a_0 incompatible with p . Let $p_0 = \Sigma_{\text{II}}^{h_0}(h_0, a_0)$. Starting with the history (h_0, a_0) , the players are to follow the play p_0 .

Consider a sequence $((h_0, a_0), p_0, \dots, (h_k, a_k), p_k)$, where $(h_0, a_0), \dots, (h_k, a_k)$ are successive deviations and for each m , the play p_m is the play that σ prescribes the players to follow starting with the history (h_m, a_m) . Suppose

that the next deviation occurs at a history h_{k+1} , with $(h_k, a_k) \leq h_{k+1} < p_k$, and let a_{k+1} be player $i(h_{k+1})$'s action at h_{k+1} not compatible with p_k . Let ℓ be the least element of $\{0, \dots, k+1\}$ such that $i(h_\ell) = \dots = i(h_{k+1})$. Define $p_{k+1} = \sum_{\Pi}^{h_\ell} ((h_\ell, a_\ell), p_\ell, \dots, (h_k, a_k), p_k, (h_{k+1}, a_{k+1}))$. Then, the players are to follow the play p_{k+1} starting with the history (h_{k+1}, a_{k+1}) . This completes the definition of the strategy profile σ .

By construction, $\pi(\sigma, \emptyset) = p$. We argue that σ is an SPE in Ω .

First notice that for every history $g \in H$ we have $\pi(\sigma, g) \in P_{\xi^*}(g)$. To see this, first suppose that g is a prefix of p . Then $\pi(\sigma, g) = p$. Since p is an element of $P_{\xi^*}(\emptyset)$ and g is its prefix, p is also an element of $P_{\xi^*}(g)$, as desired. Suppose now that g is not a prefix of p . Then to g there corresponds a sequence $((h_0, a_0), p_0, \dots, (h_k, a_k), p_k)$ as above such that (h_k, a_k) is a prefix of g and $\pi(\sigma, g) = p_k$. Since p_k is an element of $P_{\xi^*}(h_k, a_k)$ and g is its prefix, p_k is also an element of $P_{\xi^*}(g)$, as desired.

Now, take any $g \in H$. As $P_{\xi^*}(g) = P_{\xi^*+1}(g)$, we have $u_{i(g)}(\pi(\sigma, g)) \geq \beta_{\xi^*+1}(g) = \alpha_{\xi^*}(g)$. Thus player $i(g)$ obtains at least $\alpha_{\xi^*}(g)$ by playing in accordance with σ in the subgame of Ω starting with g . On the other hand, following a deviation by player $i(g)$ at g , the strategy σ essentially replicates the behavior of Player II in the game $G_{\xi^*}(g)$, so that the deviation yields no more than $\alpha_{\xi^*}(g)$. We conclude that σ is an SPE. $\square$

Proof that $P^*(h) \supseteq P_{\xi^*}(h)$ for each $h \in H$. Now take a history $h \in H$ of length t and a play $p \in P_{\xi^*}(h)$. We construct an SPE σ such that $\pi(\sigma, h) = p$.

An argument similar to that above could be used to show that the subgame of Ω starting at history h has an SPE σ^h that induces the play p . Likewise, for each history $h' \in H$ of length t , the subgame of Ω starting at history h' has an SPE $\sigma^{h'}$. Finally, we can use backward induction starting at period t to build a strategy profile σ that is an SPE of Ω and that agrees with the strategy profiles $\sigma^{h'}$ for every h' of length t in the subgame following h' . In particular, $\pi(\sigma, h) = p$. $\square$

6. The Proof of Theorem 4.3

Throughout this section, we assume that for each $i \in I$ the function u_i has finite range and is upper semicontinuous. The crucial consequence of this assumption is that the sets $P_{\xi}(h)$ are closed.

Lemma 6.1 (Closedness). *Let ξ be an ordinal such that the statement $\Phi(\xi)$ is true. Then, for each history $h \in H$, the set $P_{\xi}(h)$ is closed.*

Proof. The set $P_{\xi}(h)$ can be expressed as an intersection of the set $P_0(h)$ and sets of the form $\{p \in P: u_i(p) \geq x\}$ where i is a player and x a real number. Since the functions u_i are upper semicontinuous, all of such sets are closed. As the intersection of closed sets is closed, the result follows. $\square$

For an ordinal ξ let $\Psi(\xi)$ denote the statement that for each $h \in H$ there exists an action $\sigma_{\xi}(h)$ such that $(h, \sigma_{\xi}(h)) \in H$ and

$$P_{\xi}(h) \supseteq P_{\xi}(h, \sigma_{\xi}(h)). \quad (9)$$

We shall prove that the statements $\Phi(\xi)$ and $\Psi(\xi)$ are true for every ordinal ξ simultaneously by transfinite induction. As remarked earlier, $\Phi(0)$ is true. Clearly $\Psi(0)$ is also true. As $\sigma_0(h)$ one could take an arbitrary action a such that $(h, a) \in H$.

Lemma 6.2. *Let ξ be an ordinal such that the statements $\Phi(\xi)$ and $\Psi(\xi)$ are both true. Then, for each $h \in H$, the play $\pi(\sigma_{\xi}, h)$ is an element of $P_{\xi}(h)$.*

Proof. Take a history $h \in H$. Let $p = \pi(\sigma_{\xi}, h)$ and let $h_0, h_1, \dots$ be the enumeration of the successive prefixes of p starting with $h_0 = h$. By $\Phi(\xi)$, the set $P_{\xi}(h_t)$ is nonempty for each $t \in \mathbb{N}$. By (9), we know that $P_{\xi}(h_0) \supseteq P_{\xi}(h_1) \supseteq \dots$. Now take any $p_t \in P_{\xi}(h_t)$ for every $t \in \mathbb{N}$. Then $p_t \in P_{\xi}(h)$. Clearly, the sequence $\{p_t\}$ converges to the play p . Since $P_{\xi}(h)$ is a closed set by Lemma 6.1, we conclude that $p \in P_{\xi}(h)$, as desired. $\square$

Lemma 6.3. *Let ξ be a limit ordinal such that the statements $\Phi(\lambda)$ and $\Psi(\lambda)$ are both true for each $\lambda < \xi$. Then, $\Phi(\xi)$ and $\Psi(\xi)$ are also true.*

Proof. Take a history $h \in H$. Define recursively the sequence $h_0, h_1, \dots$ of histories and the sequence of ordinals $\lambda_0, \lambda_1, \dots$ as follows. Let $h_0 = h$. Suppose that the history h_k has been defined. Let $\lambda_k < \xi$ be a successor ordinal such that

$$\beta_{\lambda_k}(h_k) = \beta_{\xi}(h_k).$$

If $k > 0$, we require in addition that $\lambda_{k-1} \leq \lambda_k$. To see that such an ordinal exists, notice that by (2), the sequence $(\beta_{\lambda}(h_k))_{\lambda < \xi}$ is nondecreasing, and $\max_{\lambda < \xi} \beta_{\lambda}(h_k) = \beta_{\xi}(h_k)$. Hence, there is an ordinal $\lambda_* < \xi$ such that $\beta_{\lambda}(h_k) = \beta_{\xi}(h_k)$ for each λ with $\lambda_* \leq \lambda < \xi$. This implies that there does exist an ordinal λ with the desired properties.

Let $h_{k+1} = (h_k, \sigma_{\lambda_k}(h_k))$. This completes the recursive definition.

Now for each $k \in \mathbb{N}$, we have the inclusions

$$P_{\lambda_k}(h_k) \supseteq P_{\lambda_k}(h_{k+1}) \supseteq P_{\lambda_{k+1}}(h_{k+1}), \quad (10)$$

where the first inclusion is by (9), and the second inclusion holds by $\lambda_k \leq \lambda_{k+1}$ and by Lemma 3.2. Hence, for $k, \ell \in \mathbb{N}$

$$P_{\lambda_k}(h_k) \supseteq P_{\lambda_\ell}(h_\ell), \quad \text{whenever } k \leq \ell. \quad (11)$$

Let p be the unique play extending the histories $h_0 < h_1 < \dots$. We argue that $p \in P_\xi(h)$.

For each $\ell \in \mathbb{N}$, let p_ℓ be an arbitrary element of the set $P_{\lambda_\ell}(h_\ell)$. Since h_ℓ is prefix of p_ℓ , the sequence $\{p_\ell\}_{\ell \in \mathbb{N}}$ of plays converges to the play p . Now fix a $k \in \mathbb{N}$. We know by (11) that $p_\ell \in P_{\lambda_k}(h_k)$ for all $\ell \geq k$. As $P_{\lambda_k}(h_k)$ is a closed set by Lemma 6.1, we conclude that $p \in P_{\lambda_k}(h_k)$. Thus $u_{\iota(h_k)}(p) \geq \beta_{\lambda_k}(h_k) = \beta_\xi(h_k)$. Since this is true for each $k \in \mathbb{N}$, we conclude that $p \in P_\xi(h)$.

This proves the statement $\Phi(\xi)$.

To prove the statement $\Psi(\xi)$, let $\sigma_\xi(h_0) = \sigma_{\lambda_0}(h_0)$. We show that $P_\xi(h_0) \supseteq P_\xi(h_0, \sigma_\xi(h_0))$. Take a $q \in P_\xi(h_0, \sigma_\xi(h_0))$. To prove $q \in P_\xi(h_0)$, it suffices to show that $u_{\iota(h_0)}(q) \geq \beta_\xi(h_0)$. We have $P_{\lambda_0}(h_0) \supseteq P_{\lambda_0}(h_1) \supseteq P_\xi(h_1) = P_\xi(h_0, \sigma_\xi(h_0))$, where the first inclusion comes from (10), and the second one is by monotonicity. So $q \in P_{\lambda_0}(h_0)$ and hence $u_{\iota(h_0)}(q) \geq \beta_{\lambda_0}(h_0) = \beta_\xi(h_0)$, as desired.

This proves the statement $\Psi(\xi)$. $\square$

We now proceed with the more challenging part of the proof: showing that if $\Phi(\xi)$ and $\Psi(\xi)$ are both true for some ordinal ξ , then $\Phi(\xi + 1)$ and $\Psi(\xi + 1)$ are also true.

Let ξ be an ordinal such that the statements $\Phi(\xi)$ and $\Psi(\xi)$ are true. Take an $h \in H$. Fix some $\sigma_\xi(h)$ satisfying (9). Also fix some optimal strategy Σ_ξ^h for Player I in the game $G_\xi(h)$. Suppose that (h, a) is the first move Player I makes under the strategy Σ_ξ^h . Define $\sigma_{\xi+1}(h) = a$.

Given a set $F \subseteq I$ of players, let $\sigma_\xi^F(h)$ denote $\sigma_\xi(h)$ if $\iota(h) \notin F$ and denote $\sigma_{\xi+1}(h)$ if $\iota(h) \in F$. Thus, in particular, $\sigma_\xi^\emptyset(h) = \sigma_\xi(h)$ and $\sigma_\xi^I(h) = \sigma_{\xi+1}(h)$. Define also

$$P_\xi^F(h) = \left\{ p \in P(h): \begin{array}{ll} \text{for every } g \in H \text{ with } h \leq g < p & \\ \text{if } \iota(g) \in F \text{ then } u_{\iota(g)}(p) \geq \alpha_\xi(g) & \\ \text{if } \iota(g) \notin F \text{ then } u_{\iota(g)}(p) \geq \beta_\xi(g) & \end{array} \right\}.$$

Thus, in particular, $P_\xi^\emptyset(h) = P_\xi(h)$. In view of Claim (3) of Lemma 3.2, $P_\xi^I(h) = P_{\xi+1}(h)$. By the same claim, if $F \subseteq E \subseteq I$ then $P_\xi^F(h) \supseteq P_\xi^E(h)$. As in Lemma 6.1 the sets $P_\xi^F(h)$ can be shown to be closed.

Lemma 6.4. *Let ξ be an ordinal such that the statements $\Phi(\xi)$ and $\Psi(\xi)$ are true. Then, for each $h \in H$ and each $F \subseteq I$, it holds that $P_\xi^F(h) \supseteq P_\xi^F(h, \sigma_\xi^F(h))$.*

Proof. Take a $p \in P_\xi^F(h, \sigma_\xi^F(h))$. We distinguish two cases, depending on whether $\iota(h)$ is an element of F or not.

Suppose first that $\iota(h) \notin F$. In order to show that $p \in P_\xi^F(h)$, it suffices to prove that $u_{\iota(h)}(p) \geq \beta_\xi(h)$. We have $\sigma_\xi^F(h) = \sigma_\xi(h)$. Moreover, $P_\xi(h) \supseteq P_\xi(h, \sigma_\xi(h)) \supseteq P_\xi^F(h, \sigma_\xi(h))$, where the first inclusion is (9). Hence $p \in P_\xi(h)$ and thus $u_{\iota(h)}(p) \geq \beta_\xi(h)$.

Suppose now that $\iota(h) \in F$. In order to show that $p \in P_\xi^F(h)$, it suffices to prove that $u_{\iota(h)}(p) \geq \alpha_\xi(h)$. We have $\sigma_\xi^F(h) = \sigma_{\xi+1}(h)$. Now recall that Σ_ξ^h is an optimal strategy for Player I in $G_\xi(h)$ such that its first move is $(h, \sigma_{\xi+1}(h))$.

Suppose that, following Player I's move $(h_0, a_0) = (h, \sigma_{\xi+1}(h))$, Player II proposes the play $p_0 = p$. This is a legitimate move by Player II since $p_0 \in P_\xi^F(h_0, a_0) \subseteq P_\xi(h_0, a_0)$. Consider Player I's reaction to p_0 .

If, under the strategy Σ_ξ^h , Player I accepts p_0 , we must have $u_{\iota(h)}(p_0) \geq \alpha_\xi(h)$, for otherwise Σ_ξ^h would not be optimal. Now suppose that, under the strategy Σ_ξ^h , Player I responds to p_0 by making a move $(h_1, a_1) \in H$. From Player II's point of view, what follows then is strategically equivalent to what happens in the game $G_\xi(h_1)$ after Player I makes the move (h_1, a_1) . Hence, following (h_1, a_1) , Player II can make sure that I's payoff is not greater than $\alpha_\xi(h_1)$. Since Σ_ξ^h is optimal, we must then have $\alpha_\xi(h_1) \geq \alpha_\xi(h)$. Since $(h_0, a_0) \leq h_1$ and since $p_0 \in P_\xi^F(h_0, a_0)$, we also have $u_{\iota(h)}(p_0) \geq \alpha_\xi(h_1)$. It follows that $u_{\iota(h)}(p_0) \geq \alpha_\xi(h)$, as desired. $\square$

Lemma 6.5. *Let ξ be an ordinal such that the statements $\Phi(\xi)$ and $\Psi(\xi)$ are true. Let $F \subseteq I$ be a finite set. Then, for each $h \in H$, the play $\pi(\sigma_\xi^F, h)$ is an element of $P_\xi^F(h)$.*

Proof. The proof is inductive on the cardinality of F . Lemma 6.2 implies that the result holds for the set $F = \emptyset$. Suppose that the statement has been proven for each set $F \subseteq I$ of cardinality n . Consider a set $E \subseteq I$ of cardinality $n + 1$. We prove that for each $h \in H$ the play $\pi(\sigma_\xi^E, h)$ is an element of $P_\xi^E(h)$.

Take a history $h \in H$. Let $p = \pi(\sigma_\xi^E, h)$ and let $h_0, h_1, \dots$ be the enumeration of the successive prefixes of p starting with $h_0 = h$. In view of Lemma 6.4, we have the inclusions $P_\xi^E(h_0) \supseteq P_\xi^E(h_1) \supseteq \dots$. Therefore, it suffices to show that there is a $t \in \mathbb{N}$ such that $p \in P_\xi^E(h_t)$.

Take an arbitrary player $i \in E$. We consider two cases, depending on whether player i controls finitely or infinitely many prefixes of the play p .

Suppose first that player i controls only finitely many prefixes of the play p . Let $K \in \mathbb{N}$ be such that none of the prefixes $h_K, h_{K+1}, \dots$ are controlled by player i . Then $p = \pi(\sigma_\xi^E, h_K) = \pi(\sigma_\xi^F, h_K)$, where $F = E \setminus \{i\}$. By the induction hypothesis the play $p = \pi(\sigma_\xi^F, h_K)$ is an element of $P_\xi^F(h_K)$. Again using the fact that none of the prefixes of p after time K are controlled by player i , we conclude that $p \in P_\xi^E(h_K)$.

Suppose now that player i controls infinitely many prefixes of the play p . Let $h_{t_0}, h_{t_1}, \dots$ be the subsequence of the sequence $h_0, h_1, \dots$ of the prefixes of p that are controlled by player i . We let $g_k = h_{t_k}$ for each $k \in \mathbb{N}$ and we let a_k denote the action of p following the prefix g_k , that is, $(g_k, a_k) = (h_{t_k}, a_k) = h_{t_{k+1}}$. Define

$$\alpha^* = \limsup_{k \rightarrow \infty} \alpha_\xi(g_k) \quad \text{and} \quad \beta^* = \limsup_{k \rightarrow \infty} \beta_\xi(g_k).$$

Since the sequences $\alpha_\xi(g_k)$ and $\beta_\xi(g_k)$ only assume finitely many values, there exists a $K \in \mathbb{N}$ such that $\alpha^* \geq \alpha_\xi(g_k)$ and $\beta^* \geq \beta_\xi(g_k)$ for each $k \geq K$.

We show that $p \in P_\xi^E(g_K)$.

It holds that $u_i(p) \geq \alpha_\xi(g_k)$ for every $k \geq K$. To see this, notice that for each k there is a play $q_k > g_k$ such that $u_i(q_k) \geq \alpha_\xi(g_k)$. Such a play can be obtained as an outcome in the game $G_\xi(g_k)$ where Player I follows an optimal strategy. Notice that the sequence $\{q_k\}$ converges to p as k tends to infinity, and that upper semicontinuity of the function u_i implies that $u_i(p) \geq \alpha^*$. The claim follows.

Thus, it remains to be shown that $p \in P_\xi^F(g_K)$ with $F = E \setminus \{i\}$.

Notice that for each $k \in \mathbb{N}$, we have

$$P_\xi^F(g_k, a_k) \supseteq P_\xi^F(g_{k+1}). \quad (12)$$

To see this, recall that $(g_k, a_k) = h_{t_{k+1}}$ and that $g_{k+1} = h_{t_{k+1}}$. Among the histories $h_{t_{k+1}}, \dots, h_{t_{k+1}-1}$, none belongs to player i . Hence, for each $t \in \{t_k + 1, \dots, t_{k+1} - 1\}$, it holds that $h_{t+1} = (h_t, \sigma_\xi^E(h_t)) = (h_t, \sigma_\xi^F(h_t))$. Thus, Lemma 6.4 implies that $P_\xi^F(h_t) \supseteq P_\xi^F(h_{t+1})$. By unraveling the inclusions we obtain $P_\xi^F(h_{t_k+1}) \supseteq P_\xi^F(h_{t_{k+1}})$, which is exactly (12).

Define

$$B = \{q \in P: u_i(q) \geq \beta^*\}.$$

Notice that the set B is closed since u_i is upper semicontinuous. We next argue that

$$P_\xi^F(g_k) \cap B \supseteq P_\xi^F(g_k, a_k) \cap B, \quad \text{for each } k \geq K. \quad (13)$$

To see this, take a $k \geq K$ and a $q \in P_\xi^F(g_k, a_k) \cap B$. As $u_i(q) \geq \beta^* \geq \beta_\xi(g_k)$, and since $i(g_k) = i$ is not an element of F , we have $q \in P_\xi^F(g_k)$, as desired.

Combining (12) and (13), we obtain the inclusions

$$P_\xi^F(g_k) \cap B \supseteq P_\xi^F(g_{k+1}) \cap B, \quad \text{for each } k \geq K. \quad (14)$$

Finally, observe that $\beta_\xi(g_k) = \beta^*$ for infinitely many k , and that for each such k it holds that $P_\xi^F(g_k) \subseteq B$. As $P_\xi^F(g_k)$ is nonempty by the induction hypothesis, we conclude that $P_\xi^F(g_k) \cap B$ is nonempty for infinitely many $k \geq K$, and hence in view of (14) for each $k \geq K$. For each $k \geq K$, take an arbitrary element $q_k \in P_\xi^F(g_k) \cap B$. Then $q_k \in P_\xi^F(g_k) \cap B$ by (14). Since both $P_\xi^F(g_k)$ and B are closed sets and the sequence $\{q_k\}$ converges to p , we conclude that $p \in P_\xi^F(g_K) \cap B$. This completes the proof of the induction step. $\square$

Lemma 6.6. Let ξ be an ordinal such that the statements $\Phi(\xi)$ and $\Psi(\xi)$ are true. Then, for each $h \in H$, the play $\pi(\sigma_\xi^I, h)$ is an element of $P_\xi^I(h)$.

Proof. Let $p = \pi(\sigma_\xi^I, h)$ and let $h_0, h_1, \dots$ be the enumeration of the successive prefixes of p starting with $h_0 = h$. For $k \in \mathbb{N}$, let $I_k = \{i(h_0), \dots, i(h_k)\}$ and let $p_k = \pi(\sigma_\xi^{I_k}, h)$. By Lemma 6.5, $p_k \in P_\xi^{I_k}(h)$. Notice moreover that $h_k < p_k$ for each $k \in \mathbb{N}$. Hence, the sequence $\{p_k\}$ converges to the play p . Now fix an $m \in \mathbb{N}$. For each $k \geq m$ it holds that $I_k \supseteq I_m$, and hence $P_\xi^{I_k}(h) \subseteq P_\xi^{I_m}(h)$, so that $p_k \in P_\xi^{I_m}(h)$. As $P_\xi^{I_m}(h)$ is a closed set, we conclude that $p \in P_\xi^{I_m}(h)$. In particular, $u_{i(h_m)}(p) \geq \alpha_\xi(h_m)$. Since this is true for each $m \in \mathbb{N}$, it follows that $p \in P_\xi^I(h)$. $\square$

Lemma 6.7. Let ξ be an ordinal such that the statements $\Phi(\xi)$ and $\Psi(\xi)$ are true. Then, $\Phi(\xi + 1)$ and $\Psi(\xi + 1)$ are also true.

Proof. Recall that $P_\xi^1(h) = P_{\xi+1}(h)$ for each $h \in H$. Now the statement $\Phi(\xi + 1)$ follows by Lemma 6.6 and the statement $\Psi(\xi + 1)$ by Lemma 6.4. $\square$

Combining Lemmata 6.3 and 6.7, and the fact that $\Phi(0)$ and $\Psi(0)$ are true, we obtain the following result.

Lemma 6.8. For each ordinal ξ , the statements $\Phi(\xi)$ and $\Psi(\xi)$ are true.

Thus, in particular $P_{\xi^*}(\emptyset) \neq \emptyset$. Theorem 4.3 now follows from Theorem 4.1.

7. Special Cases

In this section we discuss a number of interesting special cases of our model. We assume throughout that each player's payoff function has finite range.

Games satisfying the one-deviation principle. We say that a strategy profile is *one-deviation immune* (ODI) if no player can improve his payoff by deviating at a single history. Formally, the strategy profile σ is ODI if for each $(h, a) \in H$ it holds that $u_{i(h)}(\sigma, h) \geq u_{i(h)}(\sigma, (h, a))$. Each SPE is clearly an ODI strategy profile, but in general the converse is false. A game is said to satisfy the *one-deviation principle* if the set of its SPE equals the set of ODI strategy profiles.

The importance of ODI strategy profiles for our work stems from the results of Flesch et al. [7], where the authors define an algorithm terminating with the set of plays that could be supported by an ODI strategy profile. This result thus provides a characterization of SPE plays in any game satisfying the one-deviation principle. Furthermore Flesch et al. prove⁶ existence of an ODI strategy profile under the assumption that the number of players is finite, and show, by means of an example in Section 4.3, that this assumption cannot be dropped.

The one-deviation principle is satisfied if each player i 's payoff function is *lower semicontinuous in player i 's own actions*. This means that $\liminf_{n \rightarrow \infty} u_i(p_n) \geq u_i(p)$ whenever $\{p_n\}$ is a sequence converging to p such that for each $n \in \mathbb{N}$ either $p_n = p_{n+1}$ or the longest common prefix of p_n and p_{n+1} is controlled by player i . In particular, player i 's payoff function is lower semicontinuous in player i 's own actions if (1) it is lower semicontinuous in the sense of Section 2, or (2) if for each play p the set of prefixes of p where i is active is finite. Thus the one-deviation principle is satisfied in a game where the player set is $I = \mathbb{N}$, and player t is active at each history of length t .

We presently turn to two important subclasses of games satisfying the one-shot deviation principle: the class of games with continuous payoff functions (falling under condition (1) above), and the class of centipede games where each player acts exactly once (falling under (2) above).

Games with continuous payoffs. Of particular importance for economic applications are games where the payoff functions are continuous. Any such game satisfies the one-deviation property. Notice that a function having a finite range is continuous if and only if each play p has a prefix h such that the function is constant of $P(h)$. Consequently, when the number of players is finite, one could prove existence of SPE by backward induction starting with subgames where every player's payoff function is constant.⁷ This technique runs into difficulties when the number of players is not finite, since then there might be no such subgame. However, whether the player set is finite or not, a slight modification of the algorithm in Flesch et al. [7] lends itself to a proof of the existence of SPE (see Flesch and Predtetchinski [11]).

We remark that payoff continuity, as defined here, is weaker⁸ than the condition of continuity at infinity as defined in Fudenberg and Levine [13] and in Carmona [5]. The two conditions are equivalent if the set of actions is finite.

Infinite centipede games. Cingiz et al. [6] consider infinite centipede games where each player acts exactly once. Assuming that the payoff functions are upper semicontinuous, the authors show that the game admits an SPE. In this setup the method of finite truncations of the game tree suffices.

Games with upper semicontinuous payoffs and finitely many actions. Finally, we consider the class of games where each player's payoff function is upper semicontinuous and the set of actions is finite. In this class, one could give an alternative proof of the existence of SPE, taking the concept of ODI as a starting point. While this result is subsumed by Theorem 4.3, we believe that the proof warrants some attention. The main idea is to consider ODI strategy profiles with respect to a certain refinement of the original payoff functions.

For simplicity, we assume that each player's payoff function takes integer values (recall our assumption that the range of the payoff function is finite). For each $m \in \mathbb{Z}$, let $V_i(m) = \{p \in P: u_i(p) \geq m\}$. We use the metric d on the set of plays as defined in Section 2. For a play $p \in P$ and a nonempty subset of plays $Q \subseteq P$ we let $d(p, Q)$ denote the distance from p to a set Q , i.e., the infimum of $d(p, q)$ over $q \in Q$ and we let $d(p, \emptyset) = 0$.

For a player $i \in I$ define

$$v_i(p) = u_i(p) - \frac{1}{2}d(p, V_i(u_i(p) + 1)). \quad (15)$$

The payoff function v_i could be thought of as a “refinement” of u_i that breaks indifferences by assigning a higher rank to plays that are closer to the set of plays with the higher utility. Indeed, notice that $v_i(p) < v_i(q)$ precisely when either (A) $u_i(p) < u_i(q)$ or (b) $u_i(p) = u_i(q) = m$ and $d(p, V_i(m + 1)) > d(q, V_i(m + 1))$.

Furthermore, v_i attains a minimum over any nonempty subset of plays, since u_i has a finite range of integer numbers, while d takes values in the set $\{0\} \cup \{2^{-n} : n \in \mathbb{N}\}$. Finally, v_i is upper semicontinuous. This follows from the upper semicontinuity of u_i and from the continuity of $d(p, Q)$ in p .

It is easy to prove that, in the game Ω' with the payoff functions $(v_i)_{i \in I}$, the algorithm in Flesch et al. [7] is well defined, and returns a nonempty subset of plays at every step. It is at this step of the proof that we need to assume the set of actions to be finite. Consequently, the game admits an ODI strategy profile. Moreover, each ODI strategy profile in Ω' is an SPE in the original game Ω , with the additional term involving the distance in Equation (15) making sure that deviations at infinitely many histories are not profitable.

8. On the Assumption of Finite Range

Until now we have performed our analysis under the assumption that every player's payoff function only takes finitely many values. In this section, we turn to a richer class of games, the games where the payoff function is only required to be bounded and Borel measurable. In Section 8.1 we reconsider the characterization of SPE plays, and in Section 8.2 we study discretizations of the payoff functions.

8.1. Extending the Characterization of SPE Plays

To accommodate for the payoff functions taking infinitely many values, we need a slight modification to our definitions in Section 3: replace the formula (2) by

$$\beta_\varepsilon(h) = \sup_{\lambda < \varepsilon} \alpha_\lambda(h).$$

Theorem 8.1. *Consider a perfect information game Ω . Assume that the payoff function of each player is bounded and Borel measurable. Then $P^*(\emptyset) \subseteq P_{\varepsilon^*}(\emptyset)$. Conversely, for each $p \in P_{\varepsilon^*}(\emptyset)$ and each $\varepsilon > 0$ there exists an ε -SPE σ such that $p = \pi(\sigma, \emptyset)$.*

The proof of the first claim of the theorem is similar to the first part of the proof of Theorem 4.1 in Section 5. The proof of the second claim is similar to the second part of the proof of Theorem 4.1, with the difference that one lets Σ_{Π}^ε be any Player II's strategy that is ε -optimal in every subgame of $G_{\varepsilon^*}(g)$, where $\varepsilon > 0$. Such a strategy exists in view of Lemma 2.1.

Theorem 8.1 asserts that $P_{\varepsilon^*}(\emptyset)$ contains all SPE plays and is contained in the set of ε -SPE plays for each $\varepsilon > 0$. The discrepancy between $P_{\varepsilon^*}(\emptyset)$ and $P^*(\emptyset)$ stems from the fact that in the game $G_{\varepsilon^*}(g)$ Player II might have no optimal strategy. This point is further illustrated by the example below. The game considered in the example has no SPE, so that $P^*(\emptyset) = \emptyset$, whereas $P_{\varepsilon^*}(h)$ is not empty for each history h .

Example 8.2. This is a modification of Example 3.3. There are two players, and the action set is $A = \{1, 2\}$. Player 1 starts the game. The active player decides who the next active player is by choosing the corresponding action. Let S_0 denote the set of plays where both actions 1 and 2 are played infinitely many times, S_1 the set of plays of the form $h1^\infty = (h, 1, 1, 1, \dots)$ for $h \in H$, and S_2 the set of plays of the form $h2^\infty = (h, 2, 2, 2, \dots)$ for $h \in H$.

The payoffs are $(-1, 2)$ for any play in S_1 , and $(-2, 1)$ for any play in S_2 . Player 2's payoff on S_0 is identically 0. Player 1's payoff function on S_0 is defined as follows. For a play $p = (a_0, a_1, \dots)$ in S_0 consider all blocks of consecutive 2's, and denote the lengths of these blocks by $m_0, m_1, \dots$. Thus for instance, if $p = (1, 2, 2, 2, 2, 1, 2, 2, 1, \dots)$ then $m_0 = 4$, $m_1 = 2$, etc. If the sequence $\{m_n\}_{n \in \mathbb{N}}$ is not eventually constant we let $u_1(p) = 0$. If it is eventually constant we let $u_1(p) = -1 + 2^{-m}$ where $m_n = m$ for n large.

The essential feature of this game is that within the set S_0 player 2 can force player 1's payoff arbitrarily close to -1 . To do so, every time that player 2 becomes active, he should take action 2 exactly m times before playing action 1, where m is sufficiently large. Notice however that player 2 cannot make sure that 1's payoff is exactly -1 .

The values $P_\varepsilon(h)$ and $\alpha_\varepsilon(h)$ are given by the table below, where $S_1(h) = S_1 \cap P(h)$, and $h2^\infty$ denotes the play $(h, 2, 2, \dots)$. The algorithm terminates after three rounds, with $P_{\varepsilon^*}(h) = S_1(h)$ for each history h .

For Steps 0, 1, and 2 the computation is similar to that in Example 3.3.

Table 3. The iterations in Example 8.2.

ξ	$P_\xi(h)$		$\alpha_\xi(h)$	
	$\iota(h) = 1$	$\iota(h) = 2$	$\iota(h) = 1$	$\iota(h) = 2$
0	$P(h)$	$P(h)$	-2	0
1	$P(h)$	$P(h)$	-1	1
2	$S_1(h)$	$\{h2^\infty\} \cup S_1(h)$	-1	2
3	$S_1(h)$	$S_1(h)$	-1	2

Step 3. We obtain $P_3(h)$ exactly as in Example 3.3. Consider player 1's history h and the corresponding game $G_3(h)$. Player II, for every $m \in \mathbb{N}$, has a strategy that ensures that Player I's payoff is not greater than $-1 + 2^{-m}$: namely, proposing at stage k the play $p_k = (h_k, a_k)2^m 1^\infty$, i.e., the play where after (h_k, a_k) there follows m times action 2, after which only action 1 is played. Consequently $\alpha_3(h) = -1$. Notice that Player II has no optimal strategy.

We thus conclude that $\xi^* = 3$ and $P_{\xi^*}(h) = S_1(h)$ for all $h \in H$.

The game has no SPE. For suppose to the contrary and let σ be an SPE. Applying Theorem 8.1 to the subgame of Ω starting at history h , we deduce that $\pi(\sigma, h)$ is an element of $P_{\xi^*}(h) = S_1(h)$ for each history h . This implies that under the strategy σ_2 , player 2 eventually plays action 1. Let η_1 be player 1's strategy whereby action 2 is played at each of player 1's histories. The profile of strategies (η_1, σ_2) induces a play in S_0 , and hence gives player 1 a payoff strictly greater than -1 . We conclude that a deviation to η_1 from σ_1 is profitable, contradicting the equilibrium hypothesis. $\triangleleft$

8.2. On Discretizations of Payoff Functions

This section is motivated by the technique of discretizing payoff functions: approximating a given bounded payoff function by a function with finite range.

Let $\epsilon > 0$. Consider a payoff function $u_i: P \rightarrow \mathbb{R}$. A function $v_i: P \rightarrow \mathbb{R}$ is said to be an ϵ -discretization of u_i if the function v_i only takes finitely many values and for each $p \in P$ we have $|u_i(p) - v_i(p)| \leq \epsilon$. A family $v = (v_i)_{i \in I}$ of payoff functions is an ϵ -discretization of the family $u = (u_i)_{i \in I}$ if for each $i \in I$ the function v_i is an ϵ -discretization of u_i . We let $\Omega(u)$ and $\Omega(v)$ be the games with the family of payoff functions u , respectively v .

The result below could be interpreted as saying that passing to discretizations is without loss of generality, at least if one is only interested in the question of existence of subgame-perfect ϵ -equilibrium for every positive ϵ .

Theorem 8.3. Consider a family of bounded Borel measurable payoff functions $u = (u_i)_{i \in I}$. The following conditions are equivalent:

- (a) For each $\epsilon > 0$ the game $\Omega(u)$ has an ϵ -SPE.
- (b) For each $\epsilon > 0$ there exists a Borel measurable ϵ -discretization v of u such that the game $\Omega(v)$ has an SPE.

The fact that condition (a) implies condition (b) is straightforward: if v is an ϵ -discretization of u , each SPE of $\Omega(v)$ is 2ϵ -SPE of $\Omega(u)$. We prove the converse implication.

Proof. Fix a subgame-perfect $\epsilon/2$ -equilibrium σ of the game $\Omega(u)$. We construct a Borel measurable ϵ -discretization v of u such that σ is an SPE of the game $\Omega(v)$.

For a player $i \in I$ define the binary relation $\triangleleft_i$ on P as follows: $p \triangleleft_i q$ if there exists a history $h \in H$ with $\iota(h) = i$ and a strategy η_i for player i such that $\sigma_i(h) \neq \eta_i(h)$, $p = \pi(\sigma, h)$, and $q = \pi(\sigma/\eta_i, h)$. Roughly speaking $p \triangleleft_i q$ means that player i can induce a deviation from p to q .

Claim 1. The binary relation $\triangleleft$ is transitive and well founded.

Proof of Claim 1. We omit the subscript i for simplicity.

To prove transitivity of $\triangleleft$, take the plays $p, q, r \in P$ such that $p \triangleleft q$ and $q \triangleleft r$. It is clear that $p \neq q$ and $q \neq r$. Let the history h and the strategy η_i witness the relation $p \triangleleft q$, and let the history h' and the strategy η'_i witness $q \triangleleft r$. Both h and h' are then prefixes of the play q . Hence either $h' \leq h$ or $h < h'$. If $h' \leq h < q = \pi(\sigma, h')$, we have $\pi(\sigma, h') = \pi(\sigma, h)$, implying that $q = p$, a contradiction. Hence $h < h'$. Now let η''_i be the strategy that agrees with η_i on each history g such that $g < h'$, and agrees with η'_i on all other histories. Then $r = \pi(\sigma/\eta''_i, h)$ and $\eta''_i(h) = \eta'_i(h) \neq \sigma_i(h)$, implying that $p \triangleleft r$.

Suppose that there is an infinite $\triangleleft$ -descending sequence of plays, that is, a sequence $\{p_n\}_{n \in \mathbb{N}}$ such that $p_{n+1} \triangleleft p_n$ for each $n \in \mathbb{N}$. For each $n \in \mathbb{N}$ let h_n be the history and η^n_i be player i 's strategy that witness the relation $p_{n+1} \triangleleft p_n$. Using the fact that $p_{n+2} \triangleleft p_{n+1}$ and $p_{n+1} \triangleleft p_n$, one can argue exactly as in the previous paragraph

that $h_{n+1} < h_n$. Since this holds for each $n \in \mathbb{N}$, we obtain a $<$ -decreasing sequence of histories $\{h_n\}_{n \in \mathbb{N}}$, which is clearly impossible. $\square$

Define the functions

$$\begin{aligned} u_i^-(p) &= \max\{n\epsilon: n \in \mathbb{Z}, n\epsilon \leq u_i(p)\} \\ u_i^+(p) &= \min\{n\epsilon: n \in \mathbb{Z}, u_i(p) \leq n\epsilon\}. \end{aligned}$$

It is not difficult to see that u_i^- and u_i^+ are both Borel measurable ϵ -discretizations of u_i . In particular, they take only finitely many values because the function u_i is bounded.

We define the function v_i . Let $S = \{\pi(\sigma, h): h \in H, \iota(h) = i\}$. On the complement of the set S we let v_i be identically equal to u_i^- . On the set S we define v_i by recursion on the relation $\triangleleft$ (see Kunen [17], p. 103, Theorem 5.6), as follows:

$$v_i(q) = \begin{cases} u_i^-(q) & \text{if there exists a } p \in S \text{ with } p \triangleleft_i q, \text{ such that } v_i(p) < u_i^+(q), \\ u_i^+(q) & \text{otherwise.} \end{cases} \quad (16)$$

Claim 2. The function v_i is a Borel measurable ϵ -discretization of u_i .

Proof of Claim 2. The fact that v_i is an ϵ -discretization of u_i follows since both u_i^- and u_i^+ are ϵ -discretizations of u_i . We show that v_i is Borel measurable.

Recall that a subset D of P is said to be discrete if each point of D is isolated, or equivalently, if each point $p \in D$ has a prefix h such that $P(h) \cap D = \{p\}$. A subset D of P is said to be sigma-discrete if it is a countable union of discrete sets. Each discrete set and consequently each sigma-discrete set is Borel. Moreover, each subset of a sigma-discrete set is sigma-discrete. It follows that any function on a sigma-discrete set is Borel measurable.

We argue that S is a sigma-discrete set. To see this, for each $k \in \mathbb{N}$ let $S_k = \{\pi(\sigma, h): h \in H, \iota(h) = i, h \text{ has length } k\}$. The set S_k is discrete. Indeed, if h is the prefix of a play $p \in S_k$ of length k , then $P(h) \cap S_k = \{p\}$. It follows that $S = \bigcup_{k \in \mathbb{N}} S_k$ is a sigma-discrete set.

Thus, the restriction of v_i to S is automatically Borel measurable. The restriction of v_i to the complement of S is Borel measurable since it coincides with the function u_i^- on that set. We conclude that v_i is Borel measurable. $\square$

Since σ is a subgame-perfect $\epsilon/2$ -equilibrium of the game $\Omega(u)$, we know that for every $p, q \in P$ and every $i \in I$ if $p \triangleleft_i q$ then $u_i(q) \leq u_i(p) + \epsilon/2$. Similarly, to show that σ is a subgame-perfect 0-equilibrium of the game $\Omega(v)$, it suffices to prove that for every $p, q \in P$ and every $i \in I$ if $p \triangleleft_i q$ then $v_i(q) \leq v_i(p)$.

Claim 3. The strategy profile σ is an SPE of the game $\Omega(v)$.

Proof of Claim 3. Fix a player $i \in I$ and a play $q \in P$. Let $E = \{p \in P: p \triangleleft_i q\}$. Notice that $E \subseteq S$. For a play $p \in E$ let $\Theta(p)$ denote the formula $v_i(q) \leq v_i(p)$. We prove that $\Theta(p)$ holds for each $p \in E$ by induction on $\triangleleft_i$.

Take a play $p \in E$. Suppose that $\Theta(r)$ holds for every $r \in E$ with $r \triangleleft_i p$. We now show that $\Theta(p)$ also holds. Suppose $\Theta(p)$ is false, that is, $v_i(p) < v_i(q)$. Since both $v_i(p)$ and $v_i(q)$ are multiples of ϵ , we have $v_i(p) \leq v_i(q) - \epsilon$.

We have $v_i(q) = u_i^-(q)$. Indeed, if $q \in P \setminus S$, this follows automatically by the definition of v_i , and if $q \in S$, this follows by (16) since $v_i(p) < v_i(q)$ implies that $v_i(p) < u_i^+(q)$.

Furthermore, $v_i(p) = u_i^+(p)$. To prove the equality, we have to argue that $u_i^+(p) \leq v_i(r)$ for each $r \in S$ with $r \triangleleft_i p$. Thus take any such r . We have $r \in E$, and $\Theta(r)$ implies that $v_i(q) \leq v_i(r)$. Since by assumption $v_i(p) < v_i(q)$, we have $v_i(p) < v_i(r)$ and consequently $u_i^+(p) \leq v_i(r)$, as desired.

We obtain the inequalities

$$u_i(p) \leq u_i^+(p) = v_i(p) \leq v_i(q) - \epsilon = u_i^-(q) - \epsilon \leq u_i(q) - \epsilon,$$

contradicting the fact that σ is a subgame-perfect $\epsilon/2$ -equilibrium of $\Omega(u)$. $\square$

This completes the proof of Theorem 8.3. $\square$

Endnotes

- ¹ Thus H is a pruned tree and P is the set of branches of H .
- ² Suppose first that H is uncountable. Let H_k denote the set of histories $h \in H$ of length k . Since $H = \bigcup_{k \in \mathbb{N}} H_k$, there is some $k \in \mathbb{N}$ such that H_k is uncountable. Consider the cylinder sets $P(h)$, for $h \in H_k$. They are nonempty, mutually disjoint, and there are uncountably many of them, so $(P, \mathcal{T})$ is not separable. Now suppose that H is countable. For every history $h \in H$, take an arbitrary play $p_h \in P(h)$. Then, the set $\{p_h \mid h \in H\}$ is countable and dense in $(P, \mathcal{T})$.
- ³ Throughout the paper we maintain the following convention: The players of the original game Ω are referred to as players with a lower case p . The two players of the auxiliary zero-sum games $G_\varepsilon(h)$ are referred to as Player I and Player II with a capital P .
- ⁴ Consider P_ε as a function from H to 2^P and take any ordinal ξ_* with cardinality larger than that of the set of such functions.
- ⁵ We remark that a much stronger result is proven in Flesch et al.: for small but positive ε , the game has no ε -SPE, even if one allows for mixed strategies.
- ⁶ The result is not explicitly stated in the paper. However, most of the proof of Theorem 2.3 is in fact the proof of the existence of an ODI strategy profile.
- ⁷ More precisely, consider the set E of nonterminal histories h , i.e., histories h for which at least one player's payoff function is not constant on $P(h)$. The usual $<$ relation on E is well founded. Backward induction could therefore be formalized as a recursion on the relation $<$.
- ⁸ 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.

References

- [1] Alós-Ferrer C, Ritzberger K (2016) Equilibrium existence for large perfect information games. *J. Math. Econom.* 62:5–18.
- [2] Alós-Ferrer C, Ritzberger K (2017) Characterizing existence of equilibrium for large extensive form games: A necessity result. *Econom. Theory* 63(2):407–430.
- [3] Balbus L, Jaśkiewicz A, Nowak AS (2014) Existence of stationary Markov perfect equilibria in stochastic altruistic growth economies. *J. Optim. Theory Appl.* 165(1):295–315.
- [4] Brihaye T, Bruyère V, Meunier N, Raskin J-F (2015) Weak subgame perfect equilibria and their application to quantitative reachability. *Proc. 24th EACSL Annual Conf. Comput. Sci. Logic* 504–518.
- [5] Carmona G (2005) On games of perfect information: Equilibria, ε -equilibria and approximation by simple games. *Internat. Game Theory Rev.* 7(4):491–499.
- [6] Cingiz K, Flesch J, Herings JJ, Predtetchinski A (2016) Doing it now, later, or never. *Games Econom. Behav.* 97:174–185.
- [7] Flesch J, Kuipers J, Mashiah-Yaakovi A, Schoenmakers G, Solan E, Vrieze K (2010a) Perfect-information games with lower semicontinuous payoffs. *Math. Oper. Res.* 35(4):742–755.
- [8] Flesch J, Kuipers J, Schoenmakers G, Vrieze K (2010b) Subgame-perfection in positive recursive games with perfect information. *Math. Oper. Res.* 35(1):193–207.
- [9] Flesch J, Kuipers J, Mashiah-Yaakovi A, Shmaya E, Schoenmakers G, Solan E, Vrieze K (2014) Non-existence of subgame-perfect epsilon-equilibrium in perfect information games with infinite horizon. *Internat. J. Game Theory* 43(4):945–951.
- [10] Flesch J, Predtetchinski A (2016a) Subgame-perfect ε -equilibria in perfect information games with common preferences at the limit. *Math. Oper. Res.* 41(4):1208–1221.
- [11] Flesch J, Predtetchinski A (2016b) On refinements of subgame-perfect ε -equilibrium. *Internat. J. Game Theory* 45(3):523–542.
- [12] Flesch J, Kuipers J, Schoenmakers G, Vrieze K (2013) Subgame-perfection in free transition games. *Eur. J. Oper. Res.* 228:201–207.
- [13] Fudenberg D, Levine D (1983) Subgame-perfect equilibria of finite and infinite-horizon games. *J. Econom. Theory* 31(2):251–268.
- [14] Harris C (1985a) Existence and characterization of perfect equilibrium in games of perfect information. *Econometrica* 53(3):613–628.
- [15] Harris C (1985b) A Characterisation of the perfect equilibria of infinite horizon games. *J. Econom. Theory* 37(1):99–125.
- [16] Kuipers J, Flesch J, Schoenmakers G, Vrieze K (2016) Subgame-perfection in perfect information games where each player controls one state. *Internat. J. Game Theory* 45(1):205–237.
- [17] Kunen K (1980) *Set Theory: An Introduction to Independence Proofs* (North-Holland, Amsterdam).
- [18] Le Roux S (2016) Infinite subgame perfect equilibrium in the Hausdorff difference hierarchy. Hajiaghayi TM, Mousavi RM, eds. *Topics in Theoretical Computer Science* (Springer International, Cham, Switzerland), 147–163.
- [19] Le Roux S, Pauly A (2014) Infinite sequential games with real-valued payoffs. Henzinger TA, Miller D, eds. *Proc. Joint Meeting Twenty-Third EACSL Annual Conf. Comput. Sci. Logic (CSL) and Twenty-Ninth Annual ACM/IEEE Sympos. Logic Comput. Sci. (LICS), CSL-LICS '14*, Vol. 62 (ACM, New York), 1–10.
- [20] Martin DA (1975) Borel determinacy. *Ann. Math.* 102(2):363–371.
- [21] Mashiah-Yaakovi A (2014) Subgame perfect equilibria in stopping games. *Internat. J. Game Theory* 43(1):89–135.
- [22] Mertens J-F (1987) Repeated Games. *Proc. Internat. Congress Mathematicians*, Vol. 986 (AMS, Providence, RI), 1528–1577.
- [23] Meunier N (2016) Multi-Player quantitative games: Equilibria and algorithms. PhD thesis, University of Mons, Mons, Belgium.
- [24] Purves RA, Sudderth WD (2011) Perfect information games with upper semicontinuous payoffs. *Math. Oper. Res.* 36(3):468–473.
- [25] Solan E (2005) Subgame-perfection in quitting games with perfect information. *Math. Oper. Res.* 30(1):51–72.
- [26] Solan E, Vieille N (2003) Deterministic multi-player Dynkin games. *J. Math. Econom.* 39(8):911–929.