Perfect-information games with lower-semicontinuous payoffs

Citation for published version (APA):

Flesch, J., Kuipers, J., Mashiah-Yaakovi, A., Schoenmakers, G., Solan, E., & Vrieze, K. (2010). Perfect-information games with lower-semicontinuous payoffs. *Mathematics of Operations Research*, 35(4), 742-755. <https://doi.org/10.1287/moor.1100.0469>

Document status and date:

Published: 01/11/2010

DOI:

[10.1287/moor.1100.0469](https://doi.org/10.1287/moor.1100.0469)

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.

Mathematics of Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Perfect-Information Games with Lower-Semicontinuous Payoffs

János Flesch<http://www.personeel.unimaas.nl/j-flesch/>, Jeroen Kuipers<http://www.maastrichtuniversity.nl/web/Schools/DKE/Thema/Staff/JeroenKuipers.htm>, Ayala Mashiah-Yaakovi<http://www.math.tau.ac.il/~ayalam/>, Gijs Schoenmakers<http://www.maastrichtuniversity.nl/web/Schools/DKE/Thema/Staff/GijsSchoenmakers.htm>, Eilon Solan<http://www.math.tau.ac.il/~eilons/>, Koos Vrieze,

To cite this article:

János Flesch<http://www.personeel.unimaas.nl/j-flesch/>, Jeroen Kuipers<http://www.maastrichtuniversity.nl/web/Schools/DKE/Thema/Staff/JeroenKuipers.htm>, Ayala Mashiah-Yaakovi<http://www.math.tau.ac.il/~ayalam/>, Gijs Schoenmakers<http://www.maastrichtuniversity.nl/web/Schools/DKE/Thema/Staff/GijsSchoenmakers.htm>, Eilon Solan<http://www.math.tau.ac.il/~eilons/>, Koos Vrieze, (2010) Perfect-Information Games with Lower-Semicontinuous Payoffs. *Mathematics of Operations Research* 35(4):742-755. <https://doi.org/10.1287/moor.1100.0469>

Full terms and conditions of use: <https://pubsonline.informs.org/Publications/Librarians-Portal/PubsOnLine-Terms-and-Conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2010, INFORMS

Please scroll down for article—it is on subsequent pages

With 12,500 members from nearly 90 countries, INFORMS is the largest international association of operations research (O.R.) and analytics professionals and students. INFORMS provides unique networking and learning opportunities for individual professionals, and organizations of all types and sizes, to better understand and use O.R. and analytics tools and methods to transform strategic visions and achieve better outcomes.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Perfect-Information Games with Lower-Semicontinuous Payoffs

János Flesch

Department of Quantitative Economics, Maastricht University, PO Box 616, 6200 MD Maastricht, The Netherlands,
j.flesch@maastrichtuniversity.nl, <http://www.personeel.unimaas.nl/j-flesch/>

Jeroen Kuipers

Department of Knowledge Engineering, Maastricht University, PO Box 616, 6200 MD Maastricht, The Netherlands,
kuipers@maastrichtuniversity.nl, <http://www.maastrichtuniversity.nl/web/Schools/DKE/Thema/Staff/JeroenKuipers.htm>

Ayala Mashiah-Yaakovi

School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel,
ayalam@post.tau.ac.il, <http://www.math.tau.ac.il/~ayalam/>

Gijs Schoenmakers

Department of Knowledge Engineering, Maastricht University, PO Box 616, 6200 MD Maastricht, The Netherlands, gm.schoenmakers@maastrichtuniversity.nl,
<http://www.maastrichtuniversity.nl/web/Schools/DKE/Thema/Staff/GijsSchoenmakers.htm>

Eilon Solan

School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel,
eilons@post.tau.ac.il, <http://www.math.tau.ac.il/~eilons/>

Koos Vrieze

Department of Knowledge Engineering, Maastricht University, PO Box 616, 6200 MD Maastricht, The Netherlands,
koos@ictvrieze.nl

We prove that every multiplayer perfect-information game with bounded and lower-semicontinuous payoffs admits a subgame-perfect ε -equilibrium in pure strategies. This result complements Example 3 in Solan and Vieille [Solan, E., N. Vieille. 2003. Deterministic multi-player Dynkin games. *J. Math. Econom.* **39** 911–929], which shows that a subgame-perfect ε -equilibrium in pure strategies need not exist when the payoffs are not lower-semicontinuous. In addition, if the range of payoffs is finite, we characterize in the form of a Folk Theorem the set of all plays and payoffs that are induced by subgame-perfect 0-equilibria in pure strategies.

Key words: perfect information; subgame-perfect equilibrium; lower-semicontinuous payoffs.

MSC2000 subject classification: Primary: 91A06, 91A10; secondary: 91A44.

OR/MS subject classification: Primary: games/group decisions; secondary: noncooperative.

History: Received March 16, 2010; revised September 11, 2010.

1. Introduction. A multiplayer perfect-information game is a sequential game with perfect information and without chance moves. The payoff of each player is a function of the infinite sequence of actions that the players choose. Perfect-information games were introduced by Gale and Stewart [3], who studied two-player zero-sum games where the payoff function is the indicator of some set. In other words, player 1 wins if the play generated by the players is in a given set of plays, and player 2 wins otherwise. Martin [10] proved that if the winning set of player 1 is Borel measurable, then the game is determined: either player 1 has a winning strategy or player 2 has a winning strategy. This result implies that every two-player zero-sum perfect-information game has a value, provided the payoff function is bounded and Borel measurable.

Mertens and Neyman (see Mertens [12]) used the existence of the value in two-player zero-sum perfect-information games to prove that for every $\varepsilon > 0$, every multiplayer non-zero-sum perfect-information game has a pure ε -equilibrium, provided the payoff functions are bounded and Borel measurable. Roughly, the ε -equilibrium strategies constructed by Mertens and Neyman are as follows: each player i starts by following an $\varepsilon/2$ -optimal strategy in an auxiliary two-player zero-sum game G_i , where the payoff is that of player i , player i is the maximizer, and the other players try to minimize player i 's payoff. This goes on as long as no player deviates. Once some player, say, player i , deviates, the other players switch to an $\varepsilon/2$ -optimal strategy of the minimizers in the game G_i .

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

To deal with such noncredible threats of punishment, Selten [16, 17] introduced the concept of subgame-perfect equilibrium. A strategy vector is a subgame-perfect ε -equilibrium if it induces an ε -equilibrium after any possible finite history of actions. Ummels [22] proved the existence of a subgame-perfect equilibrium in pure

strategies for multiplayer perfect-information games when the payoff function of each player is the indicator of some Borel set (for a more general result, see Grädel and Ummels [5]). His proof is based on the following recursive construction. First, one identifies all finite histories that are a winning position to at least one of the players; that is, if this finite history occurs, one of the players can ensure that his payoff is 1. After such finite histories, one instructs every winning player to play a winning strategy. This leads to a pruned game where all moves that are excluded by these winning strategies are eliminated. One subsequently identifies winning positions to the players in this new game and prunes it in a similar way. The process repeats itself, until it reaches a stable state. Ummels proves that a combination of remaining strategies is a subgame-perfect equilibrium of the original game.

In this paper we show that every multiplayer perfect-information game with bounded and lower-semicontinuous payoffs admits a subgame-perfect ε -equilibrium in pure strategies, for every $\varepsilon > 0$. This result complements Example 3 in Solan and Vieille [20], which shows that when the payoff function of at least one player is not lower-semicontinuous, a subgame-perfect ε -equilibrium in pure strategies need not exist.¹ Our proof makes use of transfinite induction; as we use the axiom of choice, our proof is valid within the ZFC framework. A different type of transfinite construction was used by Maitra and Sudderth [8] to prove the existence of the value in a certain class of stochastic games. In §4.2 we point at another possible application of our technique.

The determinacy of perfect-information games has attracted a lot of attention in descriptive set theory (see, e.g., Schilling and Vaught [15] and Kechris [6]). A rich literature identifies winning positions for the two players in the class of games that are played on graphs (see Grädel [4] for a survey). Two-player zero-sum perfect-information games were used in the computer science literature to study reactive nonterminating programs (see, e.g., Thomas [21]) and model checking in μ -calculus (see, e.g., Emerson et al. [1]) and in economics to show that measurable tests are manipulable (Shmaya [18]).

Our result also relates to the game theoretic literature that studies the existence of a subgame-perfect ε -equilibrium in various classes of infinite games, see, e.g., Mertens and Parthasarathy [13], Solan and Vieille [20], Maitra and Sudderth [9], Mashiah-Yaakovi [11], Kuipers et al. [7], or Flesch et al. [2]. In particular, our result generalizes some of the results in Flesch et al. [2]. Recently, Purves and Sudderth [14], using ideas different from ours, complemented our result by showing that a subgame perfect ε -equilibrium exists if the payoff processes are bounded and upper-semicontinuous.

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

2. The model and the main result

DEFINITION 2.1. An n -player *perfect-information game* is a quadruple $(I, A, i, (u^j)_{j \in I})$ where $I = \{1, 2, \dots, n\}$ is the set of players; A is a nonempty set² of actions; $i: \bigcup_{t \in \mathbb{N}} A^{t-1} \rightarrow I$ is a function³ that assigns an active player to each finite sequence of actions; and $u^j: A^\mathbb{N} \rightarrow \mathbf{R}$ is the payoff function, for every player $j \in I$.

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

REMARK 2.1. The assumption that the action set is the same for all players and for all stages is made for simplicity of notations only. Nothing that is said below would be affected if the action sets were to depend on the player, on the stage, or even on the whole past play.

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

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

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

DEFINITION 2.2. A (*pure*) *strategy* for player j is a function $\sigma^j: \mathcal{H}^j \rightarrow A$. A (*pure*) *strategy profile* is a vector of strategies $\sigma = (\sigma^j)_{j \in I}$.

¹ The game presented in Solan and Vieille [20] is the following two-player perfect-information game in which the players play alternately. The set of actions of each player is $A = \{c, s\}$. If both players always choose c , the payoff to each player is 0. Otherwise, let k be the first player who plays action s . If $k = 1$, the payoff vector is $(-1, 2)$, while if $k = 2$ the payoff vector is $(-2, 1)$.

² The set of actions A may be finite or infinite.

³ By convention, the initial history is the empty history $h_1 = \emptyset$, and $A^0 = \{\emptyset\}$.

In the present paper we discuss only pure strategies, and by a strategy or by a strategy profile we will always mean a pure one. Note that there are no measurability considerations in the definition of a strategy. We denote by Σ^j the strategy space of player j , and by $\Sigma := \times_{j \in I} \Sigma^j$ the set of all strategy profiles.

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

$$a_t := \sigma^{i(h_t)}(h_t), \quad \text{where } h_t := (a_1, a_2, \dots, a_{t-1}), \quad \forall t \in \mathbb{N}.$$

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

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

DEFINITION 2.3. Let $\varepsilon \geq 0$. A strategy profile $\sigma_* = (\sigma_*^j)_{j \in I}$ is an ε -equilibrium if $u^j(\sigma_*) \geq u^j(\sigma_*^{-j}, \sigma^j) - \varepsilon$ for every player $j \in I$ and every strategy $\sigma^j \in \Sigma^j$.

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

A two-player, perfect-information game is called *zero-sum* if $u^1(p) + u^2(p) = 0$ for every $p \in A^{\mathbb{N}}$. The result of Martin [10] implies that in zero-sum games, an ε -equilibrium exists for every $\varepsilon > 0$ under quite general conditions.

THEOREM 2.1. *If the game is zero-sum, and if u^1 is bounded and Borel measurable, then an ε -equilibrium exists for every $\varepsilon > 0$.*

This result implies the existence of an ε -equilibrium in every multiplayer perfect-information game with bounded and Borel measurable payoffs.

THEOREM 2.2 (MERTENS AND NEYMAN; SEE MERTENS [12]). *If u^j is bounded and Borel measurable for every player $j \in I$, then an ε -equilibrium exists for every $\varepsilon > 0$.*

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

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

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

DEFINITION 2.4. Let $\varepsilon \geq 0$. A strategy profile $\sigma_* = (\sigma_*^j)_{j \in I}$ is a *subgame-perfect ε -equilibrium* if for every finite history $h \in \mathcal{H}$, one has

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

In other words, a strategy profile is a subgame perfect ε -equilibrium if it induces an ε -equilibrium in all subgames. Here, a subgame is a game played after a finite history h with payoff function $u^j(\cdot | h)$ for player $j \in I$.

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

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

DEFINITION 2.5. The payoff function u^j is *lower-semicontinuous* if for every sequence $(p^k)_{k \in \mathbb{N}}$ of plays in $A^{\mathbb{N}}$ that converges to a limit p one has

$$\liminf_{k \rightarrow \infty} u^j(p^k) \geq u^j(p). \quad (1)$$

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

THEOREM 2.3. *If the payoff function u^j is bounded and lower-semicontinuous for every player $j \in I$, then the game admits a subgame-perfect ε -equilibrium in pure strategies for every $\varepsilon > 0$.*

This result is tight, in the sense that if the payoff function of one of the players is not lower-semicontinuous, then the game need not admit a subgame-perfect ε -equilibrium for every $\varepsilon > 0$ (see Solan and Vieille [20, Example 3] or Footnote 1).

Theorem 2.3 was recently complemented by Purves and Sudderth [14], who proved that the statement remains valid if lower-semicontinuity is replaced by upper-semicontinuity (i.e., when the inequality in (1) is reversed for every player j).

3. Proof of Theorem 2.3 and a Folk Theorem. We first argue that we can assume w.l.o.g. that the range of the payoff functions $(u^j)_{j \in I}$ is finite. Indeed,⁴ let $\hat{u}^j(p)$ be the highest multiple of ε that is *strictly* smaller than $u^j(p)$:

$$\hat{u}^j(p) := \varepsilon \left\lfloor \frac{u^j(p)}{\varepsilon} \right\rfloor.$$

Note that if u^j is bounded, then $\hat{u}^j$ has finite range, and if u^j is lower-semicontinuous, then so is $\hat{u}^j$. Moreover, every subgame-perfect ε -equilibrium in the game with payoff functions $(\hat{u}^j)_{j \in I}$ is a subgame-perfect 2ε -equilibrium in the game with payoff functions $(u^j)_{j \in I}$. Therefore, for the proof of Theorem 2.3, we may assume w.l.o.g. that the payoff functions have finite range.

From now on, we assume that the payoff functions $(u^j)_{j \in I}$ have finite range and are lower-semicontinuous. Under these assumptions we will prove the existence of a subgame-perfect 0-equilibrium. In the proof, we use the finiteness of the range of the payoffs to have a maximal payoff and a minimal payoff in every nonempty subset of payoffs. The lower-semicontinuity of the payoff functions, on the other hand, will be used to obtain the following property: when the players are supposed to play according to a strategy profile $\sigma = (\sigma^\ell)_{\ell \in I}$, if some player j cannot deviate profitably by not playing the action prescribed by σ^j finitely many times, then he cannot deviate profitably by disobeying σ^j infinitely many times either.

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

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

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

$$P_0(h) := \{p \in A^N : h < p\}, \quad (2)$$

$$\alpha_0(h) := \min_{p \in P_0(h)} u^{i(h)}(p). \quad (3)$$

The set $P_0(h)$ consists of all plays that extend h , and the quantity $\alpha_0(h)$ is a naive lower bound to the set of subgame-perfect 0-equilibrium payoffs in the subgame that starts at h . The minimum in (3) exists because the payoff functions have finite range.

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

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

$$\alpha_{\xi+1}(h) := \max_{a \in A} \min_{p \in P_\xi(h, a)} u^{i(h)}(p), \quad (4)$$

$$P_{\xi+1}(h) := \left\{ p \in \bigcup_{a \in A} P_\xi(h, a) : u^{i(h)}(p) \geq \alpha_{\xi+1}(h) \right\}. \quad (5)$$

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

⁴ For every real number x , we denote by $\lfloor x \rfloor$ the largest integer that is *strictly* smaller than x .

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

$$P_\xi(h) := \bigcap_{\lambda < \xi} P_\lambda(h), \quad (6)$$

$$\alpha_\xi(h) := \min_{p \in P_\xi(h)} u^{i(h)}(p). \quad (7)$$

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

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

LEMMA 3.1. *Let ξ be an ordinal and let h be a finite history.*

- (i) $\alpha_\xi(h) = \min_{p \in P_\xi(h)} u^{i(h)}(p)$. In particular, there is a play $p \in P_\xi(h)$ such that $u^{i(h)}(p) = \alpha_\xi(h)$.
- (ii) Let $a \in A$ be an action that achieves the maximum in the right-hand side of (4). If $p \in P_\xi(h, a)$, then $p \in P_{\xi+1}(h)$.
- (iii) Suppose that $\xi = 0$ or ξ is a limit ordinal. Let $p \in P_\xi(h)$ and let h' be a finite history that satisfies $h \preceq h' < p$. Then $p \in P_\xi(h')$.

PROOF. Part (i) follows from the definition (3) for $\xi = 0$, from the definitions (4) and (5) for successor ordinals, and from definition (7) for limit ordinals.

Part (ii) follows from the definitions (4) and (5).

We now prove Part (iii). For $\xi = 0$, the claim follows from definition (2). So, assume that ξ is a limit ordinal and let $p \in P_\xi(h)$. We will show that $p \in P_\xi(h, a)$, where $a \in A$ is the action in p right after h . Let λ be an ordinal such that $\lambda < \xi$. As $\lambda + 1 < \xi$, it follows from definition (6) that $p \in P_{\lambda+1}(h)$. Hence, by definition (5), we have $p \in P_\lambda(h, a)$. Because $\lambda < \xi$ was arbitrary, it follows that $p \in P_\xi(h, a)$. The proof for any finite history h' that satisfies $h < h' < p$ follows by induction. $\square$

The following theorem states additional properties of the sequences $(\alpha_\xi(h))_\xi$ and $(P_\xi(h))_\xi$, which play a crucial role in the proof of Theorem 2.3.

THEOREM 3.1. *The following holds for every $h \in \mathcal{H}$:*

- (i) *The set $P_\xi(h)$ is not empty for every ordinal ξ .*
- (ii) *The sequence $(P_\xi(h))_\xi$ is monotonic nonincreasing (by inclusion).*
- (iii) *The sequence $(\alpha_\xi(h))_\xi$ is monotonic nondecreasing.*

Before proving the theorem we define another property of plays, ξ -monotonicity, that will be used in the proof of Theorem 3.1.

DEFINITION 3.1. Let h be a finite history and p a play that extends h . The play p is called ξ -monotonic at h if the sequence $(P_\xi(h'))_{h \preceq h' < p}$ is nonincreasing:

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

PROOF [PROOF OF THEOREM 3.1]. The proof will follow once we prove by transfinite induction that the following properties hold for every ordinal ξ :

PROPERTY Q1. $P_\xi(h)$ is not empty for every history $h \in \mathcal{H}$. Moreover, if $\xi = 0$ or ξ is a limit ordinal, there is a play $p \in P_\xi(h)$ that is ξ -monotonic at h .

PROPERTY Q2. $P_\xi(h) \subseteq P_\lambda(h)$ for every ordinal $\lambda < \xi$ and for every history $h \in \mathcal{H}$.

PROPERTY Q3. $\alpha_\xi(h) \geq \alpha_\lambda(h)$ for every ordinal $\lambda < \xi$ and for every history $h \in \mathcal{H}$.

As we will see, most of the properties easily follow from the definitions. The proof of Q1 for a limit ordinal is the most challenging part of the whole proof.

For $\xi = 0$: Property Q1 follows from the definition (2). Properties Q2 and Q3 hold because there is no ordinal smaller than 0.

For a successor ordinal $\xi + 1$: Property Q1 follows from the definitions (4) and (5) and from the induction hypothesis (Property Q1).

The proof that Properties Q2 and Q3 hold is divided into two cases. Assume first that $\xi = \eta + 1$ is itself a successor ordinal. By the induction hypothesis Q2, $P_\xi(h, a) \subseteq P_\eta(h, a)$ for every history $h \in \mathcal{H}$ and every $a \in A$, so that by (4) it follows that Property Q3 holds, and therefore by (5) Property Q2 holds as well.

Assume now that either $\xi = 0$ or ξ is a limit ordinal. By the induction hypothesis (Property Q1) there is a play $\hat{p} \in P_\xi(h)$ that is ξ -monotonic at h . Let a_0 be the first action in $\hat{p}$ after h . Then $P_\xi(h) \supseteq P_\xi(h, a_0)$. Therefore, by (4) and (7),

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

and Q3 holds.

We now prove that Property Q2 holds. For $\xi = 0$ this follows trivially from (2). We therefore assume that ξ is a limit ordinal. Let $p \in P_{\xi+1}(h)$. By (5) we have $p \in P_\xi(h, a_0)$, where a_0 is the first action in p after h . By (6), $p \in P_\lambda(h, a_0)$ for every ordinal $\lambda < \xi$. Because we already proved that Property Q3 holds for $\xi + 1$,

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

so that by (5), $p \in P_{\lambda+1}(h)$ for every $\lambda < \xi$. It follows by (6) that $p \in P_\xi(h)$, as desired.

For a limit ordinal ξ : Property Q2 follows from definition (6), and hence Lemma 3.1(1) implies that Property Q3 holds as well. It remains to prove that Property Q1 holds for every limit ordinal ξ . The rest of this section is dedicated to this proof.

Fix a limit ordinal ξ . It will be sufficient to prove that if Properties Q1-Q3 hold for every ordinal λ smaller than ξ , and Properties Q2 and Q3 hold for ξ as well, then Property Q1 also holds for the ordinal ξ . The following lemma follows from Properties Q2 and Q3.

LEMMA 3.2. *For every finite history h , the sequence $(\alpha_\lambda(h))_{\lambda \leq \xi}$ is monotonic nondecreasing, and the sequence $(P_\lambda(h))_{\lambda \leq \xi}$ is monotonic nonincreasing (by inclusion).*

For every finite history h define

$$\tilde{\alpha}_\xi(h) := \max_{\lambda < \xi} \alpha_\lambda(h). \quad (8)$$

The maximum in (8) exists because the payoff functions have finite range. By Lemma 3.2, $\tilde{\alpha}_\xi(h) \leq \alpha_\xi(h)$.

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

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

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

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

(ii) Let a_1 be an action of player $i(h_1)$ that achieves the maximum in (4) for h_1 and ξ_1 , that is,

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

Set $h_2 = (h_1, a_1)$.

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

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

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

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

(v) Continue this way to create a sequence $(h_1, \xi_1, a_1, h_2, \xi_2, a_2, \dots)$. The iteration ends when either $\xi_m = 0$ or ξ_m is a limit ordinal; in this case, the output of the odd iteration is the finite history $h_m = (h_1, a_1, a_2, \dots, a_{m-1})$. If $\xi_m > 0$ is a successor ordinal for every $m \in \mathbb{N}$, the iteration never ends.

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

LEMMA 3.3. *There is $m \in \mathbb{N}$ such that either $\xi_m = 0$ or ξ_m is a limit ordinal.*

PROOF. Assume that the algorithm never terminates: $\xi_m > 0$ is a successor ordinal for every $m \in \mathbb{N}$, so that the algorithm generates an infinite sequence $(h_1, \xi_1, a_1, h_2, \xi_2, a_2, \dots)$.

We first argue that for every $m \in \mathbb{N}$ one has

$$P_{\xi_m}(h_m) \supseteq P_{\xi_m-1}(h_{m+1}) \supseteq P_{\xi_{m+1}}(h_{m+1}). \quad (9)$$

Indeed, the first inclusion holds by Lemma 3.1(ii), whereas the second inclusion holds by Lemma 3.2 and because $\xi_m - 1 \leq \xi_{m+1}$.

By (9), for every player j

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

Because the payoffs are discrete, the inequalities in (10) can be strict only finitely many times, for every player j . That is, there is $M \in \mathbb{N}$ sufficiently large such that for every player $j \in I$ and every $m \geq M$,

$$\min_{p \in P_{\xi_m}(h_m)} u^j(p) = \min_{p \in P_{\xi_m-1}(h_{m+1})} u^j(p) = \min_{p \in P_{\xi_{m+1}}(h_{m+1})} u^j(p). \quad (11)$$

Let m, m' be two integers satisfying (a) $M \leq m < m'$ and (b) $i(h_m) = i(h_{m'})$. By repeated use of Equation (11),

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

Hence by Lemma 3.1(i) and by $i(h_m) = i(h_{m'})$

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

thus it follows that $\xi_{m'} = \xi_{m'} - 1$. Because this equality holds for every m' sufficiently large, there is m such that either $\xi_m = 0$ or ξ_m is a limit ordinal, as desired. $\square$

Even iterations: Let h_1 be the finite history that is the output of the previous odd iteration and denote by λ the last ordinal ξ_m generated by the previous odd iteration. In particular, either $\lambda = 0$ or λ is a limit ordinal, and $\lambda < \xi$. Moreover, $\tilde{\alpha}_\xi(h_1) = \alpha_\lambda(h_1)$.

By the induction hypotheses of Property Q1 (for either $\lambda = 0$ or a limit ordinal $0 < \lambda < \xi$), there is a play $p \in P_\lambda(h_1)$ that is λ -monotonic at h_1 .

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

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

Let $(h_m)_{m \in \mathbb{N}}$ denote all finite prefixes of p_* that extend h , so that $h_1 = h$. We partition these prefixes into the sets $\mathcal{H}^{\text{odd}}$ and $\mathcal{H}^{\text{even}}$, depending on whether the action after the prefix is induced in an odd respectively even iteration. Denote by ξ_m the ordinal that is attached to h_m in the construction of p_* ; it is a successor ordinal if $h_m \in \mathcal{H}^{\text{odd}}$ and a limit ordinal or 0 if $h_m \in \mathcal{H}^{\text{even}}$.

Note that if $h_m \in \mathcal{H}^{\text{even}}$ and $h_{m+1} \in \mathcal{H}^{\text{odd}}$, i.e., when we switch from an even iteration to an odd iteration, we have

$$\alpha_{\xi_m}(h_{m+1}) < \tilde{\alpha}_\xi(h_{m+1}) = \alpha_{\xi_{m+1}}(h_{m+1}). \quad (12)$$

LEMMA 3.4. *For any $m \in \mathbb{N}$, we have*

$$P_{\xi_m}(h_m) \supseteq P_{\xi_{m+1}}(h_{m+1}). \quad (13)$$

PROOF. Let $m \in \mathbb{N}$. We distinguish three cases. Assume first that $h_m \in \mathcal{H}^{\text{odd}}$. Then ξ_m is a successor ordinal, and (13) follows from

$$P_{\xi_m}(h_m) \supseteq P_{\xi_m-1}(h_{m+1}) \supseteq P_{\xi_{m+1}}(h_{m+1}). \quad (14)$$

Indeed, the first inclusion holds by Lemma 3.1(ii), whereas the second inclusion holds by Lemma 3.2 and because $\xi_m - 1 \leq \xi_{m+1}$.

Assume now that $h_m, h_{m+1} \in \mathcal{H}^{\text{even}}$. Then (13) follows because $\xi_m = \xi_{m+1}$ (both are equal to the ordinal λ of this even iteration) and the part of the play added in this even iteration is ξ_m -monotonic.

Assume finally that $h_m \in \mathcal{H}^{\text{even}}$ and $h_{m+1} \in \mathcal{H}^{\text{odd}}$. Then by construction $\xi_m < \xi_{m+1}$. Hence, (13) follows from the ξ_m -monotonicity of the part of the play added in the even iteration and Lemma 3.2. $\square$

LEMMA 3.5. *The play p_* is ξ -monotonic at h .*

PROOF. Let $m \geq 1$, and let $p \in P_{\xi}(h_{m+1})$. We will prove that $p \in P_{\xi}(h_m)$. By (6) it follows that $p \in P_{\tau}(h_{m+1})$, for every $\tau < \xi$, and in particular $p \in P_{\xi_{m+1}}(h_{m+1})$. Then, by Lemma 3.4, we have $p \in P_{\xi_m}(h_m)$. Hence, we have

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

for every ordinal $\tau < \xi$. Because $p \in P_{\tau}(h_{m+1})$, definition (5) implies that $p \in P_{\tau+1}(h_m)$ for every $\tau < \xi$, so that by (6) we have $p \in P_{\xi}(h_m)$. $\square$

We are now ready to prove Property Q1 for a limit ordinal ξ .

LEMMA 3.6. $p_* \in P_{\xi}(h)$.

PROOF. Suppose first that the number of iterations is finite, so that the last even iteration is infinite. Denote by h_m the history at the beginning of the last even iteration, i.e., $h_{m-1} \in \mathcal{H}^{\text{odd}}$ and $h_{m'} \in \mathcal{H}^{\text{even}}$ for all $m' \geq m$. Then $\xi_m = \xi_{m+1} = \dots =: \lambda$, where λ is a limit ordinal. Moreover, $p_* \in P_{\lambda}(h_m)$ by the properties of an even iteration. We will show that $p_* \in P_{\xi}(h_m)$, so that by the ξ -monotonicity of p_* (Lemma 3.5) it will follow that $p_* \in P_{\xi}(h)$, as desired. Note that by the definition of even iterations, $\tilde{\alpha}_{\xi}(h_{m'}) = \alpha_{\lambda}(h_{m'})$ for every $m' \geq m$.

Assume to the contrary that $p_* \notin P_{\xi}(h_m)$. Let τ be the smallest ordinal such that $p_* \notin P_{\tau}(h_{m'})$ for some $m' \geq m$. Note that $\tau > \lambda$: because $p_* \in P_{\lambda}(h_m)$, by Lemma 3.1(iii) we have $p_* \in P_{\lambda}(h_{m'})$ for every $m' \geq m$. By definition (6), τ cannot be a limit ordinal, so that τ is a successor ordinal. It follows that $p_* \in P_{\tau-1}(h_{m'})$ for every $m' \geq m$. To derive a contradiction we argue that $p_* \in P_{\tau}(h_{m'})$ for every $m' \geq m$. Indeed, for every $m' \geq m$, because $p_* \in P_{\lambda}(h_{m'})$, $\tilde{\alpha}_{\xi}(h_{m'}) = \alpha_{\lambda}(h_{m'})$, and $\xi > \tau$, it follows that

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

so that by definition (5) we have $p_* \in P_{\tau}(h_{m'})$, as claimed.

We now show that the number of iterations cannot be infinite. From Lemma 3.4 it follows that for every $m \in \mathbb{N}$ and every player j ,

$$\min_{p \in P_{\xi_m}(h_m)} u^j(p) \leq \min_{p \in P_{\xi_{m+1}}(h_{m+1})} u^j(p). \quad (15)$$

Because the range of the payoffs is finite, the inequality (15) can be strict only finitely many times, for every player j .

Assume that $h_m \in \mathcal{H}^{\text{even}}$ and $h_{m+1} \in \mathcal{H}^{\text{odd}}$. Then

$$\min_{p \in P_{\xi_m}(h_m)} u^{i(h_{m+1})}(p) \leq \min_{p \in P_{\xi_m}(h_{m+1})} u^{i(h_{m+1})}(p) < \min_{p \in P_{\xi_{m+1}}(h_{m+1})} u^{i(h_{m+1})}(p), \quad (16)$$

where the weak inequality holds because the part of p_* generated by the even iteration is ξ_m -monotonic, and the strict inequality holds because of (12) and Lemma 3.1(i). In particular, by Equation (15), because the range of the payoffs is finite and because there are finitely many players, Equation (16) can hold only finitely many times, so there can be only finitely many even iterations. $\square$

The proof of Theorem 3.1 is now complete. $\square$

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

THEOREM 3.2. *There is an ordinal ξ_* such that for every finite history $h \in \mathcal{H}$ we have $\alpha_{\xi_*}(h) = \alpha_{\xi_*+1}(h)$, and $P_{\xi_*}(h) = P_{\xi_*+1}(h) \neq \emptyset$.*

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

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

By Lemma 3.1(i) it follows that for each such ordinal ξ , $\alpha_\xi(h) = \alpha_{\xi+1}(h)$ for every $h \in \mathcal{H}$, and the result follows. $\square$

3.3. Proof of Theorem 2.3. We now construct a strategy profile $\sigma_* = (\sigma_*^j)_{j \in I}$ and show that it is a subgame-perfect 0-equilibrium.

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

$$u^{i(h^-)}(p(h)) = \min_{p \in P_{\xi_*}(h)} u^{i(h^-)}(p). \quad (17)$$

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

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

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

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

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

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

We complete the proof by showing that

$$u^j(p) \leq u^j(p_*). \quad (18)$$

It is sufficient to show that

$$u^j(p_k) \leq u^j(p_*), \quad \forall k. \quad (19)$$

Indeed, if σ^j and σ_*^j differ only finitely many times along p , Equation (18) follows from Equation (19) applied to the last k ; if σ^j and σ_*^j differ infinitely many times along p , then the sequence $(p_k)_{k \in \mathbb{N}}$ converges to p , so that Equation (18) follows from Equation (19) and the lower-semicontinuity of u^j . This is the only place in the proof where the lower-semicontinuity of the payoff functions is used.

The proof of (19) is by induction on k . Because of the construction and Lemma 3.1(iii), we have $p_* \in P_{\xi_*}(h)$. Hence, Lemma 3.1(iii) and Equation (17) imply that $u^j(p_1) \leq u^j(p_*)$. For every $k \geq 1$, because $p_k \in P_{\xi_*}(h_{t_k})$, and by Lemma 3.1(iii) and Equation (17), $u^j(p_{k+1}) \leq u^j(p_k)$, which is at most $u^j(p_*)$ by the induction hypothesis. The proof is now complete.

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

THEOREM 3.3. *A play p is induced by some subgame-perfect 0-equilibrium if and only if $p \in P_{\xi_*}(\emptyset)$.*

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

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

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

Because every play that extends h is in $P_0(h)$, the claim holds for $\xi = 0$.

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

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

3.5. An example. In this section we illustrate the construction presented in the proof of Theorem 2.3. Consider the following two-player perfect-information game, where the set of actions is $A = \{c, s\}$ and the players play alternately. If both players always choose c , the payoff to each player is 0. Otherwise, let k be the first player who plays action s . If $k = 1$, the payoff vector is $(1, 2)$, while if $k = 2$, the payoff vector is $(2, 1)$. Note that the payoff functions are lower-semicontinuous.

We now construct the sequences $(\alpha_{\xi}(h))_{\xi}$ and $(P_{\xi}(h))_{\xi}$. Note that if $h \in \mathcal{H}$ contains the action s , then the payoff is independent of the continuation play, and therefore for every ξ , $\alpha_{\xi+1}(h) = \alpha_{\xi}(h) \neq 0$ and $P_{\xi+1}(h) = P_{\xi}(h)$. Moreover, each profile in the subgame that starts at h is a subgame perfect 0-equilibrium.

Assume next that $h = (c, c, \dots, c)$. For $\xi = 0$, $P_0(h)$ is the set of all plays that extends h , including the play in which both players always play c . It follows that $\alpha_0(h) = 0$.

We continue with the ordinal $\xi = 1$. If player $i(h)$ plays action s , then he receives payoff 1. On the other hand, if he plays action c , then play continues in $P_0(h, c)$, with 0 as the worst possible payoff. Therefore, $\alpha_1(h) = 1$, and $P_1(h)$ consists of all plays that extend h in which at least one of the players plays s .

We now turn to the ordinal $\xi = 2$. Again, action s yields 1 to player $i(h)$, whereas action c results in a continuation play in $P_1(h, c)$ with 1 as the worst possible payoff. Hence $\alpha_2(h) = \alpha_1(h) = 1$ and $P_2(h) = P_1(h)$.

Consequently, for every $\xi \geq 1$, $\alpha_{\xi+1}(h) = \alpha_{\xi}(h) = 1$ and $P_{\xi+1}(h) = P_{\xi}(h)$. It follows that $\xi_* = 2$, and for every finite history $h \in \mathcal{H}$ one has $\alpha_{\xi_*}(h) = 1$, and $P_{\xi_*}(h)$ consists of all plays that extend h where at least one of the players plays s (at least once). Thus, each play $p \in P_{\xi_*}(h)$ is induced by the following subgame perfect 0-equilibrium in pure strategies: follow p as long as no player deviates. If a player deviates to action s , then switch to an arbitrary continuation play. If, on the other hand, a player deviates to action c , then switch to a continuation play in which this very same player is the first one to play action s .

4. Concluding remarks

4.1. How large can ξ_* be? A natural question is whether the ordinal ξ_* of Theorem 3.2 is at most the first infinite ordinal ω , or whether it can be larger than ω . In this section we present an example that demonstrates that ξ_* can be any ordinal. Moreover, a variant of this game shows that ξ_* can be larger than ω , even when the number of actions is finite.

4.1.1. ξ_* can be any ordinal. Let τ be any ordinal. We will now construct a game G^{τ} for which $\xi_* = \tau$. For simplicity of exposition, in this game the set of actions is history dependent. Roughly speaking, G^{τ} is a two-player game in which players I and II choose a nonincreasing sequence $(a_t)_{t=0}^{\infty}$ of ordinals, with $a_0 = \tau$, according to the following rules. If the current ordinal a_{t-1} is a successor ordinal, then player I chooses a_t from $\{a_{t-1}, a_{t-1} - 1\}$. If the current ordinal a_{t-1} is a limit ordinal, then player II chooses any ordinal a_t smaller than a_{t-1} . Finally, if $a_{t-1} = 0$, then all further choices $(a_m)_{m \geq t}$ will be 0 as well. The payoff for player I equals 1 if the sequence $(a_t)_t$ eventually reaches 0 and equals 0 otherwise. The payoff for player II equals 0 for every play.

The idea of this game G^{τ} is that player I can make sure that the sequence $(a_t)_{t=0}^{\infty}$ eventually reaches 0 and thereby obtain the best payoff 1, but the number of stages needed to reach 0 depends on player II. More precisely, if player I, whenever he is the active player, always lowers the current ordinal by 1, then $a_t < a_{t-1}$ holds as long as $a_{t-1} > 0$. Because there is no infinite decreasing sequence of ordinals, this strategy of player I guarantees that $a_t = 0$ for some $t \in \mathbb{N}$, regardless the actions that player II chooses. Still, if λ_1 and λ_2 are two limit ordinals such that $\lambda_2 < \lambda_1 \leq \tau$, there is no bound on the number of stages needed to descend from λ_1 to λ_2 , as player II can choose the ordinal $\lambda_2 + k$ for any $k \in \mathbb{N}$ when the current ordinal is λ_1 . As we will show, one needs τ steps in our iterative method to realize that the sequence $(a_t)_t$ eventually reaches 0 whatever ordinals player II chooses.

Now we provide a formal definition of the two-player perfect-information game G^τ . For any history⁵ $h = (a_0, a_1, \dots, a_{t-1})$, the active player $i(h)$ and his action set $A(h)$ are defined as follows:

- If a_{t-1} is a successor ordinal, $i(h) = \text{I}$ and $A(h) = \{a_{t-1}, a_{t-1} - 1\}$.
- If a_{t-1} is a limit ordinal, $i(h) = \text{II}$ and $A(h) = \{\text{all ordinals smaller than } a_{t-1}\}$.
- If $a_{t-1} = 0$, $i(h) = \text{I}$ and $A(h) = \{0\}$.⁶

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

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

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

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

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

For $\xi = 0$, the claim is obvious.

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

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

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

Finally, let ξ be a limit ordinal, and assume that the claim holds for all ordinals $\lambda < \xi$. If either $\xi < a_{t-1}$ or $\xi > a_{t-1}$, then the respective parts of the claim for ξ follow by (6). Suppose then that $\xi = a_{t-1}$ and take any play $p \in P_\xi(h)$. We will show that $p \in W$. Let a_t denote the action in p right after h . Because $p \in P_\xi(h)$ and $a_t < \xi$, we have $p \in P_{a_t+1}(h)$, and hence $p \in P_{a_t}(h, a_t)$. By the induction hypothesis, $p \in W$ as desired.

4.1.2. ξ_* can be larger than ω even with finitely many actions. We will now show that ξ_* can be larger than ω , the first infinite ordinal, even when the number of actions is finite. We will do so by examining the following variant of the game G^τ for $\tau = \omega + 1$, which was described in the previous section.

The action set is $A = \{c, s\}$. Again, players I and II choose a nonincreasing sequence $(a_t)_{t=0}^\infty$ of ordinals, with $a_0 = \omega + 1$, according to the following rules. If the current ordinal a_{t-1} is a successor ordinal, i.e., $a_{t-1} = \omega + 1$ or $0 < a_{t-1} < \omega$, then player I is the active player, and he can set either $a_t = a_{t-1}$ by playing action c or $a_t = a_{t-1} - 1$ by playing action s . If the current ordinal a_{t-1} is ω , say for the k th time, then player II is the active player, and he can set either $a_t = \omega$ by playing action c or $a_t = k$ by playing action s . If the current ordinal is zero, then only zero can be announced at all further stages.

The payoff for player I equals 1 if there is a T such that $a_t = 0$ for all $t \geq T$ or $a_t = \omega$ for all $t \geq T$, and his payoff is zero otherwise. The payoff for player II equals 0 for every play.

This game differs from $G^{\omega+1}$ in the sense that, if the current ordinal becomes ω , then player II is no longer able to choose all finite ordinals immediately. In the new game, if player II wants to move to finite ordinal k , then he first has to play action c precisely $k - 1$ times and then to play action s . It can be verified very similarly to $G^{\omega+1}$ that we need $\omega + 1$ iterations to achieve a fixed point, so that $\xi_* = \omega + 1$ as desired.

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

This property does not hold, e.g., for mixed equilibria in sequential games with simultaneous moves, because in such games, if the continuation payoffs change, then the set of mixed actions that form a Nash equilibrium in

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

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

the one-shot game with these continuation payoffs may change as well, and a deviation from the original mixed equilibrium may not be detected.

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

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

As the following example shows, when there are infinitely many players, a subgame-perfect 0-equilibrium need not exist, even when the range of the payoff functions is finite. Suppose that the set of players is the set $\mathbb{N}$ of natural numbers, and the set of actions is $A = \{a, b\}$. Each player $t \in \mathbb{N}$ plays only once, at stage t . The payoff of player t is 1 if he played b , 2 if he played a and some player $j > t$ played b , and 0 if he played a and every player $j > t$ also played a .

This game has a 0-Nash equilibrium, where player 1 starts by playing b , and each other player $t > 1$ plays b only if every player $j < t$ also played b ; otherwise, player t plays a . On the other hand, there is no subgame-perfect 0-equilibrium in this game. Indeed, suppose to the contrary that σ is a subgame-perfect 0-equilibrium. Because every player can guarantee 1 by playing the action b , it cannot happen in any subgame that σ prescribes all players that have not played yet to play action a . This means in particular that, with respect to σ , infinitely many players play action b and receive 1. However, then each of those players is better off by deviating to a and receiving 2.

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

4.5. Positive recursive perfect-information games. Recursive perfect-information games are games in which some finite histories are terminating, in the sense that once they occur the payoff is determined (and the play that follows them does not affect the players' payoffs), and the payoff of every infinite (nonterminating) play is 0. Various positional games that are studied in the computer science literature have this form. The significance of this class of games to game theory was exhibited in the context of stochastic games by Vieille [23, 24], who used it as a step to proving the existence of an equilibrium payoff in every two-player stochastic game. A recursive perfect-information game is called positive if the terminal payoffs are positive for both players.

Flesch et al. [2] studied positive recursive perfect-information games with finitely many states; these are positional games that are played on a finite directed graph, where each vertex is controlled by some player, and when the game reaches some vertex, the controlling player can choose whether to terminate the game or whether to continue the game by choosing one of the edges that leaves the vertex. The terminal payoff, which is positive for all players, depends only on the vertex where termination occurred, not on the whole past play.

Flesch et al. [2] prove that every such game admits a subgame-perfect 0-equilibrium.⁷ In their proof, they define for every vertex s a sequence $(\alpha_k(s))_{k \in \mathbb{N}}$ that is similar to our sequence $(\alpha_\xi(h))_\xi$, they prove that this sequence is nondecreasing and, because there are finitely many vertices, they deduce that there is $k_* \in \mathbb{N}$ such that $\alpha_{k_*+1}(s) = \alpha_{k_*}(s)$ for every vertex s . They then use a similar construction of the subgame-perfect 0-equilibrium as the one that we used.

In perfect-information games every history is a different vertex. Therefore, one needs to use a much more delicate construction, which differs from the one in Flesch et al. [2] in two respects. First, when the number of vertices is infinite, there need not be $k_* \in \mathbb{N}$ such that $\alpha_{k_*+1}(s) = \alpha_{k_*}(s)$ for every vertex s , and therefore $(\alpha_\xi(h))_\xi$ should be defined for every ordinal. Second, because play never terminates, one has to deal with plays of infinite length and introduce the sets $(P_\xi(h))_\xi$.

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

⁷ When transitions are random, Flesch et al. [2] prove the existence of a subgame-perfect ε -equilibrium, for every $\varepsilon > 0$.

4.6. Perfect information games with general payoffs. Example 3 in Solan and Vieille [20, see also Footnote 1] shows that without the condition that payoffs are lower-semicontinuous, a subgame-perfect ε -equilibrium need not exist. However, Solan and Vieille [20] show that in their example a subgame-perfect ε -equilibrium does exist if one allows behavior strategies. The existence of a subgame-perfect ε -equilibrium in behavior strategies was proved in other setups where the payoff functions are not lower-semicontinuous, see, e.g., Mertens and Parthasarathy [13], Solan [19], Maitra and Sudderth [9], and Mashiah-Yaakovi [11].

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

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

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

Note that u^1 and u^2 are not lower-semicontinuous. Playing b finitely many times is a dominant strategy for both players, so that the unique subgame-perfect 0-equilibrium payoff is $(2, 2)$. However, one can verify that for every finite history h and every ordinal ξ , $P_\xi(h)$ contains all plays in which at least one player plays b finitely many times, so that the Folk Theorem, Theorem 3.3, does not hold, and our construction of the subgame-perfect 0-equilibrium in the proof of Theorem 2.3 is invalid.

Acknowledgments. The research of Mashiah-Yaakovi and Solan was supported by the Israel Science Foundation (Grant 212/09). The research of Mashiah-Yaakovi was partially supported by the Farajun Foundation Fellowship.

References

- [1] Emerson, E. A., C. S. Jutla, A. P. Sistla. 2001. On model checking for the μ -calculus and its fragments. *Theoret. Comput. Sci.* **258** 491–522.
- [2] Flesch, J., J. Kuipers, G. Schoenmakers, K. Vrieze. 2010. Subgame perfection in positive recursive games with perfect information. *Math. Oper. Res.* **35** 193–207.
- [3] Gale, D., F. M. Stewart. 1953. Infinite games with perfect information. H. W. Kuhn, A. W. Tucker, eds. *Contributions to the Theory of Games*, Volume 28. *Annals of Mathematics Studies*. Princeton University Press, Princeton, NJ, 245–266.
- [4] Grädel, E. 2004. Positional determinacy of infinite games. *STACS 2004, Lecture Notes in Computer Science* **2996**. Springer, Berlin, 4–18.
- [5] Grädel, E., M. Ummels. 2008. Solution concepts and algorithms for Infinite multiplayer games. K. Apt, R. van Rooij, eds. *New Perspectives on Games and Interaction*, Vol. 4. *Texts in Logic and Games*. Amsterdam University Press, Amsterdam, 151–178.
- [6] Kechris, A. S. 1995. Classical descriptive set theory. J. H. Ewing, F. W. Gehring, P. R. Halmos, eds. *Graduate Texts in Mathematics*, Vol. 156. Springer, Berlin.
- [7] Kuipers, J., J. Flesch, G. Schoenmakers, K. Vrieze. 2008. Pure subgame-perfect equilibria in free transition games. *Eur. J. Oper. Res.* **99** 442–447.
- [8] Maitra, A., W. Sudderth. 1993. Borel stochastic games with limsup payoffs. *Ann. Probab.* **20** 860–885.
- [9] Maitra, A. P., W. D. Sudderth. 2007. Subgame-perfect equilibria for stochastic games. *Math. Oper. Res.* **32** 711–722.
- [10] Martin, D. A. 1975. Borel determinacy. *Ann. Math.* **2** 363–370.
- [11] Mashiah-Yaakovi, A. 2009. Subgame-perfect equilibria in stopping games. Working paper, Tel Aviv University, Tel Aviv, Israel.
- [12] Mertens, J.-F. 1987. Repeated games. *Proc. Internat. Congress Mathematicians*. American Mathematical Society, Providence, RI, 1528–1577.
- [13] Mertens, J.-F., T. E. S. Parthasarathy. 2003. Equilibria for discounted stochastic games. A. Neyman, S. Sorin, eds. Chapter 10, *Stochastic Games and Applications*, Vol. 570, *NATO Science Series C, Mathematical and Physical Sciences*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 131–172.
- [14] Purves, R. A., W. D. Sudderth. 2010. Perfect information games with upper semicontinuous payoffs. Preprint.
- [15] Schilling, K., R. Vaught. 1983. Borel games and the Baire property. *Trans. Amer. Math. Soc.* **279** 400–428.
- [16] Selten, R. 1965. Spieltheoretische Behandlung eines Oligopolmodells mit Nachfrageträgheit. *Zeitschrift für die Gesamte Staatswissenschaft* **121** 301–324, 667–689.

- [17] Selten, R. 1973. A simple model of imperfect competition, where four are few and six are many. *Internat. J. Game Theory* **2** 141–201.
- [18] Shmaya, E. 2008. Many inspections are manipulable. *Theoret. Econom.* **3** 367–382.
- [19] Solan, E. 2005. Subgame-perfection in quitting games with perfect information. *Math. Oper. Res.* **30** 51–72.
- [20] Solan, E., N. Vieille. 2003. Deterministic multi-player Dynkin games. *J. Math. Econom.* **39** 911–929.
- [21] Thomas, W. 2002. Infinite games and verification. *Proc. Internat. Conf. Comput. Aided Verification, CAV'02*, Vol. 2404. Springer, Berlin, 58–64.
- [22] Ummels, M. 2005. Rational behaviour and strategy construction in infinite multiplayer games. Ph.D. thesis, RWTH Aachen University, Aachen, Germany.
- [23] Vieille, N. 2000. Equilibrium in 2-person stochastic games I: A reduction. *Israel J. Math.* **119** 55–91.
- [24] Vieille, N. 2000. Equilibrium in 2-person stochastic games II: The case of recursive games. *Israel J. Math.* **119** 93–126.