Non-Existence of Subgame-Perfect ε -Equilibrium in Perfect Information Games with Infinite Horizon

Citation for published version (APA):

Flesch, J., Kuipers, J., Mashiah-Yaakovi, A., Schoenmakers, G., Shmaya, E., Solan, E., & Vrieze, O. (2014). Non-Existence of Subgame-Perfect ε -Equilibrium in Perfect Information Games with Infinite Horizon. *International Journal of Game Theory*, 43(4), 945-951. <https://doi.org/10.1007/s00182-014-0412-3>

Document status and date:

Published: 01/01/2014

DOI:

[10.1007/s00182-014-0412-3](https://doi.org/10.1007/s00182-014-0412-3)

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.

Non-existence of subgame-perfect ε -equilibrium in perfect information games with infinite horizon

János Flesch · Jeroen Kuipers ·
Ayala Mashiah-Yaakovi · Gijs Schoenmakers ·
Eran Shmaya · Eilon Solan · Koos Vrieze

Accepted: 21 January 2014 / Published online: 13 February 2014
© Springer-Verlag Berlin Heidelberg 2014

Abstract Every finite extensive-form game with perfect information has a subgame-perfect equilibrium. In this note we settle to the negative an open problem regarding the existence of a subgame-perfect ε -equilibrium in perfect information games with infinite horizon and Borel measurable payoffs, by providing a counter-example. We also consider a refinement called strong subgame-perfect ε -equilibrium, and show by means of another counter-example, with a simpler structure than the previous one,

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

J. Kuipers · G. Schoenmakers · K. Vrieze
Department of Knowledge Engineering, Maastricht University,
P.O. Box 616, 6200 MD Maastricht, The Netherlands
e-mail: kuipers@maastrichtuniversity.nl

G. Schoenmakers
e-mail: gm.schoenmakers@maastrichtuniversity.nl

K. Vrieze
e-mail: koosvrieze@gmail.com

A. Mashiah-Yaakovi · E. Shmaya · E. Solan (✉)
School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69978, Israel
e-mail: eilons@post.tau.ac.il

A. Mashiah-Yaakovi
e-mail: ayalam@post.tau.ac.il

E. Shmaya
Kellogg School of Management, Northwestern University,
2001 Sheridan Road, Evanston, IL 60208, USA
e-mail: e-shmaya@kellogg.northwestern.edu

that a game may have no strong subgame-perfect ε -equilibrium for sufficiently small $\varepsilon > 0$, even though it admits a subgame-perfect ε -equilibrium for every $\varepsilon > 0$.

Keywords Subgame-perfect equilibrium · Perfect-information games · Infinite horizon · Non-existence

1 Introduction

We examine perfect information games in which the active player always has a finite number of actions. The most common refinement of the concept of Nash equilibrium in such games is subgame-perfect equilibrium (cf. [Selten 1965](#)), which is a strategy profile that induces a Nash equilibrium in every subgame. This concept rules out the use of noncredible threats by the players.

As is well known, when the game tree is finite, a subgame-perfect equilibrium always exists and can be found by backward induction, going from the leaves of the tree back to the root. In perfect information games with infinite horizon, there are no leaves and this process is not applicable. Yet, under certain continuity assumptions on the payoffs, a subgame-perfect equilibrium can be shown to exist (cf. for example [Fudenberg and Levine 1983](#)). For infinite horizon games with discontinuous payoffs, this is no longer the case and a subgame-perfect equilibrium, or even a Nash equilibrium, need not exist. Indeed, consider a one-player decision problem in which at every period the player chooses whether to stop or to continue. The payoff is $1 - \frac{1}{t}$ if the player stops at period t , and the payoff is 0 if the player never stops. In this game the player wants to stop in finite time but as late as possible. Consequently, this game has no Nash equilibrium.

Discontinuous payoffs and the nonexistence of Nash equilibria is not a phenomenon that is unique to artificially constructed games. For example, in models of price competition (cf. [Bertrand 1883](#)), spatial competition (cf. [Hotelling 1929](#)), auctions (cf. [Milgrom and Weber 1982](#)), and dynamic oligopoly (cf. [Fudenberg 1983](#)), the existence of a Nash equilibrium, and therefore also a subgame-perfect equilibrium, is not guaranteed.

In games with discontinuous payoffs we often resort to the concept of ε -equilibrium, which is a strategy profile such that no player can profit more than ε by deviating. Though the concept of ε -equilibrium is less common than Nash equilibrium, it makes sense in various situations. First, if no Nash equilibrium exists, it seems reasonable to forgo a small extra profit if this ensures stability. Second, if payoffs in the model are only approximations of actual payoffs, for example, due to measurement errors or simplifying assumptions, it does not seem reasonable to change a strategy for a small extra profit below the margin of error in the payoffs. There are also other reasons to consider ε -equilibria. For example, [Kalai and Lehrer \(1993\)](#) proved that in repeated games, the process of rational learning converges to an ε -equilibrium. Mertens and Neyman (see [Mertens 1987](#)) proved that in games with perfect information, an ε -equilibrium in pure strategies exists as soon as the payoff functions of the players are bounded and Borel measurable.

A subgame-perfect ε -equilibrium is a strategy profile that induces an ε -equilibrium in every subgame. Solan and Vieille (2003) provide an example of a two-player game with perfect information that has no subgame-perfect ε -equilibrium in pure strategies, provided ε is sufficiently small, but does have a subgame-perfect ε -equilibrium in behavior strategies, for every $\varepsilon > 0$. Nevertheless, in recent years much effort has been made to prove the existence of subgame-perfect ε -equilibrium in pure strategies, for all $\varepsilon > 0$, in various classes of perfect information games (cf. for example, Flesch et al. 2010 and Purves and Sudderth 2011).

The definition of subgame-perfect ε -equilibrium has one drawback: the requirement that the strategy profile induces an ε -equilibrium in a certain subgame does not rule out the possibility that a player plays, with small probability, an action that leads to a low continuation payoff. If one interprets a mixed action as a lottery that the player uses for the choice of an action, and the lottery picks an action with low payoff, then the player may become reluctant to execute this action. Weaker versions of subgame-perfect ε -equilibrium have also been examined in the literature, such as ε -consistent equilibrium by Lehrer and Sorin (1998), but this drawback still appears in those definitions. This is the motivation to examine the following refinement of subgame-perfect ε -equilibrium. A strategy profile is called a *strong subgame-perfect ε -equilibrium* if the following property holds at every decision node and every action that is played with positive probability at this node: for the active player, this action in combination with his follow-up strategy is an ε -best reply against the strategies of the opponents.

In this note we settle to the negative the open problem of whether every perfect information game with infinite horizon and Borel measurable payoffs has a subgame perfect ε -equilibrium. We provide an example of such a game that does not have a subgame-perfect ε -equilibrium, and a second example of a game that has a subgame-perfect ε -equilibrium but no strong subgame-perfect ε -equilibrium.

2 Model

We consider n -player perfect information games in extensive form given by an infinite directed tree with a root, in which each node has finitely many outgoing edges. Each node in the tree is associated with one of the players, the player who acts at that node. The outgoing edges at a node are interpreted as the available actions of the player associated with that node. An infinite path in the tree, starting at the root, is called a *play*. Each player i has a payoff function u_i , which assigns a payoff $u_i(p) \in \mathbb{R}$ to each play p . The interpretation is that player i receives the amount $u_i(p)$ in case the play p is realized. We assume that each u_i is bounded and measurable with respect to the sigma-algebra $\mathcal{P}$ that is generated by the cylinder sets on the set of plays. Let $u = (u_1, \dots, u_n)$ be the vector of payoff functions.

The set of all decision nodes that are associated with player i is denoted by Z_i . A (behavior) strategy for player i is a function β_i that assigns to each node $z \in Z_i$ a probability distribution over the available actions at z . A strategy is called *pure* if it selects at each decision node one of the available actions with probability 1. Every strategy profile $\beta = (\beta_1, \dots, \beta_n)$ induces a probability distribution ρ_β on $\mathcal{P}$. The

corresponding expected payoff is given by $u(\beta) = \int u(p) \, d\rho_\beta(p)$. For every strategy profile β and every player i , we will use the notation $\beta_{-i} = (\beta_j)_{j \neq i}$ for the strategy profile of the opponents of player i .

Definition 1 Let $\varepsilon \geq 0$. A strategy profile β^* is an ε -equilibrium if $u_i(\beta^*) \geq u_i(\beta_i, \beta_{-i}^*) - \varepsilon$ for every player i and every strategy β_i of player i .

Thus, a strategy profile β^* is an ε -equilibrium if no player can gain more than ε by a unilateral deviation. Mertens and Neyman (cf. Mertens 1987) proved that a pure ε -equilibrium exists for every $\varepsilon > 0$.

Because the game has perfect information, every decision node z in the game tree naturally defines a game $\Gamma(z)$, the subgame that starts at z . For every strategy β_i of player i and every decision node z , denote by $\beta_i|z$ the strategy that β_i induces in the subgame $\Gamma(z)$. For every strategy profile β denote $\beta|z := (\beta_i|z)_{i=1}^n$. We denote by $u_i(\beta|z)$ player i 's expected payoff in $\Gamma(z)$ under the strategy profile $\beta|z$.

We now define the concept of subgame-perfect ε -equilibrium (cf. Selten 1975), which is a refinement of ε -equilibrium.

Definition 2 Let $\varepsilon \geq 0$. A strategy profile β^* is a *subgame-perfect ε -equilibrium*, if for every decision node z in the game tree, the strategy profile $\beta^*|z$ is an ε -equilibrium of the subgame $\Gamma(z)$.

As we have already mentioned, this definition does not rule out the possibility that a player plays with a small probability an action that leads to a low payoff. Hence, we also examine the following refinement of this concept.

Definition 3 Let $\varepsilon \geq 0$. A strategy profile β^* is a *strong subgame-perfect ε -equilibrium* if for every player i , for every decision node $z \in Z_i$ of player i , for every action a_i that is chosen with positive probability at z according to β_i^* , and for every strategy β_i of player i , we have

$$u_i((\beta_i^*[z : a_i], \beta_{-i}^*) | z) \geq u_i((\beta_i, \beta_{-i}^*) | z) - \varepsilon, \quad (1)$$

where $\beta_i^*[z : a_i]$ denotes the strategy for player i which follows β_i^* except at node z where action a_i is chosen.

In words, a strategy profile is a strong subgame-perfect ε -equilibrium if the following property holds at every decision node and every action that is played with positive probability at this node: for the active player, this action in combination with his follow-up strategy is an ε -best reply against the strategies of the opponents. It is clear that every strong subgame-perfect ε -equilibrium is a subgame-perfect ε -equilibrium. A sufficient condition for the converse is given in the following lemma, whose proof follows from the definitions.

Lemma 4 Suppose that β^* is a subgame-perfect ε -equilibrium, and let $\delta \geq 0$. If for every player i , for every decision node $z \in Z_i$ of player i , and for every action a_i that is chosen with positive probability at z according to β_i^* we have

$$u_i((\beta_i^*[z : a_i], \beta_{-i}^*) | z) \geq u_i(\beta^* | z) - \delta,$$

then β^* is a strong subgame-perfect $(\varepsilon + \delta)$ -equilibrium. Here again, $\beta_i^*[z : a_i]$ denotes the strategy for player i which follows β_i^* except at node z where action a_i is chosen.

Note that if β^* is a strong subgame-perfect 0-equilibrium, and for some player i , β_i is a strategy that, at every node $z \in Z_i$, assigns probability 1 to an action that is used with positive probability by β_i^* , then it is not necessarily true that β_i is a best reply to β_{-i}^* . This is illustrated by the following one-player game: at every period Player 1 can choose between L and R , and his payoff is 1 if he chooses L infinitely often, and 0 otherwise. In this game, the strategy that always chooses both actions with probability $\frac{1}{2}$ constitutes a strong subgame-perfect 0-equilibrium, but the strategy that always chooses R is not a strong subgame-perfect 0-equilibrium.

3 The examples

In this section we present two examples. In Example 5, which is based on Solan and Vieille (2003), the game admits a subgame-perfect ε -equilibrium for every $\varepsilon > 0$ but not a strong subgame-perfect ε -equilibrium for sufficiently small ε . In Example 6 the game does not even admit a subgame-perfect ε -equilibrium for sufficiently small ε .

Example 5 Consider the following two-player perfect information game. Alice can stop or continue at even periods $0, 2, 4, \dots$, and Bob can stop or continue at odd periods $1, 3, 5, \dots$. As soon as a player stops, the game terminates (one could assume that play goes on without any further influence on the payoffs, so as to make the horizon infinite). If Alice stops then the payoff is -1 to Alice and 2 to Bob. If Bob stops then the payoff is -2 to Alice and 1 to Bob. If nobody ever stops then the payoff is zero to both players.

As Solan and Vieille (2003) pointed out, this game has a subgame-perfect ε -equilibrium for every $\varepsilon \in (0, 1)$, given by the following strategies $\beta^* = (\beta_A^*, \beta_B^*)$: Under β_A^* , Alice stops with probability 1 at every even period, and under β_B^* Bob stops with probability ε at every odd period. These strategies induce payoffs $u(\beta^* | z) = (-1, 2)$ in every subgame $\Gamma(z)$ that starts at one of Alice's decision nodes, and payoffs $u(\beta^* | z) = (-1 - \varepsilon, 2 - \varepsilon)$ in every subgame $\Gamma(z)$ that starts at one of Bob's decision nodes. Profitable deviations exist only in subgames that start at one of Bob's decision nodes; in such a subgame, Bob increases his payoff only if he stops with probability less than ε , but the improvement is bounded from above by ε .

Notice that the above subgame-perfect ε -equilibrium is not strong, since Bob's strategy uses the noncredible threat of stopping (inequality (1) fails for Bob's action "stop"). In fact, we now prove that this game has no strong subgame-perfect ε -equilibrium for a sufficiently small ε .

Let $0 < \varepsilon < \frac{1}{7}$ and assume that $\beta^* = (\beta_A^*, \beta_B^*)$ is a subgame-perfect ε -equilibrium. Our goal is to prove that β^* cannot be strong. Note that in any subgame $\Gamma(z)$, Bob's payoff satisfies $u_B(\beta^* | z) \geq 1 - \varepsilon$, since he can guarantee at least 1 by the strategy that stops at the next odd period. Assume first that there is a subgame $\Gamma(z)$ in which (1) Alice is the active player at the initial node z , and (2) under $\beta^*|z$, the probability that Alice eventually stops is at most ε . Since $u_B(\beta^* | z) \geq 1 - \varepsilon$, the probability that Bob eventually stops under $\beta^*|z$ is at least $1 - 3\varepsilon$, and therefore Alice's payoff

satisfies $u_A(\beta^* | z) \leq -2(1 - 3\varepsilon) < -1 - \varepsilon$. But then Alice would gain more than ε by stopping at node z , which is a contradiction. Therefore, there is no subgame $\Gamma(z)$ as above.

Thus, if Alice is active at node z , then the probability that she eventually stops in the subgame $\Gamma(z)$ is at least ε . It follows that for every decision node z in the game tree, it holds with probability 1 under $\beta^*|_z$ that one of the players eventually stops. Consequently, the payoff $u(\beta^* | z)$ is a convex combination of $(-1, 2)$ and $(-2, 1)$. For every decision node z in which Alice is the active player, we have $u_A(\beta^* | z) \geq -1 - \varepsilon$. Consequently, $u_B(\beta^* | z) \geq 2 - \varepsilon$ for each such node z . It follows that for every decision node z in which Bob is the active player, he can get a payoff of at least $2 - \varepsilon$ by continuing, and therefore $u_B(\beta^* | z) \geq 2 - 2\varepsilon$ for every such decision node z .

Now assume by way of contradiction that the subgame-perfect ε -equilibrium β^* is strong. Since $u_B(\beta^* | z) \geq 2 - 2\varepsilon$ for every decision node z in which Bob is the active player, β_B^* always has to assign probability zero to stopping. But then Alice could gain 1 by always continuing, which is a contradiction.

The following example shows that a subgame-perfect ε -equilibrium need not exist for small $\varepsilon > 0$.

Example 6 Consider the following perfect information game with infinite horizon, played by Alice and Bob. At each node of the game tree, the active player has two actions, A and B . If action A is played then Alice is the next active player, whereas if action B is played then Bob is the next active player. The payoff function is as follows:

- If there is a period t such that Alice is the active player at all periods from t onwards, then the payoff is $(-1, 2)$. When this happens, we say that play *absorbs at Alice*.
- If there is a period t such that Bob is the active player at all periods from t onwards, then the payoff is $(-2, 1)$. When this happens, we say that play *absorbs at Bob*.
- Otherwise, the payoff is $(0, 0)$.

Whereas in Example 5 punishment involves stopping once, the main idea of this example is that Bob can only punish Alice by always playing B from some period onward. We will show that the game has no subgame-perfect ε -equilibrium for $\varepsilon \in [0, \frac{1}{7})$. Assume to the contrary that such a subgame-perfect ε -equilibrium β^* exists. Then, just as for Example 5 with “absorbing” replacing “stopping”, we can prove that $u_B(\beta^* | z) \geq 2 - 2\varepsilon$ for every decision node z at which Bob is the active player. Thus, the probability to ever absorb at Bob under $\beta^*|_z$ is at most 2ε , for every such node z . This implies that under $\beta^*|_z$ the probability of absorption at Bob is 0. Hence, Alice can get payoff 0 in any subgame by always playing action B . Thus, $u_A(\beta^* | z) \geq -\varepsilon$ for every decision node z . This is however in contradiction with the fact that $u_B(\beta^* | z) \geq 2 - 2\varepsilon$ for every decision node z at which Bob is the active player. Thus, there exists no subgame-perfect ε -equilibrium for $\varepsilon \in [0, \frac{1}{7})$, as claimed.

Remark 1 Notice that the game of Example 6 possesses the following pure equilibrium: If play starts at Alice, then she always plays A and Bob always plays B . If play starts at Bob, then Alice always plays A and Bob initially plays A and afterwards he always plays B .

Remark 2 The structure of the game of Example 6 is more complicated than that of Example 5: the game of Example 5 can be seen as a two-player stochastic game with finitely many states and actions in which the players receive stage payoffs and each player maximizes his long-run average stage payoff (rather than having each player receive a final payoff that he maximizes). This is however not possible for the game of Example 6. Indeed, in the game of Example 6, after every history there are plays that yields each of the payoffs $(-1, 2)$, $(-2, 1)$, and $(0, 0)$. This implies that if there were such a stochastic game, then there would have been a play along which the sequence of average payoffs fluctuates between, say, $(-1, 2)$ and $(0, 0)$, which has no equivalent in the game of Example 6.

Moreover, regarding the payoff structure, let $U(-1, 2)$ (resp. $U(-2, 1)$, $U(0, 0)$) be the set of all plays that induce the payoff $(-1, 2)$ (resp. the payoff $(-2, 1)$, the payoff $(0, 0)$). In Example 5, $U(-1, 2)$ and $U(-2, 1)$ are open sets, whereas $U(0, 0)$ is closed (when considering the topology generated by the cylinder sets on the space of all plays). In contrast, in Example 6, $U(-1, 2)$ and $U(-2, 1)$ are F_σ sets, i.e., countable unions of closed sets, whereas $U(0, 0)$ is a G_δ set, i.e., a countable intersection of open sets.

Acknowledgments The research of Mashiah-Yaakovi has received funding from the European Research Council under the European Community's Seventh Framework Programme (FP7/2007-2013)/ ERC Grant agreement n [249159]. The research of Solan was supported in part by the Israel Science Foundation (Grant number 212/09) and by the Google Inter-university center for Electronic Markets and Auctions.

References

- Bertrand J (1883) Theorie mathematique de la richesse sociale. J des Savants 67:499–508
- Flesch J, Kuipers J, Mashiah-Yaakovi A, Schoenmakers G, Solan E, Vrieze K (2010) Perfect-information games with lower-semicontinuous payoffs. Math Operations Res 35:742–755
- Fudenberg D, Levine D (1983) Subgame-perfect equilibria of finite and infinite-horizon games. J Econ Theory 31:251–268
- Fudenberg D, Tirole J (1983) Dynamic models of oligopoly, Stanford University Workshop discussion paper
- Hotelling H (1929) The stability of competition. Econ J 39:41–57
- Kalai E, Lehrer E (1993) Rational learning leads to Nash equilibrium. Econometrica 61:1019–1045
- Lehrer E, Sorin S (1998) ϵ -Consistency in repeated games. Int J Game Theory 27:231–244
- Mertens, J-F (1987) Repeated games. In: Proceedings of the international congress of mathematicians, (Berkeley, California, 1986), American Mathematical Society, Providence, RI, 1987, pp 528–577
- Milgrom P, Weber R (1982) A theory of auctions and competitive bidding. Econometrica 50:1089–1122
- Purves RA, Sudderth WD (2011) Perfect information games with upper-semicontinuous payoffs. Math Operations Res 36:468–473
- Selten R (1965) Spieltheoretische behandlung eines oligopolmodells mit nachfragentragheit. Zeitschrift fur die gesamte Staatswissenschaft 12:201–324
- Selten R (1975) Reexamination of the perfectness concept for equilibrium points in extensive games. Int J Game Theory 4:25–55
- Solan E, Vieille N (2003) Deterministic multi-player Dynkin games. J Math Econ 39:911–929