

Individual upper semicontinuity and subgame perfect ϵ -equilibria in games with almost perfect information

Citation for published version (APA):

Flesch, J., Herings, P. J.-J., Maes, J., & Predtetchinski, A. (2019). Individual upper semicontinuity and subgame perfect ϵ -equilibria in games with almost perfect information. Maastricht University, Graduate School of Business and Economics. GSBE Research Memoranda No. 002
<https://doi.org/10.26481/umagsb.2019002>

Document status and date:

Published: 14/01/2019

DOI:

[10.26481/umagsb.2019002](https://doi.org/10.26481/umagsb.2019002)

Document Version:

Publisher's PDF, also known as Version of record

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.

János Flesch,
P. Jean-Jacques Herings,
Jasmine Maes,
Arkadi Predtetchinski

**Individual upper
semicontinuity and subgame
perfect ϵ -equilibria in games
with almost perfect
information**

RM/19/002

GSBE

Maastricht University School of Business and Economics
Graduate School of Business and Economics

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

Individual upper semicontinuity and subgame perfect ϵ -equilibria in games with almost perfect information

János Flesch* P. Jean-Jacques Herings† Jasmine Maes ‡
Arkadi Predtetchinski§

January 2, 2019

Abstract

We study games with almost perfect information and an infinite time horizon. In such games, at each stage, the players simultaneously choose actions from finite action sets, knowing the actions chosen at all previous stages. The payoff of each player is a function of all actions chosen during the game. We define and examine the new condition of individual upper semicontinuity on the payoff functions, which is weaker than upper semicontinuity. We prove that a game with individual upper semicontinuous payoff functions admits a subgame perfect ϵ -equilibrium for every $\epsilon > 0$, in eventually pure strategy profiles.

Keywords: Almost perfect information, subgame perfect ϵ -equilibrium, individual upper semicontinuity.

JEL: C62, C65, C72, C73.

1 Introduction

Games with almost perfect information play a prominent role in the theory of dynamic games. In these games, at each stage, the players simultaneously choose actions, knowing the actions chosen at all previous stages. The payoff of each player is determined by the sequence of actions that have been chosen during the entire game. The notion of subgame perfect equilibrium and the weaker

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

†Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. E-mail: p.herings@maastrichtuniversity.nl

‡Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. E-mail: jasmine.maes@maastrichtuniversity.nl

§Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. E-mail: a.predtetchinski@maastrichtuniversity.nl

We thank Jérôme Renault, whose remarks inspired us to investigate this research question.

notion of subgame perfect ϵ -equilibrium, where $\epsilon > 0$ is an error term, are two main solution concepts in these games. We continue the discussion of these games under the assumptions that the number of players is finite and the payoffs are bounded.

A major body of the literature on games with almost perfect information examines the existence of subgame perfect equilibrium while assuming some kind of continuity of the payoff functions and compactness of the action spaces. Fudenberg and Levine (1983) showed that if the payoff functions are continuous and the action spaces are finite, then there exists a subgame perfect equilibrium. With infinite but compact action spaces the existence question becomes very subtle, see for instance Harris, Reny and Robson (1995), Mertens and Parthasarathy (2003), or Maitra and Sudderth (2007). For a detailed overview, we refer the reader to the recent survey by Jaśkiewicz and Nowak (2016).

In this paper we continue the investigation of the existence of subgame perfect ϵ -equilibria in games with almost perfect information. We assume that the set of available actions is finite. The sequence of actions chosen in the game is referred to as a play. The contribution of this paper is twofold. First, we introduce a new condition, called individual upper semicontinuity, on the payoff functions. Individual upper semicontinuity is reminiscent of upper semicontinuity, but it uses a new notion of convergence on the set of plays for each player separately.

Consider a player i . We say that a sequence of plays $(p_m)_{m \in \mathbb{N}}$ is i -convergent to a play p if (1) the length of the common history between the play p_m and p tends to infinity, and (2) for large $m \in \mathbb{N}$, either $p_m = p$ or the first deviation from p to p_m involves only player i . This notion of convergence of plays is stronger than the conventional notion, which would only require property (1). That is, each i -convergent sequence of plays is also convergent in the conventional sense, but not vice versa. Based on the notion of i -convergence, we say that the payoff function of player i is i -upper semicontinuous if for each sequence of plays $(p_m)_{m \in \mathbb{N}}$ that is i -convergent to a play p , player i weakly prefers the payoff along p compared to the payoff along p_m for large m . If player i 's payoff function is upper semicontinuous in the conventional sense, then it is also i -upper semicontinuous, but not vice versa. We analyze the topological properties of i -convergence and the properties of i -upper semicontinuous payoff functions in detail.

Secondly, we show that if each player i 's payoff function is i -upper semicontinuous, then the game admits a subgame perfect ϵ -equilibrium for every $\epsilon > 0$. Moreover, these strategy profiles are eventually pure, meaning that in each subgame randomization is only used at finitely many stages with probability 1. The proof of our existence result is constructive. The key idea is that if each player i 's payoff function is i -upper semicontinuous, then each play eventually reaches a history after which no player has an incentive to deviate from this play. This enables us to essentially cut the time horizon of the game at some finite stages and to apply backward induction on the earlier stages. This existence result generalizes the result of Purves and Sudderth (2011), who study perfect information games with upper semicontinuous payoff functions, in two directions. We reduce the topological restrictions on the payoff functions of the players and we allow the players

to move simultaneously.¹ In normal form games with totally ordered compact strategy spaces, Prokopovych and Yannelis (2017) show that it is possible to ensure the existence of a pure Nash equilibrium, while significantly relaxing the requirements relating to the upper semicontinuity and single crossing properties of the payoff functions.

Games with perfect information are an important special case of games with almost perfect information. Conditions related to semicontinuity of the payoff functions play an important role in several results for perfect information games. Not only do perfect information games where every player has an upper semicontinuous payoff function admit a subgame perfect ϵ -equilibrium for every $\epsilon > 0$ (Purves and Sudderth, 2011), so do perfect information games where every player has a lower semicontinuous payoff function as proven in Flesch et al. (2010). Moreover, these results were extended and unified in Flesch and Predtetchinski (2016). Questions related to semicontinuity are also studied in Le Roux and Pauly (2016) and Bruyère, Le Roux, Pauly and Raskin (2017). A counterexample by Flesch et al. (2014) shows that perfect information games with finite action sets and Borel measurable payoff functions do not always have subgame perfect ϵ -equilibria. This further illustrates the importance of topological properties of payoff functions in equilibrium analysis. For an overview on subgame perfect ϵ -equilibria in perfect information games, we refer to Jaśkiewicz and Nowak (2016) and Bruyère (2017). We remark that it follows from a result of Secchi and Sudderth (2001) that if the players have upper semicontinuous payoff functions in a stochastic game with a countable state space, then the game admits a Nash ϵ -equilibrium for every $\epsilon > 0$.

The remainder of this paper is structured as follows. In Section 2, we describe the model and introduce the notions of i -convergence for plays and i -upper semicontinuity for payoff functions of player i . In Section 3, we present four illustrating examples. In Section 4, we prove the above-mentioned existence result for subgame perfect ϵ -equilibria. In Section 5, we study the topology induced by i -convergence in detail. Finally, in Section 6, we provide some concluding remarks.

2 The model

In this section we describe the model and define the notion of individual upper semicontinuity of payoff functions. Let $\mathbb{N} = \{1, 2, \dots\}$ and $\mathbb{N}^* = \{0, 1, 2, \dots\}$.

The game: We consider games with almost perfect information and an infinite time horizon. We assume that the set of players and the sets of available actions are finite. Such a game is given by a tuple $G = \{I, (A_i)_{i \in I}, H, (u_i)_{i \in I}\}$, where

1. I is a non-empty and finite set of players.
2. For each player $i \in I$, A_i is a non-empty finite set of actions. Let $A = \times_{i \in I} A_i$. The set A corresponds to the set of stage game outcomes.

¹Purves and Sudderth (2011) allow arbitrary action sets in perfect information games. In Section 6 we argue that for perfect information games, we can dispense with the assumption that the action sets are finite.

3. $H \subseteq \cup_{t \in \mathbb{N}^*} A^t$ is a non-empty set of histories, where A^0 is the singleton $\{\emptyset\}$. At each history $h \in H$, let $A(h) = \{a \in A | ha \in H\}$ and for each player $i \in I$ let $A_i(h)$ denote the projection of $A(h)$ on A_i . The set $A(h)$ corresponds to the set of stage game outcomes at history h and $A_i(h)$ to the set of available actions for player i at history h . We assume that: (i) H is closed under initial segments: if for some $t \in \mathbb{N}^*$ and $ha \in A^{t+1}$ we have $ha \in H$, then we also have $h \in H$, (ii) for each history $h \in H$ and each player $i \in I$, the set $A_i(h)$ is non-empty, (iii) the set of stage game outcomes at history h is a product set: for each $h \in H$ it holds that $A(h) = \times_{i \in I} A_i(h)$.
4. Let P denote the set of all sequences $(a^1, a^2, \dots) \in A^{\mathbb{N}}$ such that for every $t \in \mathbb{N}$ we have $(a^1, \dots, a^t) \in H$. The set P is called the set of plays.² For each $i \in I$, $u_i : P \rightarrow \mathbb{R}$ is the payoff function of player i .

The game is played at stages in \mathbb{N} . At stage 1, each player $i \in I$ chooses an action a_i^1 from the set $A_i(\emptyset)$, independently of the other players. This yields a stage game outcome $a^1 = (a_i^1)_{i \in I}$. At a general stage $t \in \mathbb{N}$, knowing the previous stage game outcomes a^1, \dots, a^{t-1} , each player $i \in I$ chooses an action a_i^t from the set $A_i(a^1, \dots, a^{t-1})$, independently of the other players. This yields a stage game outcome $a^t = (a_i^t)_{i \in I}$. The payoff of each player $i \in I$ is given by $u_i(a^1, a^2, \dots)$.

Notations for histories and plays: For every $t \in \mathbb{N}^*$, for every history $h = (a^1, \dots, a^t) \in H$, let $\|h\| = t$ denote the number of stage games played during h . We refer to $\|h\|$ as the length of history h . For $k \in \mathbb{N}^*$ such that $k \leq t$, let $h_{|k} = (a^1, \dots, a^k)$ denote the truncated history available after stage game k . If $h, h' \in H$ are such that there exists $t \in \mathbb{N}^*$ for which $h'_{|t} = h$, then h is called an initial segment of h' , denoted by $h \preceq h'$. We write $h \prec h'$ if $h \preceq h'$ and $h \neq h'$. Furthermore if $h \preceq h'$ then we can define the maximum of these histories as $\max(h, h') = h'$ and the minimum of these histories as $\min(h, h') = h$.

For every $t \in \mathbb{N}^*$, let $p_{|t}$ denote the history that arises by restricting p to the first t stages. Let $p_{[t],i} = a_i^t$ denote the action that player i played at stage t along the play p . For a history $h \in H$ and a play $p \in P$ we write $h \prec p$ if p is an extension of h , i.e. for some $t \in \mathbb{N}^*$ we have $p_{|t} = h$. We call h a prefix of p . For any two distinct plays $p, q \in P$, let $\min(p, q)$ denote the longest common history shared by those plays and let $\ell(p, q)$ denote its length. Furthermore, define $\min(p, p) = p$ and $\ell(p, p) = \infty$. For any two distinct plays $p, q \in P$, let $I(p, q) \subseteq I$ denote the subset of players who deviated first from the play p to the play q , i.e. if $\ell(p, q) = t$ then $I(p, q) = \{i \in I | p_{[t+1],i} \neq q_{[t+1],i}\}$. We formally define $I(p, p) = \emptyset$. Note that I, ℓ , and \min are all symmetric in their two arguments.

Strategies: A mixed action for player $i \in I$ at history $h \in H$ is a probability measure on $A_i(h)$. A *strategy* for player $i \in I$ is a mapping σ_i that assigns to each history $h \in H$ a mixed action $\sigma_i(h)$ for player i at history h . Let \mathcal{S}_i denote the set of all strategies of player i . If for every history $h \in H$ the mixed action $\sigma_i(h)$ places probability 1 on one action, then σ_i is called a *pure strategy*.

²Properties (3.i) and (3.ii) mean that H is a pruned tree on A . Thus, P can be seen as the set of all infinite branches of H .

A tuple of strategies $\sigma = (\sigma_i)_{i \in I}$ is called a strategy profile. The set of strategy profiles is denoted by $\mathcal{S} = \times_{i \in I} \mathcal{S}_i$. For each player $i \in I$, let $\sigma_{-i} = (\sigma_j)_{j \in I \setminus \{i\}}$ denote the profile of strategies of player i 's opponents. The strategy profile σ is called pure if each player i 's strategy σ_i is pure.

A pure strategy profile $\sigma \in \mathcal{S}$ induces a unique play from every history h , which we denote by $\pi(\sigma; h)$. For a general strategy profile $\sigma \in \mathcal{S}$, let H_σ denote the set of histories h with the following property: there is a play $p = (h, a^{\ell(h)+1}, a^{\ell(h)+2}, \dots)$ such that for each stage $k \geq \ell(h)$ and for each player $i \in I$, the mixed action $\sigma_i(h, a^{\ell(h)+1}, \dots, a^k)$ places probability 1 on action a_i^{k+1} . Intuitively, this means that at history h , the strategy profile σ induces the play p with probability 1. The strategy profile σ is called *eventually pure in every subgame* if for each history $h \in H$ and each play $p \succ h$ there is a history h' such that $h \preceq h' \prec p$ and $h' \in H_\sigma$. This means that starting at any history, any continuation play reaches a history after which the strategy profile σ induces a unique play.

Topology on the set of plays: We now define the standard cylinder topology on the set of plays P .³ For each history $h \in H$, let $P(h)$ denote the set of all plays extending h , i.e. $P(h) = \{p \in P | p \succ h\}$. Let \mathcal{T} denote the topology on the set of plays P induced by the collection $\{P(h) | h \in H\}$. It is easy to see that the collection $\{P(h) | h \in H\}$ forms a basis of the topology \mathcal{T} . That is, for a set $Q \subseteq P$, we have $Q \in \mathcal{T}$ exactly when Q can be written as a union of sets belonging to $\{P(h) | h \in H\}$. The topological space (P, \mathcal{T}) is metrizable. For example, the metric $d : P \times P \rightarrow [0, \infty)$ defined by $d(p, q) = 2^{-\ell(p, q)}$ induces the topology \mathcal{T} . Moreover, \mathcal{T} coincides with the relative topology inherited from the product topology on $A^\mathbb{N}$. Since the latter space is compact and P is a closed subset of it, we conclude that (P, \mathcal{T}) is compact.

Let Σ denote the Borel sigma-algebra corresponding to the topology \mathcal{T} . Then, (P, Σ) is a measurable space. Similarly, for every history $h \in H$, we can construct the measurable space $(P(h), \Sigma_h)$, where Σ_h is the sigma-algebra generated by all sets $P(h')$ with $h' \succeq h$. We have $\Sigma_h = \{Q \cap P(h) | Q \in \Sigma\}$.

As a consequence of the Ionescu Tulcea extension theorem, every strategy profile $\sigma \in \mathcal{S}$ induces for every history $h \in H$ a probability measure $\mathbb{P}_{h, \sigma}$ on the measurable space $(P(h), \Sigma_h)$. Note that σ is pure exactly when for each history $h \in H$, $\mathbb{P}_{h, \sigma}$ is a Dirac measure. Also, σ is eventually pure in every subgame exactly when for each history $h \in H$ and each play $p \succ h$ there exists a history h' such that $h \preceq h' \prec p$ and $\mathbb{P}_{h', \sigma}$ is a Dirac measure.

Convergence and upper semicontinuity: In the topology \mathcal{T} , a sequence of plays $(p_m)_{m \in \mathbb{N}}$ is convergent to the play p if $\lim_{m \rightarrow \infty} \ell(p_m, p) = \infty$. The payoff function u_i of player i is called upper semicontinuous if for each play $p \in P$ and each sequence $(p_m)_{m \in \mathbb{N}}$ that is convergent to p , we have $\limsup_{m \rightarrow \infty} u_i(p_m) \leq u_i(p)$. If u_i is upper semicontinuous, then it is also Borel measurable.

We now strengthen the notion of convergence. Consider a player $i \in I$. We say that a sequence of plays $(p_m)_{m \in \mathbb{N}}$ is *i -convergent* to the play p and write $\lim_{m \rightarrow \infty}^{(i)} p_m = p$ if $\lim_{m \rightarrow \infty} \ell(p_m, p) = \infty$ and

³For further reading, we refer to Kechris (1995).

there exists $M \in \mathbb{N}$ such that for every $m \geq M$, $I(p_m, p) \subseteq \{i\}$. The notion of i -convergence thus strengthens the notion of convergence by additionally imposing that eventually, if the plays p_m and p differ, player i is the only player who causes the first difference between p_m and p . Clearly, if a sequence of plays is i -convergent then it is also convergent.

In turn, this leads to a weakening of upper semicontinuity that we refer to as individual upper semicontinuity. The payoff function u_i of player i is called *i -upper semicontinuous* if for each play $p \in P$ and each sequence $(p_m)_{m \in \mathbb{N}}$ that is i -convergent to p , i.e. $\lim_{m \rightarrow \infty}^{(i)} p_m = p$, we have $\limsup_{m \rightarrow \infty} u_i(p_m) \leq u_i(p)$. As each i -convergent sequence of plays is also convergent, it holds that each upper semicontinuous payoff function is also i -upper semicontinuous. The converse however is not true and there are even i -upper semicontinuous payoff functions that are not Borel measurable. These issues are illustrated by Example 1 in Section 3.

Subgame perfect ϵ -equilibrium: Assume that for every player $i \in I$ the payoff function u_i is bounded and Borel measurable. Then, for $\epsilon \geq 0$, a strategy profile $\sigma \in \mathcal{S}$ is called a *subgame perfect ϵ -equilibrium* if for each history $h \in H$, each player $i \in I$, and each strategy $\sigma'_i \in \mathcal{S}_i$, we have

$$\mathbb{E}_{h, \sigma} [u_i] \geq \mathbb{E}_{h, \sigma'_i, \sigma_{-i}} [u_i] - \epsilon.$$

In other words, σ is a subgame perfect ϵ -equilibrium if at each history $h \in H$ it induces a Nash ϵ -equilibrium. When $\epsilon = 0$, the concept of subgame perfect 0-equilibrium coincides with the usual concept of subgame perfect equilibrium.

3 Examples

In this section we discuss a few illustrative examples.

The first example demonstrates that i -upper semicontinuity of u_i does not imply that u_i is upper semicontinuous. In fact, it does not even imply that u_i is Borel measurable.

Example 1. Consider the following game with two players. At each stage, player 1 chooses an action from the set $\{1, 2\}$. Player 2 is a dummy player, who can only choose action 0 at each stage. The set of plays is thus $P = (\{1, 2\} \times \{0\})^{\mathbb{N}}$. Let Q be a non-Borel set of P .⁴ Let the payoff function u_2 of player 2 be defined as $u_2(p) = 1$ if $p \in Q$ and $u_2(p) = 0$ if $p \in P \setminus Q$.

The payoff function u_2 is clearly not Borel measurable and hence not upper semicontinuous either. However, u_2 is 2-upper semicontinuous. Indeed, take any sequence of plays $(p_m)_{m \in \mathbb{N}}$ that is 2-converging to a play p . Since player 2 is a dummy player, we have $p_m = p$ for large m . Hence, $u_2(p_m) = u_2(p)$ for large m , proving that u_2 is 2-upper semicontinuous.

Of course, as long as the payoff function u_1 is bounded and Borel measurable, the game admits a subgame perfect ϵ -equilibrium for each $\epsilon > 0$. ◇

⁴The set P is an uncountable Polish space (essentially the Cantor space $\{1, 2\}^{\mathbb{N}}$), so P contains a non-Borel set, see Corollary 6.7.11 in Bogachev (2007).

The next example belongs to the class of so-called quitting games, see, for instance, Flesch et al. (1997) and Solan and Vieille (2001).

Example 2. Consider the following game with two players. The set of actions is $\{c_1, q_1\}$ for player 1 and $\{c_2, q_2\}$ for player 2, with c_1 and c_2 standing for “continue” and q_1 and q_2 for “quit”. The players choose actions simultaneously. For a play $p \in P$, let $K_1(p)$ denote the first stage at which player 1 chooses action q_1 . If player 1 always chooses action c_1 , then $K_1(p) = \infty$. We define $K_2(p)$ similarly for player 2. The payoffs $(u_1(p), u_2(p))$ for the players are defined as follows. They are equal to $(0, 0)$ if $K_1(p) = K_2(p) = \infty$, i.e. if the players always choose “continue”, $(-1, 1)$ if $K_1(p) < K_2(p)$, $(1, -1)$ if $K_1(p) > K_2(p)$, and $(2, 2)$ if $K_1(p) = K_2(p) < \infty$. Note that u_1 and u_2 are symmetric and continuous everywhere, except at the play $p^c = ((c_1, c_2), (c_1, c_2), \dots)$.

We focus on player 1. The payoff function u_1 is not upper semicontinuous. Indeed, for each $m \in \mathbb{N}$, consider the play p_m in which (c_1, c_2) is chosen at the first m stages and (q_1, q_2) at stage $m + 1$. The payoff for player 1 is 2 at each p_m . However, the sequence $(p_m)_{m \in \mathbb{N}}$ converges to the play p^c , which only gives payoff 0 to player 1.

Nevertheless, u_1 is 1-upper semicontinuous. Indeed, consider a sequence of plays $(p_m)_{m \in \mathbb{N}}$ that is 1-convergent to the play p^c . Then for large m we have either $p_m = p^c$ or $K_1(p_m) < K_2(p_m)$. Hence, for large m we obtain $u_1(p_m) \leq 0 = u_1(p^c)$. So, u_1 is 1-upper semicontinuous.

A subgame perfect equilibrium in the game is obtained if player 1 always chooses action q_1 and player 2 always chooses action q_2 . Another subgame perfect equilibrium is for example to always choose actions c_1 and c_2 . \diamond

We now illustrate the usefulness of i -upper semicontinuity by examining a class of games that includes the game in Example 2 as a special case.

Example 3. Consider the two-player game of Example 2, where the action sets for player 1 and 2 is given by $\{c_1, q_1\}$ and $\{c_2, q_2\}$ respectively. At every stage both players pick an action from their action space. The game ends when at least one player picks quit. Let $u_{i,\{j\}}^t$ denote the payoff of player i when player j unilaterally decides to quit at stage $t \in \mathbb{N}$ and let $u_{i,\emptyset}$ denote the payoff player i receives when quitting never takes place. For the sake of simplicity, assume that all payoffs are integers. Note that a payoff function of this game is continuous everywhere, except possibly at the play $p^c = ((c_1, c_2), (c_1, c_2), \dots)$.

The payoff function u_i of player i is i -upper semicontinuous in this game if the payoff player i receives when being the only one who quits at a late stage is at most the payoff player i would receive if the game goes on indefinitely, i.e. if there exists a stage $K \in \mathbb{N}$ such that for each $k \geq K$ we have $u_{i,\{i\}}^k \leq u_{i,\emptyset}$. This follows from the fact that the play p^c is the only possible point of discontinuity of u_i and the assumption that all payoffs are integers.

We now explain why the game admits a subgame perfect equilibrium provided that for each player $i \in I$ the payoff function u_i is i -upper semicontinuous. A generalization of the main idea behind the construction constitutes the basis for the proof of Theorem 5. Assume that for each player $i \in I$ the payoff function u_i is i -upper semicontinuous, and define a strategy profile $\sigma \in \mathcal{S}$ as

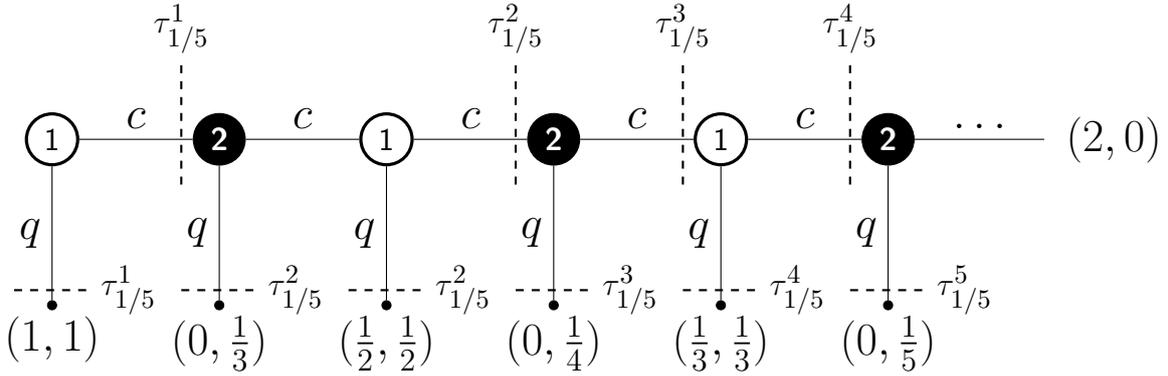


Figure 1: The non-existence of a subgame perfect equilibrium and the stopping times $\tau_{1/5}^m$.

follows. Since for each player $i \in I$ the payoff function u_i is i -upper semicontinuous, there exists a stage K such that, for each stage $k \geq K$, for each player $i \in I$, we have $u_{i,\{i\}}^k \leq u_{i,\emptyset}$. Consequently, after stage K , there is a play that is most preferred by both players: the play p^c . Thus, after stage K , we define σ_1 and σ_2 to continue indefinitely with probability 1. This implies that if stage no one quits before stage K , σ gives player i a payoff of $u_{i,\emptyset}$. We complete the strategy profile σ by a backward induction argument on the stages $K - 1, K - 2, \dots, 1$. By construction, σ constitutes a subgame perfect equilibrium of the game. \diamond

The following example shows that games with bounded and upper semicontinuous payoff functions may have no subgame perfect equilibrium. This motivates the use of the concept of subgame perfect ϵ -equilibrium.

Example 4. Consider the two-player centipede game depicted in Figure 1. Before we proceed, we spell out some of our notational conventions.

First, whenever an action is taken from a singleton action set, it will be omitted from our notation. Notice that, the game in Figure 1 having perfect information, at any given history only one of the players has a non-singleton action set. In accordance with our convention, we will only be recording the action taken by the active player. Thus for each $t \in \mathbb{N}^*$ we write c^t to denote the history of length t obtained if the players, whenever it is their turn to move, take the action c . In particular c^0 is the empty history. The symbol $c^t q$ denotes the history of length $t + 1$ obtained if the action q is chosen for the first time in period $t + 1$.

Second, we assume that, once either player takes the action q , both players' action sets become singletons for the rest of the game. This implies that the game only has countably many plays, which we denote by $c^0 q, c^1 q, c^2 q, \dots$ and c^∞ . Here $c^t q$ denotes the play where action q is chosen for the first time in period $t + 1$. The symbol c^∞ is the play such that no one ever plays q . With these conventions in place, we turn to the analysis of the game.

The payoff function of player 1 is upper semicontinuous and the payoff function of player 2 is continuous. We show that the game does not have a subgame perfect equilibrium.

Suppose that $\sigma = (\sigma_1, \sigma_2)$ is a subgame perfect equilibrium. For each $t \in \mathbb{N}^*$, let $\mathbb{P}_{c^t, \sigma}(c^\infty)$ denote the probability that, in the subgame at history c^t , under σ the play c^∞ is realized. We

distinguish two cases and derive a contradiction in each of them.

Assume first that for each $t \in \mathbb{N}^*$ we have $\mathbb{P}_{c^t, \sigma}(c^\infty) = 0$. As σ is a subgame perfect equilibrium, it follows that at each history where player 1 is active, σ_1 must place probability 1 on action q . This implies that, at each history where player 2 is active, σ_2 places probability 1 on action c . But then player 1 would have a profitable deviation from σ_1 by always choosing the action c , which is a contradiction.

Assume now that for some $t \in \mathbb{N}^*$ we have $\mathbb{P}_{c^t, \sigma}(c^\infty) > 0$. Then, $\mathbb{P}_{c^k, \sigma}(c^\infty)$ converges to 1 as $k \rightarrow \infty$. So there exists $K \in \mathbb{N}^*$ such that, for each $k \geq K$, player 1's expected payoff at history c^k is strictly more than 1, $\mathbb{E}_{c^k, \sigma}[u_1] > 1$. As σ is a subgame perfect equilibrium, σ_1 places probability 1 on action c at each history where player 1 is active beyond stage K . This implies that σ_2 places probability 1 on action q at each history where player 2 is active beyond stage K . But this contradicts $\mathbb{P}_{c^t, \sigma}(c^\infty) > 0$. Consequently, the game does not have a subgame perfect equilibrium.

Note however that for each $\epsilon > 0$ there exists a pure subgame perfect ϵ -equilibrium. Indeed, fix $\epsilon > 0$. Then there is $K \in \mathbb{N}$ such that if neither player quits before stage K then all feasible payoffs for player 2 are at most ϵ . Now the following pure strategy profile constitutes a subgame perfect ϵ -equilibrium: player 1 quits at each history, whereas player 2 continues at each history before stage K and quits at each history after stage K . Note that this strategy profile can be found as follows: let player 2 quit at each history after stage K and, given this, use backward induction at the various parts of the game tree. This is in line with the construction that we provide in the proof of Theorem 5. \diamond

4 Existence of a subgame perfect ϵ -equilibrium

The goal of this section is to present and prove the existence result, Theorem 5, for subgame perfect ϵ -equilibria in games with almost perfect information as defined in Section 2.

Section 4 is structured as follows. In the first subsection, we present and discuss this existence result. In the second subsection, we examine some implications of individual upper semicontinuity in more detail. In the third subsection, we provide the proof.

4.1 The existence result

The main result of this section is the following theorem.

Theorem 5. *Consider a game G with almost perfect information as defined in Section 2. If for every player $i \in I$ the payoff function u_i is bounded, Borel measurable, and i -upper semicontinuous, then, for each $\epsilon > 0$, the game admits a subgame perfect ϵ -equilibrium, which is eventually pure.*

Theorem 5 generalizes the result of Purves and Sudderth (2011) on the existence of subgame perfect ϵ -equilibria in perfect information games with upper semicontinuous payoff functions in two directions. First, as each upper semicontinuous function u_i is also Borel measurable and i -upper semicontinuous, we reduce the topological restrictions on the payoff functions of the players. Second, we allow the players to move simultaneously.

We provide a constructive proof for Theorem 5 in the next subsections. We remark that our proof is different from the one in Purves and Sudderth (2011), who after discretization of the payoffs, make use of an induction argument on the cardinality of the available payoffs in the subgames.

We discuss extensions of Theorem 5 and the necessity of the conditions in Section 6. Note Example 4 demonstrates that the conditions in Theorem 5 do not guarantee the existence of a subgame perfect equilibrium.

4.2 Optimal plays

Consider a player $i \in I$. For every play $p \in P$ and stage $t \in \mathbb{N}^*$, we define the set $O_i(p, t)$ as the set consisting of the play p and all other plays $q \neq p$ such that q coincides with p until at least stage t and the first deviation from p to q is caused by player i alone. That is,

$$O_i(p, t) = \{p\} \cup \{q \in P \mid \ell(p, q) \geq t \text{ and } I(p, q) = \{i\}\}. \quad (1)$$

Note that $O_i(p, t)$ is non-empty as it contains the play p .

Now we formulate a condition such that a player does not wish to deviate from a play after a certain stage.

Definition 6. Consider a player $i \in I$ and some $\epsilon > 0$. A play $p \in P$ is (ϵ, i) -optimal after stage t if for each play $q \in O_i(p, t)$

$$u_i(p) \geq u_i(q) - \epsilon.$$

A play $p \in P$ is ϵ -optimal after stage t if it is (ϵ, i) -optimal after stage t for each player $i \in I$. A play $p \in P$ is ϵ -optimal after history h if $h \prec p$ and p is ϵ -optimal after stage $\|h\|$.

The concept of ϵ -optimality is defined as a property of a play and not as a property of a strategy profile. Further, if the play p is ϵ -optimal after stage t then, for each $k \geq t$, the play p is also ϵ -optimal after stage k .

Lemma 7 and Corollary 8 relate individual upper semicontinuity and ϵ -optimality.

Lemma 7. Consider a player $i \in I$ and assume that player i 's payoff function u_i is bounded. Then u_i is i -upper semicontinuous if and only if for each play $p \in P$ and for each $\epsilon > 0$, there exists a stage $t \in \mathbb{N}^*$ such that the play p is (ϵ, i) -optimal after stage t .

Proof. Assume that the payoff function u_i is i -upper semicontinuous. Take a play $p \in P$ and some $\epsilon > 0$.

Suppose that for each $t \in \mathbb{N}^*$ there exists a play $q_t \in O_i(p, t)$ such that

$$u_i(p) < u_i(q_t) - \epsilon.$$

By construction, the sequence of plays $(q_t)_{t \in \mathbb{N}^*}$ is i -convergent to p and $u_i(p) \leq \limsup_{t \rightarrow \infty} u_i(q_t) - \epsilon$. This contradicts the assumption that u_i is i -upper semicontinuous.

For the other direction, fix a play $p \in P$ and some $\epsilon > 0$. Then, by assumption there exists a stage $t \in \mathbb{N}^*$ such that for every play $q_t \in O_i(p, t)$, $u_i(p) \geq u_i(q_t) - \epsilon$. Take any sequence of plays $(p_m)_{m \in \mathbb{N}}$ that i -converges to the play p . Then there exists $M \in \mathbb{N}$ such that for every $m \geq M$, $p_m \in O_i(p, t)$. Consequently, for every $m \geq M$, $u_i(p) \geq u_i(p_m) - \epsilon$. We conclude that $u_i(p) \geq \limsup_{m \rightarrow \infty} u_i(p_m) - \epsilon$. Since this holds for any play $p \in P$ and any $\epsilon > 0$, it follows that u_i is i -upper semicontinuous. \square

Since the set of players is finite, we have the following corollary.

Corollary 8. *Assume that, for each player $i \in I$, the payoff function u_i is bounded and i -upper semicontinuous. Then, for each play $p \in P$ and for each $\epsilon > 0$, there exists a stage $t \in \mathbb{N}^*$ such that the play p is ϵ -optimal after stage t .*

The next definition presents the notion of an ϵ -optimal history.

Definition 9. Let $\epsilon > 0$. A history $h \in H$ is ϵ -optimal if there exists a play $p \in P$ which is ϵ -optimal after history h .

In our construction of a subgame perfect ϵ -equilibrium, strategy profiles will be such that in a subgame corresponding to an ϵ -optimal history h , an ϵ -optimal play after history h is followed.

For $\epsilon > 0$, the number of times a history $h \in H$ has an initial segment which is ϵ -optimal is denoted by $n_\epsilon(h)$, so

$$n_\epsilon(h) = |\{t \in \mathbb{N}^* \mid t \leq \ell(h) \text{ and } h|_t \text{ is } \epsilon\text{-optimal}\}|.$$

Next, for $m \in \mathbb{N}$, the stopping time τ_ϵ^m is defined by

$$\tau_\epsilon^m(p) = \min\{t \in \mathbb{N}^* \mid n_\epsilon(p|_t) = m\}, \quad p \in P.$$

In words, $\tau_\epsilon^m(p)$ is the stage at which it occurs for the m -th time that a prefix of p is ϵ -optimal. Note that this does not mean that the play p needs to be ϵ -optimal from this prefix.

The following lemma claims that all these stopping times are uniformly bounded.

Lemma 10. *For every $\epsilon > 0$, for every $m \in \mathbb{N}$, there exists $K_\epsilon^m \in \mathbb{N}^*$ such that for each play $p \in P$ we have $\tau_\epsilon^m(p) \leq K_\epsilon^m$.*

Proof. Let some $\epsilon > 0$ and some $m \in \mathbb{N}$ be given.

Suppose that $\sup_{p \in P} \tau_\epsilon^m(p) = \infty$. Then, by finiteness of the set $A(\emptyset)$ of possible stage game outcomes at stage 1, there is $a^1 \in A(\emptyset)$ such that $\sup_{p \succ (a^1)} \tau_\epsilon^m(p) = \infty$. Then, for the same reason, there is $a^2 \in A(a^1)$ such that $\sup_{p \succ (a^1, a^2)} \tau_\epsilon^m(p) = \infty$. Continuing this way, we obtain a play $p = (a^1, a^2, \dots)$ for which $\tau_\epsilon^m(p) = \infty$. This is however in contradiction with Corollary 8. \square

Example 11. Consider again the two-player centipede game depicted in Figure 1 and discussed in Example 4. For the remainder of the example take $\epsilon = 1/5$. We will now compute the stopping times $\tau_{1/5}^m(p)$ for all $m \in \mathbb{N}$ and all plays $p \in P$. To do this we start by finding all $1/5$ -optimal

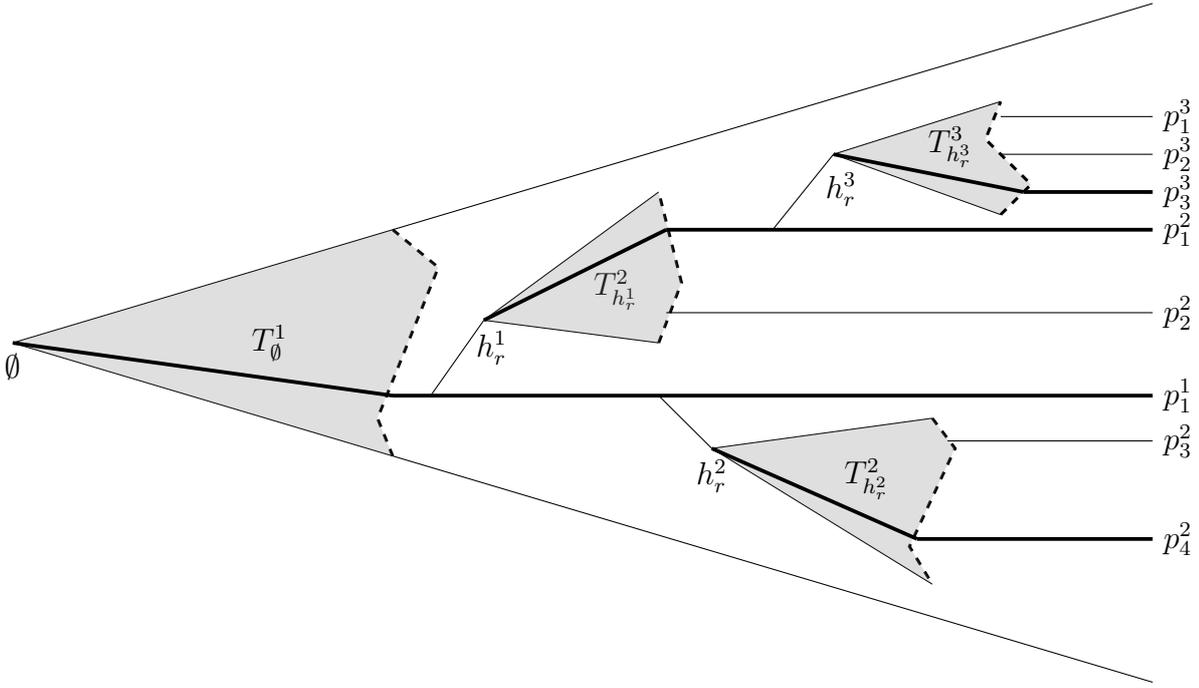


Figure 2: Construction of a subgame perfect ϵ -equilibrium.

histories.

Part 1: All histories $h \notin \{\emptyset, c^2\}$ are $1/5$ -optimal.

First note that all histories which contain a quitting action are clearly $1/5$ -optimal. All other histories have the form c^t for some $t \in \mathbb{N}^*$.

The history $c^0 = \emptyset$ is not $1/5$ -optimal. Indeed, the only $(1/5, 1)$ -optimal play after \emptyset is c^∞ . However, c^∞ is not an $(1/5, 2)$ -optimal play after \emptyset because the play cq gives player 2 a payoff strictly higher than $1/5$. Likewise the history c^2 is not $1/5$ -optimal: c^∞ is the only $(1/5, 1)$ -optimal play after history c^2 , but it is not $(1/5, 2)$ -optimal after c^2 , because the play c^3q gives player 2 a payoff strictly higher than $1/5$.

The history c is $1/5$ -optimal because the play cq is $1/5$ -optimal after history c . To see this, notice that player 1 cannot profitably deviate from cq after history c since history c is controlled by player 2. And any deviation from the play cq by Player 2 after history c can increase player 2's payoff by no more than $1/2 - 1/3 < 1/5$. Using an analogous argument one can easily show that history c^3 is $1/5$ -optimal as the play c^3q is $1/5$ -optimal after this history.

All histories c^t with $t \geq 4$ are $1/5$ -optimal because the play c^∞ becomes $1/5$ -optimal after history c^4 . Indeed, once history c^4 has occurred neither player can gain an additional payoff of more than $1/5$ by quitting.

Part 2: The stopping times $\tau^m(p)$.

From part 1 it now easily follows that $\tau_{1/5}^m(p) = m$ if $m = 1$ or if $p \in \{q, cq\}$. While $\tau_{1/5}^m(p) = m + 1$ if $m \geq 2$ and $p \notin \{q, cq\}$. \diamond

4.3 The proof of Theorem 5

In this subsection, we prove Theorem 5. Assume that for every player $i \in I$ the payoff function u_i is bounded, Borel measurable, and i -upper semicontinuous. Fix $\epsilon > 0$. Our goal is to construct a subgame perfect ϵ -equilibrium σ . The construction is illustrated graphically in Figure 2. The trees T_h^k are defined in the proof.

STEP 1: We start from the root of the game, at history \emptyset . Along each play $p \in P$, the first ϵ -optimal history is given by $p|_{\tau_\epsilon^1(p)}$. Let H^1 denote the set of all these ϵ -optimal histories:

$$H^1 = \{p|_{\tau_\epsilon^1(p)} \in H \mid p \in P\}.$$

For each history $h \in H^1$, choose a play p^h that is ϵ -optimal after h .

Let T_\emptyset^1 denote the subtree that consists of all the histories that can occur at or before this stopping time τ_ϵ^1 , so

$$T_\emptyset^1 = \{h \in H \mid \exists h^1 \in H^1 : h \preceq h^1\}.$$

In view of Lemma 10, the tree T_\emptyset^1 is finite. That is, the set of non-terminal nodes of T_\emptyset^1 is the finite set

$$Z^1 = \{h \in H \mid \exists h^1 \in H^1 : h \prec h^1\}$$

and the terminal histories of T_\emptyset^1 are exactly the histories in H^1 . At each terminal history $h^1 \in H^1$, for each player $i \in I$, we define the terminal payoff of T_\emptyset^1 to be $u_i(p^{h^1})$. Intuitively, if h^1 is reached, then player i receives the payoff corresponding to the ϵ -optimal play p^{h^1} extending h^1 . Given these terminal payoffs, we can find a subgame perfect equilibrium σ_\emptyset^1 for T_\emptyset^1 by backward induction. Note that σ_\emptyset^1 is not necessarily pure, due to the potential presence of simultaneous moves.

For each $h^1 \in H^1$, let $W^1(h^1)$ denote the prefixes of the play p^{h^1} after the subtree T_\emptyset^1 :

$$W^1(h^1) = \{h \in H \mid h^1 \preceq h \prec p^{h^1}\}.$$

Define

$$W^1 = \cup_{h^1 \in H^1} W^1(h^1).$$

We now define the strategy profile σ for histories belonging to Z^1 and W^1 . The strategy profile σ equals σ_\emptyset^1 at histories in Z^1 and follows the ϵ -optimal plays at histories in W^1 . Thus, for every $h \in Z^1$ we define $\sigma(h) = \sigma_\emptyset^1(h)$ and, for every $h^1 \in H^1$, for every $h \in W^1(h^1)$, $\sigma(h)$ puts probability 1 on $p_{[\|h\|+1]}^{h^1}$.

STEP 2: Now we proceed by considering the minimal histories outside $Z^1 \cup W^1$. These are the histories that arise outside T_\emptyset^1 when along a play p^{h^1} with $h^1 \in H^1$ a deviation occurs from p^{h^1} . That is, we are considering the set of histories

$$R^2 = \bigcup_{h^1 \in H^1} \{ha \in H \mid h \in W^1(h^1), a \in A(h)\} \setminus W^1(h^1).$$

For each history $h^\circ = ha \in R^2$, we execute the following. Similarly to step 1, by using the boundedness of stopping times, we let $H^2(h^\circ)$ be the set of minimal ϵ -optimal histories h^2 with initial segment h° such that $n_\epsilon(h^2) = n_\epsilon(h) + 1$, so

$$H^2(h^\circ) = \{p|_{\tau_\epsilon^{n_\epsilon(h)+1}(p)} \in H \mid p \in P(h^\circ)\}.$$

Let $T_{h^\circ}^2$ denote the finite subtree with root h° that consists of all histories that can occur at or before the stopping time $\tau_\epsilon^{n_\epsilon(h)+1}$, so

$$T_{h^\circ}^2 = \{h' \in H \mid \exists h^2 \in H^2(h^\circ) : h^\circ \preceq h' \preceq h^2\}.$$

Let $Z^2(h^\circ)$ denote the set of non-terminal histories belonging to $T_{h^\circ}^2$. The terminal histories of $T_{h^\circ}^2$ are precisely the histories in $H^2(h^\circ)$. For each history $h^2 \in H^2(h^\circ)$, choose a play p^{h^2} that is ϵ -optimal after h^2 . At each terminal history $h^2 \in H^2(h^\circ)$, for each player $i \in I$, we define the terminal payoff of $T_{h^\circ}^2$ to be $u_i(p^{h^2})$. Given these terminal payoffs, we can find a subgame perfect equilibrium $\sigma_{h^\circ}^2$ for $T_{h^\circ}^2$ by backward induction. Once again, $\sigma_{h^\circ}^2$ is not necessarily pure.

Just like in step 1, for each $h^2 \in H^2(h^\circ)$, let $W^2(h^\circ, h^2)$ denote the histories along the play p^{h^2} after the subtree $T_{h^\circ}^2$, and let $W^2(h^\circ)$ be the union of these sets. We define the strategy profile σ at histories belonging to $Z^2(h^\circ)$ and $W^2(h^\circ)$: the strategy profile σ equals $\sigma_{h^\circ}^2$ at histories in $Z^2(h^\circ)$ and follows the ϵ -optimal plays at histories in $W^2(h^\circ)$.

ALL FURTHER STEPS, AND CONCLUSION: By repeating this construction with steps in \mathbb{N} , we eventually consider each history of the game: certain histories belong to a finite subtree and certain histories are ϵ -optimal. This yields a fully specified strategy profile σ . By construction, σ is a subgame perfect ϵ -equilibrium. Indeed, a player cannot gain more than ϵ by deviating along an ϵ -optimal play, and given this, it is never profitable to deviate at histories belonging to a finite subtree.

Moreover, as the prescription for σ along the ϵ -optimal plays does not use randomization, σ is eventually pure in every subgame. This completes the proof of Theorem 5.

5 The topology induced by i -convergence

In this section, we examine the topology induced by the notion of i -convergence. We give criteria for metrizability, compactness, and separability for this topology.

5.1 The topological space (P, \mathcal{T}_i)

In this subsection, we fix a player $i \in I$ and define a topology \mathcal{T}_i on the set of plays and show that a sequence of plays converges to a play p with respect to this topology \mathcal{T}_i exactly when this sequence of plays i -converges to p . We also examine the relationship between the topology \mathcal{T}_i and the topology \mathcal{T} .

For every play $p \in P$ and stage $t \in \mathbb{N}^*$, the set $O_i(p, t)$ is defined in (1). Let \mathcal{T}_i be the topology on P that is induced by the collection of sets $\mathcal{O}_i = \{O_i(p, t) | p \in P, t \in \mathbb{N}^*\}$. That is, \mathcal{T}_i is the smallest topology on P that contains each set belonging to \mathcal{O}_i . As the next lemma shows, the collection \mathcal{O}_i forms a basis of the topology \mathcal{T}_i . This means that for a set $O \subseteq P$, we have $O \in \mathcal{T}_i$ exactly when O can be written as a union of sets belonging to \mathcal{O}_i .

Lemma 12. *The collection \mathcal{O}_i is a basis for the topology \mathcal{T}_i .*

Proof. We only need to show (cf. Aliprantis and Border, page 25) that, for any two sets $O_i(p, s) \in \mathcal{O}_i$ and $O_i(q, t) \in \mathcal{O}_i$, the intersection $O_i(p, s) \cap O_i(q, t)$ can be written as a union of sets in \mathcal{O}_i . Let $O = O_i(p, s) \cap O_i(q, t)$. We can assume without loss of generality that $s \leq t$. We distinguish a number of cases and show in each case that O is indeed such a union.

Case 1: $p = q$. In this case, $O = O_i(q, t)$, so we are done.

Case 2: $p \neq q$. Then $\ell(p, q) < \infty$. We divide this case into three subcases.

Subcase 2.1: $\ell(p, q) < s \leq t$. This case is trivial, as $O = \emptyset$.

Subcase 2.2: $s \leq \ell(p, q) < t$. If $I(p, q) = \{i\}$ then $O = O_i(q, t)$. If $I(p, q) \neq \{i\}$ then $O = \emptyset$.

Subcase 2.3: $s \leq t \leq \ell(p, q)$. If $I(p, q) = \{i\}$ then $O = O_i(q, t)$. If $I(p, q) \neq \{i\}$, then O is the union of the sets $P(ha)$ where the history $h \in H$ and the stage outcome $a \in A$ have the properties: (a) h is a common prefix of p and q with length $t \leq \|h\| \leq \ell(p, q) - 1$, (b) $ha \in H$, (c) at the history h , the stage game outcome $a \in A$ deviates from the common prefix of p and q , but the only difference is caused by player i . Note that each such set $P(ha)$ is a union of sets in \mathcal{O}_i , hence we are done. \square

The next lemma gives some basic properties of the topological space (P, \mathcal{T}_i) .

Lemma 13. *The topological space (P, \mathcal{T}_i) is:*

1. *Hausdorff.*
2. *First countable.*
3. *Sequential.*

Proof.

1. If $p \neq q$ then $\ell(p, q) < \infty$ and hence $O_i(p, \ell(p, q) + 1)$ and $O_i(q, \ell(p, q) + 1)$ are disjoint open sets containing p and q , respectively.
2. Take a play $p \in P$. Then the sequence $O_i(p, 1) \supseteq O_i(p, 2) \supseteq \dots$ is a countable neighborhood basis for p .
3. Follows immediately from (2). \square

Because the topological space (P, \mathcal{T}_i) is sequential, it is fully determined by its convergent sequences. In the following lemma we show that the convergent sequences of (P, \mathcal{T}_i) are precisely the i -convergent sequences, thereby proving that \mathcal{T}_i is indeed the topology induced by i -convergence.

Lemma 14. *A sequence of plays $(p_m)_{m \in \mathbb{N}}$ converges to the play p in the topological space (P, \mathcal{T}_i) if and only if the sequence $(p_m)_{m \in \mathbb{N}}$ i -converges to the play p .*

Proof. Take a sequence of plays $(p_m)_{m \in \mathbb{N}}$ that converges to the play p in the topological space (P, \mathcal{T}_i) . For each $t \in \mathbb{N}^*$, the set $O_i(p, t)$ is an open neighborhood of the play p , i.e. $p \in O_i(p, t) \in \mathcal{T}_i$. Hence, for each $t \in \mathbb{N}^*$, there exists a constant $M_t \in \mathbb{N}$ such that for all $m \geq M_t$ we have $p_m \in O_i(p, M_t)$. Therefore, $\lim_{m \rightarrow \infty} \ell(p_m, p) = \infty$ and for all $m \geq M_1$ we have $I(p_m, p) \subseteq \{i\}$. This means that the sequence $(p_m)_{m \in \mathbb{N}}$ i -converges to the play p .

Conversely, take a sequence of plays $(p_m)_{m \in \mathbb{N}}$ that i -converges to the play p . Consider an open neighborhood O of the play p , i.e. $p \in O \in \mathcal{T}_i$. By Lemma 12, there exists $t \in \mathbb{N}^*$ such that $O_i(p, t) \subseteq O$. Since we assumed that $(p_m)_{m \in \mathbb{N}}$ is i -convergent to p , there exists $M \in \mathbb{N}$ such that for each $m \geq M$ we have (a) $\ell(p_m, p) \geq t$ and (b) $I(p_m, p) \subseteq \{i\}$. Thus, for each $m \geq M$ we have $p_m \in O_i(p, t)$ and in particular $p_m \in O$. Therefore, $(p_m)_{m \in \mathbb{N}}$ converges to p in the topological space (P, \mathcal{T}_i) . \square

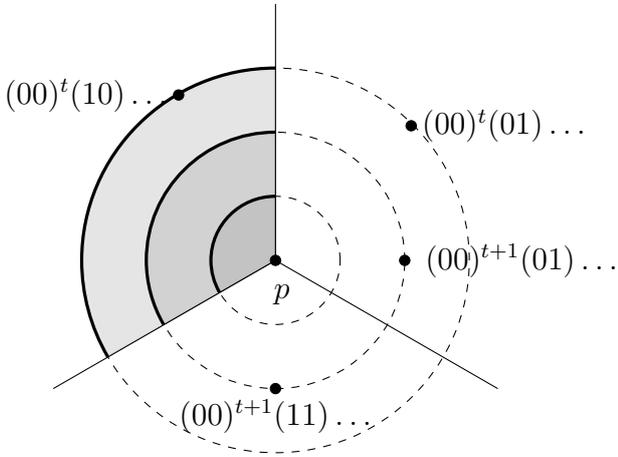


Figure 3: The topology \mathcal{T}_1 .

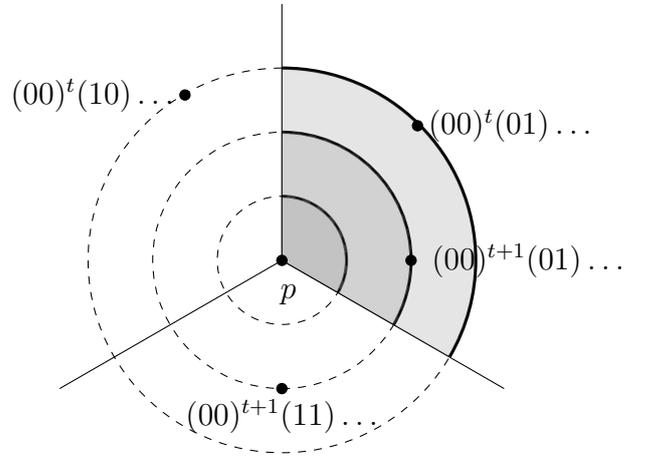


Figure 4: The topology \mathcal{T}_2 .

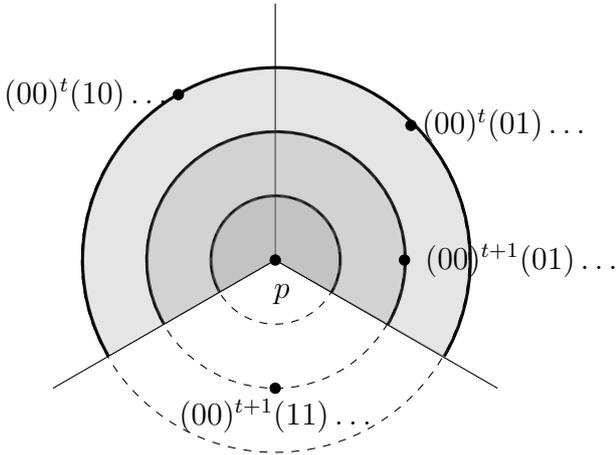


Figure 5: The topology $\mathcal{T}_1 \cap \mathcal{T}_2$.

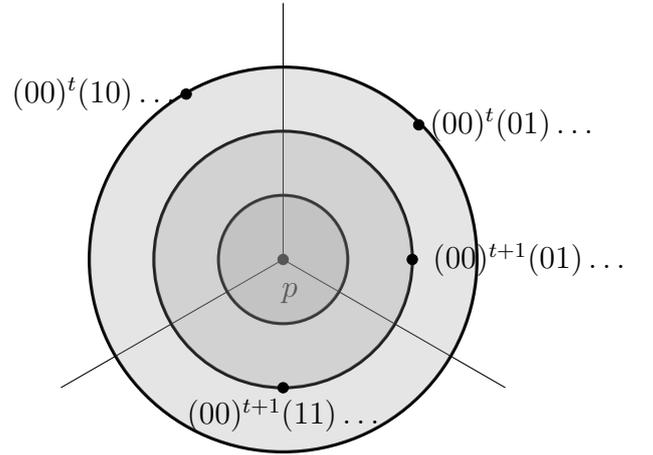


Figure 6: The topology \mathcal{T} .

Figure 7: Local basis of the play $p = (00)^\infty$ in the topology \mathcal{T}_1 , \mathcal{T}_2 , $\mathcal{T}_1 \cap \mathcal{T}_2$ and \mathcal{T} , respectively, for a 2-player multistage game where, for every $h \in H$, for $i = 1, 2$, $A_i(h) = \{0, 1\}$. The Euclidean distance between any play and the play p is inversely related to the length of the longest common history, while the circle sector indicates the subset of players who first deviated.

Lemma 15. *The topology \mathcal{T}_i is finer than the topology \mathcal{T} , that is, $\mathcal{T}_i \supseteq \mathcal{T}$.*

Proof. For every history $h \in H$ and for each play $p \in P(h)$ it holds that $p \in O_i(p, \|h\|) \subseteq P(h)$, and hence we have that $P(h) = \cup_{p \in P(h)} O_i(p, \|h\|)$. Therefore, we have that \mathcal{T}_i contains each set $P(h)$. As the topology \mathcal{T} is induced by these sets, we obtain $\mathcal{T}_i \supseteq \mathcal{T}$. \square

Note that \mathcal{T}_i is strictly larger than \mathcal{T} in certain games. Consider the game in Example 4. The set $O_1(c^\infty, 1)$ belongs to \mathcal{T}_1 and it contains the play c^∞ and all plays in which player 1 quits at stage 3 or later. However, $O_1(c^\infty, 1)$ does not belong to \mathcal{T} for the following reason. There is no history $h \in H$ such that $c^\infty \in P(h) \subseteq O_1(c^\infty, 1)$. As the sets $P(h)$, where $h \in H$, form a basis of the topology \mathcal{T} , we conclude that $O_1(c^\infty, 1) \notin \mathcal{T}$.

Now we turn to the question when \mathcal{T}_i coincides with \mathcal{T} . The answer, and also the answer to many other questions, depends on the size of the action sets along each play. For this purpose, we introduce the notions of i -finiteness and i -cofiniteness of a play.

Definition 16. A play $p \in P$ is *i -finite* if there exists a stage $K \in \mathbb{N}^*$ such that at every prefix of p with length at least K the set of available actions to player i is a singleton, so for every $k \geq K$, $|A_i(p|_k)| = 1$.

Thus, if a play p is i -finite then, after finitely many stages, it is not within the power of player i to deviate from p . Let $F_i = \{p \in P | p \text{ is } i\text{-finite}\}$ denote the subset of i -finite plays of the P .

Definition 17. A play $p \in P$ is *i -cofinite* if there exists a stage $K \in \mathbb{N}^*$ such that at every prefix of p with length at least K the set of available actions to every player $j \neq i$ is a singleton, so for every $k \geq K$, for every $j \neq i$, $|A_j(p|_k)| = 1$.

Thus, if a play p is i -cofinite then, after finitely many stages, player i has full control over the realization of the play p . Let $C_i = \{p \in P | p \text{ is } i\text{-cofinite}\}$ denote the subset of i -cofinite plays of the P .

Note that there can be plays that are neither i -finite nor i -cofinite and that a play can be both i -finite and i -cofinite.

As it turns out, $\mathcal{T}_i = \mathcal{T}$ holds if and only if the topological space (P, \mathcal{T}_i) is compact if and only if each play is i -cofinite.

Proposition 18. *The following statements are equivalent:*

- (1) $\mathcal{T}_i = \mathcal{T}$.
- (2) *The topological space (P, \mathcal{T}_i) is compact.*
- (3) *Each play $p \in P$ is i -cofinite, i.e. $P = C_i$.*

Proof. We prove that (1) and (2) are equivalent and that (1) and (3) are equivalent as well.

(1) implies (2): This is immediate because (P, \mathcal{T}) is compact.

(2) implies (1): As we know, (P, \mathcal{T}) is a compact Hausdorff space. By assumption, (P, \mathcal{T}_i) is compact as well and by Lemma 13 it is Hausdorff too. Because $\mathcal{T}_i \supseteq \mathcal{T}$ due to Proposition 15, the maximality principle of compact Hausdorff spaces implies that $\mathcal{T}_i = \mathcal{T}$.⁵

(1) implies (3): Suppose there exists a play $p \in P$ which is not i -cofinite. Then there is an infinite set $T \subseteq \mathbb{N}^*$ of stages such that at each prefix of p with length $t \in T$ there is a player $j_t \neq i$ with an action set containing at least two alternatives, so $|A_{j_t}(p|_t)| \geq 2$. For each $t \in T$, let p_t be a play in which the first difference from p is caused by the action of player j_t at stage t . Then, the sequence of plays $(p_t)_{t \in T}$ is convergent to the play p , but not i -convergent. Because both \mathcal{T}_i and \mathcal{T} are fully determined by the convergent sequences, this contradicts the assumption that $\mathcal{T}_i = \mathcal{T}$. Consequently, it follows that every play $p \in P$ is i -cofinite.

(3) implies (1): Because both \mathcal{T}_i and \mathcal{T} are determined by their convergent sequences it is sufficient to show that if all plays are i -cofinite then every convergent sequence is also i -convergent. To this end, fix a play $p \in P$ and a sequence $(p_m)_{m \in \mathbb{N}}$ that converges to the play p , i.e. $\lim_{m \rightarrow \infty} p_m = p$. Because the play p is i -cofinite, there exists a time t^* such that for every player $j \neq i$ and every $t \geq t^*$ $|A_j(p|_t)| = 1$. Because the sequence $(p_m)_{m \in \mathbb{N}}$ is convergent to the play p we have that there exists an $M \in \mathbb{N}^*$ such that, for all $m \geq M$, $\ell(p_m, p) \geq t^*$. Furthermore we have that $I(p_m, p) \subseteq \{i\}$ for all $m \geq M$. We can conclude that $\lim_{m \rightarrow \infty}^{(i)} p_m = p$. □

A topological space is called separable if it contains a countable dense subset. The separability of (P, \mathcal{T}_i) depends on the cardinality of the set of plays which are i -finite. The reason is that each i -finite play, as a singleton, is open in \mathcal{T}_i .

Proposition 19. *The topological space (P, \mathcal{T}_i) is separable if and only if the subset of i -finite plays F_i is countable.*

Proof. First notice that, for every $p \in F_i$, the singleton $\{p\}$ is open in \mathcal{T}_i . Indeed, as p is i -finite, there is a stage t such that at any prefix of p with length at least t player i 's action set is a singleton. It follows that $\{p\} = O_i(p, t) \in \mathcal{T}_i$.

Part 1: Assume that (P, \mathcal{T}_i) is separable and let D be a countable dense subset of P . For each $p \in F_i$, the singleton $\{p\}$ is open in \mathcal{T}_i , so we have $F_i \subseteq D$. It follows that F_i is countable.

Part 2: Assume that F_i is countable. We show that (P, \mathcal{T}_i) is separable by constructing a countable dense subset of P .

For every history $h \in H$, fix a play $p^h \in P(h)$. Because the set H of histories is countable and because F_i is countable by assumption, the set $D = \{p^h | h \in H\} \cup F_i$ is countable too. We claim that D is dense in P with respect to \mathcal{T}_i .

⁵If (X, \mathcal{V}) is a compact Hausdorff space and \mathcal{W} is another topology on X such that \mathcal{W} strictly includes \mathcal{V} , then the topological space (X, \mathcal{W}) is not compact. Indeed, take a set $U \in \mathcal{W} \setminus \mathcal{V}$. As U is not open in \mathcal{V} , the set $X \setminus U$ is not closed in \mathcal{V} . Because (X, \mathcal{V}) is Hausdorff, $X \setminus U$ is not compact in \mathcal{V} and hence $X \setminus U$ is not compact in \mathcal{W} . However, by the choice of U , the set $X \setminus U$ is closed in \mathcal{W} . Hence, (X, \mathcal{W}) cannot be compact.

It suffices to show that, for every $p \notin D$, for every stage $t \in \mathbb{N}^*$, $O_i(p, t) \cap D \neq \emptyset$. Take a play $p \notin D$ and a stage $t \in \mathbb{N}^*$. As $p \notin D$ we also have $p \notin F_i$, and hence there exists a prefix h of p such that $\|h\| \geq t$ and $|A_i(h)| \geq 2$. Consider a history h' of length $\|h\| + 1$ such that (a) h' coincides with h at the first $\|h\|$ stages and (b) at stage $\|h\| + 1$, h' differs from the play p only by the action of player i . Then, $p^{h'} \in O_i(p, t) \cap D$. \square

The following proposition shows that (P, \mathcal{T}_i) is not metrizable under mild conditions.

Proposition 20. *If the set $F_i \cup C_i$ is finite and $P \setminus (F_i \cup C_i)$ is infinite, then the topological space (P, \mathcal{T}_i) is not metrizable.*

Proof. Suppose $d_i : P \times P \rightarrow [0, \infty)$ is a metric which induces the topology \mathcal{T}_i .

Step 1: Construction of two sequences of plays $(p_m)_{m \in \mathbb{N}}$ and $(q_m)_{m \in \mathbb{N}}$.

In this step we inductively construct two sequences of plays $(p_m)_{m \in \mathbb{N}}$ and $(q_m)_{m \in \mathbb{N}}$ as illustrated in Figure 8.

Start with any play p_1 which is not i -finite. Because the play p_1 is not i -finite and there are infinitely many plays which are not i -cofinite there exists a play q_1 with the following three properties (1) $I(p_1, q_1) = \{i\}$, (2) q_1 is not i -cofinite, and (3) $d_i(p_1, q_1) < 1/2$. Indeed, the fact that the play p_1 is not i -finite guarantees that there exists an eventually non-constant sequence of plays which i -converges to the play p_1 . This guarantees that this sequence of plays contains an element q_1 having all the desired properties.

Now assume that for some $m \in \mathbb{N}$ the plays p_m and q_m are defined such that p_m is not i -finite and q_m is not i -cofinite. Take any play p_{m+1} that is not i -finite such that p_{m+1} coincides with the play q_m on a longer history than the play p_m and such that the first time the play p_{m+1} differs from q_m is not solely due to player i , so

$$\ell(p_m, q_m) < \ell(p_{m+1}, q_m) \text{ and } I(p_{m+1}, q_m) \setminus \{i\} \neq \emptyset. \quad (2)$$

Because the play p_{m+1} is not i -finite it is possible to choose the play q_{m+1} in such a way that q_{m+1} is not i -cofinite and

$$\ell(p_{m+1}, q_m) < \ell(p_{m+1}, q_{m+1}), \quad I(p_{m+1}, q_{m+1}) = \{i\}, \text{ and } d_i(p_{m+1}, q_{m+1}) < 2^{-(m+1)}. \quad (3)$$

Step 2: Finding a contradiction

For $m \in \mathbb{N}$, let $h_m = \min(p_m, q_m)$ denote the longest common history between the play p_m and the play q_m . Notice that by construction we have $h_1 \prec h_2 \prec \dots$. Therefore there exists a unique play p extending all histories h_m . Note that for every $m \in \mathbb{N}$ we have that

$$\ell(p, p_m) = \ell(p_m, q_m) < \ell(p_{m+1}, q_m) < \ell(p_{m+1}, q_{m+1}) = \ell(p, p_{m+1}).$$

It follows that $\lim_{m \rightarrow \infty} \ell(p, p_m) = \infty$. Furthermore, we have that $I(p, p_m) = I(p_m, q_m) = \{i\}$. Therefore, we can conclude that $\lim_{m \rightarrow \infty}^{(i)} p_m = p$ and hence $\lim_{m \rightarrow \infty} d_i(p, p_m) = 0$. Furthermore, we have by construction that $\lim_{m \rightarrow \infty} d_i(p_m, q_m) = 0$. It follows from the triangle inequality that, for every

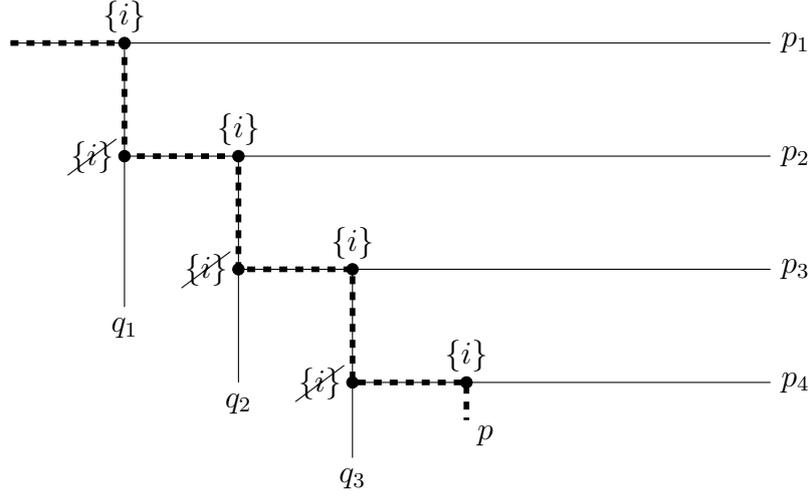


Figure 8: The non-metrizability of the topology of i -convergence.

$m \in \mathbb{N}$, $d_i(p, q_m) \leq d_i(p, p_m) + d_i(p_m, q_m)$. This yields that $\lim_{m \rightarrow \infty} d_i(p, q_m) = 0$ and consequently $\lim_{m \rightarrow \infty}^{(i)} q_m = p$. However, by construction we have that, for every $m \in \mathbb{N}$, $I(p, q_m) = I(p_m, q_m) \not\subseteq \{i\}$, so we have obtained a contradiction. \square

Example 21. Proposition 20 implies that the topological space (P, \mathcal{T}_i) induced by an infinitely repeated stage game in which all players have at least two actions is not metrizable. Indeed, every play in such a game is neither i -finite nor i -cofinite and there are infinitely many plays.

Even though in general the topological space (P, \mathcal{T}_i) is not metrizable, there are games for which the associated topological space is metrizable.

Lemma 22. *If $P = F_i \cup C_i$, then the topological space (P, \mathcal{T}_i) is metrizable.*

Proof. Let $F_{i,k}$ denote the set of i -finite plays such that $A_i(p_{|t})$ is a singleton for each $t \geq k$. Thus $F_{i,1}, F_{i,2}, \dots$ is a non-decreasing sequence of sets converging to the set of i -finite plays F_i . Let $\delta(p) = 0$ for each $p \notin F_i$ and let $\delta(p) = 2^{-k}$ for each $p \in F_{i,k} \setminus F_{i,k-1}$.

Define $d_i(p, p) = 0$ for each $p \in P$, and for two distinct plays $p, q \in P$ we let $d_i(p, q) = \max\{d(p, q), \delta(p), \delta(q)\}$, where d is the usual ultrametric defined by $d(p, q) = 2^{-\ell(p, q)}$. We now show that because d is an ultrametric, d_i is as well. By definition we have that $d_i(p, p) = 0$ for all $p \in P$, furthermore it is trivial to see that d_i is non-negative and symmetric. It remains to show that $d_i(p, q) \leq \max\{d_i(p, r), d_i(r, q)\}$ for any $p, q, r \in P$. We have

$$\begin{aligned} d_i(p, q) &\leq \max\{d(p, r), d(r, q), \delta(p), \delta(q)\} \\ &\leq \max\{d(p, r), d(r, q), \delta(p), \delta(q), \delta(r)\} \\ &= \max\{d_i(p, r), d_i(r, q)\}. \end{aligned}$$

We now show that the metric d_i induces the topology \mathcal{T}_i of i -convergence.

Part 1: If $\lim_{m \rightarrow \infty} d_i(p_m, p) = 0$ then the sequence $(p_m)_{m \in \mathbb{N}}$ i -converges to p .

Suppose the sequence $(p_m)_{m \in \mathbb{N}}$ converges to p under the metric d_i . First assume that $p \in F_i$ then because p is i -finite there exists $M \in \mathbb{N}$ such that $p_m = p$ for all $m \geq M$, which implies that $I(p_m, p) = \emptyset$ and $\ell(p_m, p) = \infty$ for all $m \geq M$. We can conclude that $(p_m)_{m \in \mathbb{N}}$ i -converges to p . Now assume that $p \in C_i$ and note that because $d \leq d_i$ we have that the sequence $(p_m)_{m \in \mathbb{N}}$ converges to p under the metric d . Furthermore observe that any sequence of plays that converges to $p \in C_i$ under the metric d also i -converges to p .

Part 2: If the sequence $(p_m)_{m \in \mathbb{N}}$ i -converges to p then $\lim_{m \rightarrow \infty} d_i(p_m, p) = 0$

If the sequence $(p_m)_{m \in \mathbb{N}}$ i -converges to p , then $\lim_{k \rightarrow \infty} d(p_k, p) = 0$ and there exists an $M \in \mathbb{N}$ such that $I(p_m, p) \subseteq \{i\}$ for every $m \geq M$. Let $n_m = \ell(p_m, p)$. Then for each $m \geq M$ we have that either $p_m = p$ or the set $A_i(p_m|n_k) = A_i(p|n_m)$ is not a singleton. From this we can conclude that p and p_m are not elements of F_{i, n_m} . Therefore $\delta(p) \leq 2^{-n_m}$ and $\delta(p_m) \leq 2^{-n_m}$. Because $\lim_{m \rightarrow \infty} n_m = \infty$ and $\lim_{m \rightarrow \infty} d(p_m, p) = 0$ we conclude that p_m converges to p under the metric d_i . \square

5.2 The topological space $(P, \cap_{i \in I} \mathcal{T}_i)$

In this section, we take a closer look at the collection $\cap_{i \in I} \mathcal{T}_i$ of sets which are open in every topological space (P, \mathcal{T}_i) . By Proposition 15 we have $\mathcal{T}_i \supseteq \mathcal{T}$ and hence $\cap_{i \in I} \mathcal{T}_i \supseteq \mathcal{T}$. It is now a natural question to ask whether in fact $\cap_{i \in I} \mathcal{T}_i = \mathcal{T}$. As we will see, this is the case for perfect information games, but not in general.

For every play $p \in P$ and stage $t \in \mathbb{N}^*$, let

$$\begin{aligned} O^*(p, t) &= \cup_{i \in I} O_i(p, t) \\ &= \{q \in P \mid \ell(p, q) \geq t \text{ and } |I(p, q)| \leq 1\}. \end{aligned}$$

Let \mathcal{T}^* be the topology on P that is induced by the collection $\mathcal{B}^* = \{O^*(p, t) \mid p \in P, t \in \mathbb{N}^*\}$. That is, \mathcal{T}^* is the smallest topology on P that contains each set belonging to \mathcal{B}^* . As the next lemma shows, the collection \mathcal{B}^* forms a basis of the topology \mathcal{T}^* . This means that for a set $O \subseteq P$, we have $O \in \mathcal{T}^*$ exactly when O can be written as a union of sets in \mathcal{B}^* .

Lemma 23. *The collection \mathcal{B}^* is a basis for the topology \mathcal{T}^* .*

Proof. The proof is similar to that of Lemma 12. We only need to show (cf. Aliprantis and Border, page 25) that, for any two sets $O^*(p, s) \in \mathcal{B}^*$ and $O^*(q, t) \in \mathcal{B}^*$, the intersection $O^*(p, s) \cap O^*(q, t)$ can be written as a union of sets in \mathcal{B}^* . Let $O^* = O^*(p, s) \cap O^*(q, t)$. We can assume without loss of generality that $s \leq t$. We distinguish a number of cases and prove in each case that O^* is indeed such a union.

Case 1: $p = q$. In this case $O^* = O^*(q, t)$ and we are done.

Case 2: $p \neq q$. It follows that $\ell(p, q) < \infty$. We divide this case into three subcases.

Subcase 2.1: $\ell(p, q) < s \leq t$. In this case it holds that $O^* = \emptyset$.

Subcase 2.2: $s \leq \ell(p, q) < t$. If $I(p, q)$ is a singleton, then $O^* = O^*(q, t)$. If $I(p, q)$ is not a singleton, so multiple players cause the first difference between p and q , then $O^* = \emptyset$.

Subcase 2.3: $s \leq t \leq \ell(p, q)$. If $I(p, q)$ is a singleton, then $O^* = O^*(p, s) = O^*(q, t)$. If $I(p, q)$ is not a singleton, then O^* is the union of the sets $P(ha)$ where the history $h \in H$ and the stage outcome $a \in A$ have the properties: (a) h is a common prefix of p and q with length $t \leq \|h\| \leq \ell(p, q) - 1$, (b) $ha \in H$, (c) at the history h , the stage game outcome $a \in A$ deviates from the common prefix of p and q and the difference is caused by exactly one player. Because for every history $h \in H$ we have $P(h) = \cup_{p \in P(h)} O_i(p, \|h\|)$, each such set $P(ha)$ is a union of sets in \mathcal{B}^* , and hence we are done. \square

We now show that the topologies \mathcal{T}^* and $\cap_{i \in I} \mathcal{T}_i$ coincide. Hence, \mathcal{B}^* is a basis for the topology $\cap_{i \in I} \mathcal{T}_i$.

Lemma 24. $\mathcal{T}^* = \cap_{i \in I} \mathcal{T}_i$.

Proof.

Part 1: $\mathcal{T}^* \supseteq \cap_{i \in I} \mathcal{T}_i$.

Let $O \in \cap_{i \in I} \mathcal{T}_i$ be given. It is sufficient to show that for every $p \in O$ there exists $t \in \mathbb{N}^*$ such that $O^*(p, t) \subseteq O$. Let some $p \in O$ be given. For every $i \in I$ there exists $t_i \in \mathbb{N}^*$ such that $O_i(p, t_i) \subseteq O$. For $t = \max_{i \in I} t_i$ it holds that, for every $i \in I$, $O_i(p, t) \subseteq O$. It follows that $O^*(p, t) = \cup_{i \in I} O_i(p, t) \subseteq O$.

Part 2: $\mathcal{T}^* \subseteq \cap_{i \in I} \mathcal{T}_i$.

Fix a play $p \in P$ and a time $t \in \mathbb{N}^*$. It is sufficient to prove that, for every $i \in I$, $O^*(p, t) \in \mathcal{T}_i$. Fix some player $i \in I$. Let some $q \in O^*(p, t)$ be given. It is sufficient to show that there exists $k \in \mathbb{N}^*$ such that $O_i(q, k) \subseteq O^*(p, t)$. If $q = p$, then let $k = t$, so we have that $O_i(p, t) \subseteq \cup_{j \in I} O_j(p, t) = O^*(p, t)$ and we are done. If $q \neq p$, then there exists $k \in \mathbb{N}$ such that $k > \ell(p, q) \geq t$. It is immediate that $O_i(q, k) \subseteq O^*(p, t)$. \square

Because in perfect information games at every stage only one player can deviate from a given play, we obtain the following result.

Proposition 25. Consider a game G as defined in Section 2 with perfect information. It holds that $\mathcal{T}^* = \mathcal{T}$.

Proof. Because G has perfect information, we have for every play $p \in P$ and stage $t \in \mathbb{N}^*$ that $O^*(p, t) = \{q \in P \mid \ell(p, q) \geq t\} = P(p_t)$. Hence \mathcal{T}^* and \mathcal{T} have the same basis and the statement follows. \square

Corollary 26. If a sequence of plays is convergent in \mathcal{T}^* , then it is convergent in \mathcal{T} as well. The converse holds for perfect information games.

6 Discussion

6.1 Perfect information games

In Theorem 5, we assumed that the set of available actions is always finite. This has two consequences. First, when applying backward induction in the proof of Theorem 5, we were guaranteed to have a subgame perfect equilibrium in each finite tree. Second, each stopping time τ_ϵ^k is bounded.

The second consequence is not essential. Even with unbounded, but finite, stopping times we would obtain subtrees in the proof of Theorem 5 that have no infinite branches. Hence, it would still be possible to apply backward induction by the following idea. Start at the root of a subtree. If there is a child history that is non-terminal, then take this history, and repeat this process as long as it is possible to choose a non-terminal child history. Since there is no infinite branch, this process will eventually stop at a history such that all of its children are terminal histories. Then, apply a step of backward induction at this history. By means of a transfinite procedure, applying a step of backward induction at each iteration step, we finally end up with a complete strategy profile for this subtree, and this strategy profile is thus a subgame perfect equilibrium for this subtree.

The first consequence is important though. Still, it would be enough to have a one-shot ϵ -equilibrium in each possible stage game, for every $\epsilon > 0$. In particular, this would be the case if the game under consideration has perfect information, even if the set of available actions is not finite. To be more precise, this leads to the following statement.

Consider a perfect information game that satisfies the assumptions of Section 2, without the requirement that the set of available actions of each player is finite. If the payoff function u_i of every player $i \in I$ is bounded and i -upper semicontinuous, then for each $\epsilon > 0$, the game admits a pure strategy profile such that in any subgame, no player can gain more than ϵ by unilaterally deviating to another pure strategy.

This statement generalizes the corresponding result in Purves and Sudderth [17], by relaxing the topological condition on the payoff functions.

6.2 Necessity of the conditions in Theorem 5

In this subsection, we discuss to which extent the conditions assumed in Theorem 5 are necessary for the existence of a subgame perfect ϵ -equilibrium, for every $\epsilon > 0$.

The assumption that the set of available actions is always finite was already discussed in the previous subsection.

We also assumed in Theorem 5 that the payoffs are bounded. This is a standard assumption, and without this assumption even very simple 1-player games fail to have ϵ -optimal solutions.

The assumption of individual upper semicontinuity is of course the main condition on the payoffs and is heavily used.

We also assumed that the payoffs are Borel measurable. The subgame perfect ϵ -equilibria that we constructed in Theorem 5 are eventually pure in every subgame. Hence, the calculation of the

expected payoffs for these strategy profiles, or even if a player deviates to a pure strategy, does not require the payoffs to be Borel measurable. Rather, Borel measurability is needed to be able to calculate the expected payoffs when a player deviates to a non-pure strategy.

The assumption that the number of players is finite was used to obtain the crucial Corollary 8.

6.3 Payoff functions that are i -lower semicontinuous

It is not known whether all games as defined in Section 2, where each player's payoff function is bounded and lower semicontinuous admits an ϵ -equilibrium, for every $\epsilon > 0$.⁶

Nevertheless, using i -convergent sequences we can define the notion of i -lower semicontinuity. A payoff function u_i is i -lower semicontinuous if for every play p and every sequence $(p_m)_{m \in \mathbb{N}}$ of plays that is i -convergent to p we have $\liminf_{m \rightarrow \infty} u_i(p_m) \geq u_i(p)$. With this concept, for perfect information games we can conclude the following statement. Flesch et al. (2010) showed that in every perfect information game, with arbitrary action spaces, if each player's payoff function is bounded and lower semicontinuous then there exists a pure strategy profile such that in any subgame, no player can gain more than ϵ by unilaterally deviating to another pure strategy. After a careful look at the proof, the only place where the lower semicontinuity of the payoff functions is used is on page 750 in order to prove that Equation (18) follows from Equation (19). Since the sequence of plays considered there is not only convergent but also i -convergent, the result of Flesch et al. (2010) remains valid even if lower semicontinuity is relaxed to i -lower semicontinuity.

6.4 Stochastic transitions

It is a natural question whether or not Theorem 5 can be extended to games with stochastic transitions. Even if this is possible, the proof would be substantially more complex, as we would not be able to work with ϵ -optimal plays any more.

References

- [1] Bruyère, V. (2017). Computer Aided Synthesis: A Game-Theoretic Approach. In *Developments in Language Theory* (pp. 3-35). Springer International Publishing.
- [2] Bruyère, V., Le Roux, S., Pauly, A., & Raskin, J. F. (2017). On the existence of weak subgame perfect equilibria. In *International Conference on Foundations of Software Science and Computation Structures* (pp. 145-161). Springer, Berlin, Heidelberg.
- [3] Bogachev, V. (2007). *Measure Theory*. Springer-Verlag.

⁶It is not even known in the context of quitting games, see for instance Solan and Solan (2017) and the references therein.

- [4] Flesch, J., Kuipers, J., Mashiah-Yaakovi, A., Schoenmakers, G., Shmaya, E., Solan, E. & Vrieze, K. (2014). Non-existence of subgame-perfect ε -equilibrium in perfect information games with infinite horizon. *International Journal of Game Theory*, 43(4), 945-951.
- [5] 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.
- [6] Flesch, J., & Predtetchinski, A. (2016). Subgame-perfect ε -equilibria in perfect information games with common preferences at the limit. *Mathematics of Operations Research*, 41(4), 1208-1221.
- [7] Flesch, J., Thuijsman, F., & Vrieze, K. (1997). Cyclic Markov equilibria in stochastic games. *International Journal of Game Theory*, 26(3), 303-314.
- [8] Fudenberg, D., & Levine, D. (1983). Subgame-perfect equilibria of finite and infinite horizon games. *Journal of Economic Theory*, 31(2), 251-268.
- [9] Harris C., Reny P., & Robson, A. (1995). The existence of subgame-perfect equilibrium in continuous games with almost perfect information: a case for public randomization. *Econometrica*, 63(3), 507-544.
- [10] Jaśkiewicz, A., & Nowak, A. (2016). Non-Zero-Sum Stochastic Games. In *Handbook of Dynamic Game Theory* (pp. 1-64). Springer International Publishing.
- [11] Kechris, A. (1995). *Classical Descriptive Set Theory*. Springer-Verlag.
- [12] Le Roux, S., & Pauly, A. (2016). Equilibria in multi-player multi-outcome infinite sequential games. Preprint: Retrieved from *arXiv:1401.3325*.
- [13] Maitra, A. P., & Sudderth, W. D. (2007). Subgame-perfect equilibria for stochastic games. *Mathematics of Operations Research*, 32(3), 711-722.
- [14] Mertens, J. F., & Parthasarathy, T. (2003). Equilibria for discounted stochastic games. In: *Stochastic Games and Applications* (pp. 131-172). Springer.
- [15] Munkres, J. R. (2000). *Topology*. Prentis Hall.
- [16] Prokopovych, P., & Yannelis, N. C. (2017). On strategic complementarities in discontinuous games with totally ordered strategies. *Journal of Mathematical Economics*, 70, 147-153.
- [17] Purves, R. A. & Sudderth, W. D. (2011). Perfect information games with upper semicontinuous payoffs. *Mathematics of Operations Research*, 36(3), 468-473.
- [18] Secchi, P., & Sudderth, W. D. (2001). N -Person stochastic games with upper semi-continuous payoffs. *International Journal of Game Theory*, 30(4), 491-502.

- [19] Solan, E., & Solan, O. (2017). Quitting games and linear complementarity problems. Preprint: Retrieved from *arXiv:1707.02598*.
- [20] Solan, E., & Vieille, N. (2001). Quitting games. *Mathematics of Operations Research*, 26(2), 265-285.