On refinements of subgame perfect -equilibrium

CrossMark

On refinements of subgame perfect ϵ-equilibrium

János Flesch · Arkadi Predtetchinski

Abstract The concept of subgame perfect ϵ-equilibrium (ϵ-SPE), where ϵ is an error-term, has in recent years emerged as a prominent solution concept for perfect information games of infinite duration. We propose two refinements of this concept: continuity ϵ-SPE and ϕ-tolerance equilibrium. A continuity ϵ-SPE is an ϵ-SPE in which, in any subgame, the induced play is a continuity point of the payoff functions. We prove that continuity ϵ-SPE exists for each $\epsilon > 0$ if the payoff functions are bounded and lower semicontinuous. A loss tolerance function ϕ is a function that assigns to each history h a positive real number $\phi(h)$. A strategy profile is said to be a ϕ-tolerance equilibrium if for each history h it is a $\phi(h)$-equilibrium in the subgame starting at h. We prove that, for each loss tolerance function ϕ, there exists a ϕ-tolerance equilibrium provided that the payoff functions are bounded and continuous. We give counterexamples to show the sharpness of the existence results.

1 Introduction

Perfect information games are routinely used in various areas of economics such as bargaining theory (e.g. Binmore et al. 1992) and coalition formation (e.g. Ray 2007),

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

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

as well as in logic and computer science (e.g. Apt and Grädel 2011). We examine a rich class of perfect information games with deterministic transitions. Such a game can be given by a finite set of players and a directed tree with a root and no terminal nodes. Each node in the tree is associated with a player, who controls this node. Play of the game starts at the root. At every node that play visits, the player who controls this node has to choose one of the outgoing arcs, which brings play to the next node. This induces an infinite path, called play, in the tree, and depending on this play, each player receives a payoff.

This setup encompasses all games of finite duration (as we can replace every terminal node with one infinite sequence of arcs). Another important special case is the situation when the players receive instantaneous payoffs at every period of the game and then aggregate them into one payoff, for example by taking the discounted sum.

In these games, subgame perfect equilibrium (cf. Selten 1965), SPE for short, is the most common refinement of Nash equilibrium. A strategy profile is called an SPE if it induces a Nash equilibrium in every subgame. More precisely, for every node in the tree, given play reaches this node from the root, the continuation strategies form a Nash equilibrium.

Various conditions have been found which guarantee the existence of an SPE (cf. for example Fudenberg and Levine 1983; Harris 1985; Maitra and Sudderth 2007). Yet, without strong conditions, an SPE can easily fail to exist, even in one-player games. Suppose for example that a player has to choose a natural number, and his payoff is $1 - \frac{1}{n}$ if he chooses $n \in \mathbb{N}$. In this game, the player has obviously no optimal strategy, and therefore the game admits no Nash equilibrium and no SPE either. Or as a similar example, suppose that a player can choose to stop or to continue at time periods $t \in \mathbb{N}$, and his payoff is $1 - \frac{1}{t}$ if he decides to stop the game at period t and his payoff is 0 if he never stops. A common feature of these two games is that, although the player has no optimal strategy, he does have approximately optimal strategies at his disposal. Indeed, for every error-term $\epsilon > 0$, if the player chooses an $n \geq \frac{1}{\epsilon}$ in the first game or stops at a period $t \geq \frac{1}{\epsilon}$ in the second game, then his payoff is at least $1 - \epsilon$. For a small error-term ϵ, these approximately optimal strategies are fairly satisfactory, and offer a remedy to the non-existence of optimal solutions.[1]

It has been discovered in recent years that there are various classes of perfect information games in which, despite the non-existence of SPE, one can construct approximate SPEs with arbitrarily small errors. The main definition of approximate SPE is called subgame perfect ϵ-equilibrium, ϵ-SPE for short, where $\epsilon > 0$, which is a strategy profile that induces an ϵ-equilibrium in every subgame.[2] Here, an ϵ-equilibrium is a strategy profile such that no player can gain more than ϵ by a unilateral deviation. The

---

[1] Approximate solutions in games are in general not only justified by the non-existence of exact solutions. In applications, we often need to make simplifying assumptions and use approximations of payoffs, either because we do not know the precise data, or just because we try to keep the model tractable. In such situations, an approximate solution with a very small error-term is arguably just as good as an exact solution. Besides, approximate solutions sometimes have an easier form than the exact ones, and sometimes they can be calculated in an easier fashion.

[2] An alteratively approach is to require perfection in a sufficiently "large" set of subgames, as is done in Himmelberg et al. (1976), Lehrer and Sorin (1998), and Mashiah-Yaakovi (2014).

concept of ϵ-SPE has appealing properties. The definition is natural, and as mentioned above, it often provides a way out when an exact SPE fails to exist.

Existence of ϵ-SPE has been shown in various classes of games with perfect information: e.g. in Carmona (2005) for games with bounded continuous at infinity payoffs, in Flesch et al. (2010) for games with bounded lower semicontinuous payoffs, in Purves and Sudderth (2011) for games with bounded upper semicontinuous payoffs, in Flesch et al. (2013) for free transition games, in Solan and Vieille (2001, 2003), Solan (2005) and Mashiah-Yaakovi (2009) for quitting games. Laraki et al. (2013) consider zero-sum games with semicontinuous payoffs, and prove that both players have ϵ-optimal strategies and one of the players even has a subgame perfect ϵ-optimal strategy.

Yet, there are reasons to impose additional requirements and look for refinements. Flesch et al. (2014) pointed out the drawback of ϵ-SPE that it does not rule out that a player chooses, with small probability, an action that leads to a bad payoff. This causes a problem if one interprets a probability distribution, which a strategy prescribes on the actions, as a lottery that the player uses when choosing an action. Indeed, in the unlikely but possible event that the lottery picks an action with low payoff, the player may become reluctant to execute this action. Flesch et al. (2014) proposed and examined a refinement, called strong ϵ-SPE, which avoids this shortcoming. As another refinement, Brihaye et al. (2013) investigated so-called secure SPE, which can be applied for assume-guarantee synthesis and model checking. Secure refers to a property that the players, in some sense, do not have to fear deviations when an opponent changes his strategy to another one which gives this opponent the same payoff. An analog of this concept would be a secure ϵ-SPE .

1.1 Our contribution

In this paper, we identify further properties that one may wish to require from a strategy profile in addition to being an ϵ-SPE, and propose two refinements of the concept ϵ-SPE: continuity ϵ-SPE and ϕ-tolerance equilibrium. We also examine their properties and their existence, mainly when the payoff functions are bounded and semicontinuous.

1.2 Continuity ϵ-SPE

When the payoff functions are not fully continuous, it can happen that an ϵ-SPE induces an infinite play which is not a continuity point of the payoff functions. This has the consequence that arbitrarily small perturbations in the strategies, even at arbitrarily far time periods, can lead to drastic changes in payoffs. This is particularly problematic if the players are not absolutely certain about their opponents strategies, or when strategies cannot be executed with exact precision.

Therefore, we examine a refinement of ϵ-SPE, called continuity ϵ-SPE, which requires from the ϵ-SPE that, in every subgame, the induced play is a continuity point of the payoff functions. This concept has an appealing finitistic nature, due to the additional continuity property: in any subgame, if such a profile is followed then eventually a time period is reached after which the payoffs are essentially fixed, i.e. any other

continuation profile would give almost the same payoff. Roughly speaking, the game is strategically over in finite time.

Of course, not all games admit a continuity ϵ-SPE for every $\epsilon > 0$. But we prove this to be the case when the payoff functions are bounded and lower semicontinuous. Bounded and upper semicontinuous payoffs are not sufficient, as a counter-example demonstrates.

1.3 ϕ-tolerance equilibrium

The mistake of a strategy profile in a certain subgame is defined as the maximal improvement in this subgame that a player can gain by a unilateral deviation from his strategy. With this terminology, an ϵ-SPE is a strategy profile for which the mistake is at most ϵ in every subgame.

It may however be reasonable to also require that the mistakes in the ensuing subgames converge to 0 as play progresses. The need for this type of refinement was already recognized by Mailath et al. (2005), who introduced the concept of contemporaneous perfect ϵ-equilibrium, in the more restricted context of games in which each player receives a payoff at every period of the game and aims at maximizing his total discounted payoff. In a subgame of such a discounted game, with a first node at a far period, the payoffs are discounted heavily, and therefore every strategy profile of the original game automatically induces an ϵ-equilibrium in this subgame. Consequently, ϵ-SPE does not guarantee acceptable outcomes and realistic behavior in far subgames in such discounted games. The main point is that not only the payoffs, but the maximally allowed mistake ϵ should also be discounted.

There is in fact another motivation to assume that the mistakes tend to zero as time progresses. Since a subgame is smaller than the original game, one could argue that in certain situations it should be easier for the players to overview the strategic possibilities in a subgame and to reason about them than in the original game. Hence, it may be reasonable to assume that the players only tolerate smaller mistakes once play reaches later subgames.

With these motivations in mind, let ϕ be a function that assigns a positive number $\phi(h)$ to every history h. Any such function is called a loss tolerance function, and the value $\phi(h)$ is called a tolerance level at h. A strategy profile is called a ϕ-tolerance equilibrium if for every history h it induces a $\phi(h)$-equilibrium in the subgame that begins at h. We obtain ϵ-SPE as a special case by taking ϕ to be a constant function assigning the tolerance level of ϵ to each history. We are most interested in cases where tolerance levels converge to zero as the length of the history increases. We prove that, if the payoff functions are bounded and continuous, then a ϕ-tolerance equilibrium exists for every loss tolerance function. Counter-examples demonstrate that continuity cannot be replaced with lower or upper semicontinuity.

1.4 Vanishing loss

At this point we have introduced two refinements of ϵ-SPE that at first glance might appear to be unrelated. Yet, common to the two concepts is an intuitive property that we call vanishing loss.

A strategy profile is said to exhibit vanishing loss if the maximal improvement any player can gain by unilaterally deviating from the given strategy approaches zero as the play proceeds. We do not consider the property of vanishing loss as an independent refinement. However, we do think that this is a noteworthy property. In fact, much of the motivation we have given for our refinements stems from vanishing loss.

1.5 Structure of the paper

In Sect. 2, we define the model precisely. Then, we discuss continuity ϵ-SPE in Sect. 3. We turn to loss tolerance equilibrium in Sect. 4, and prove the existence result in Sect. 5. Section 6 concludes with the discussion of vanishing loss.

2 The model

2.1 The game

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

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

2.2 Strategies

A pure strategy for player i is a function $\sigma_i : \iota^{-1}(i) \to A$, where $\iota^{-1}(i)$ is the set of histories where player i moves. As we only consider pure strategies, we omit the qualification "pure" in the sequel. A strategy profile is a tuple $(\sigma_1, \ldots, \sigma_n)$ where each σ_i is a strategy for player i. Given a strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ and a strategy η_i for player i we write σ/η_i to denote the strategy profile obtained from σ by replacing σ_i with η_i.

For the concatenation of histories and actions we use the following notations. For a history $h = (a_0, \ldots, a_t) \in H$ and an action $b \in A$, we denote the sequence $(a_0, \ldots, a_t, b)$ by (h, b) or simply hb, and for a sequence of actions $(b_0, b_1, \ldots)$ in A let $(h, b_0, b_1, \ldots) = (a_0, \ldots, a_t, b_0, b_1, \ldots)$.

We can identify a strategy profile σ with a function $\sigma : H \to A$. Define the play induced by σ starting from the history h, denoted by $\pi(\sigma; h)$, inductively as follows: Let $h_0 = h$. If h_t has been defined for some $t \geq 0$, then let $a_t = \sigma(h_t)$ and set

$h_{t+1} = (h_t, a_t)$. Then $\pi(\sigma; h) = (h, a_0, a_1, \dots)$. For the special case $h = \emptyset$, we simply write $\pi(\sigma) = \pi(\sigma; \emptyset)$.

2.3 Subgame-perfect ϵ-equilibrium

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

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

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

$$u_i(\pi(\sigma; h)) \geq u_i(\pi(\sigma/\sigma_i'; h)) - \epsilon.$$

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

2.4 The topological structure

Following, among others, Martin (1975), Fudenberg and Levine (1983), Flesch et al. (2010), Purves and Sudderth (2011), we endow the set A with the discrete topology and $A^{\mathbb{N}}$ with the product topology. The topology on $A^{\mathbb{N}}$ is completely metrizable,[3] and a basis of this topology is formed by the cylinder sets $O(h) = \{p \in A^{\mathbb{N}} : h < p\}$ for $h \in H$, where for a history $h \in H$ and a play $p \in A^{\mathbb{N}}$ we write $h < p$ if h is the initial segment of p. Thus a sequence of plays $(p_n)_{n \in \mathbb{N}}$ converges to a play p precisely when for every $k \in \mathbb{N}$ there exists an $N_k \in \mathbb{N}$ such that p_n coincides with p on the first k coordinates for every $n \geq N_k$.

A function $f : A^{\mathbb{N}} \to \mathbb{R}$ is said to be *continuous* at a play $p \in A^{\mathbb{N}}$ if, for every sequence of plays $(p_n)_{n \in \mathbb{N}}$ converging to p, we have $\lim_{n \to \infty} f(p_n) = f(p)$. Thus, f is continuous at p precisely when for every $\delta > 0$ there is an $N_\delta \in \mathbb{N}$ such that if a play q coincides with p on the first N_δ coordinates then $|f(p) - f(q)| \leq \delta$. Further, f is said to be continuous if it is continuous at each play in $A^{\mathbb{N}}$.

Continuity of a payoff function with respect to the chosen topology on $A^{\mathbb{N}}$ is implied by a condition known as continuity at infinity, see e.g. Fudenberg and Levine (1983). The condition only says that the actions taken at distant stages of the game have little effect on the payoff. For example, a discounted sum of bounded instantaneous payoffs automatically satisfies continuity at infinity.

---

[3] One can take for example the metric $d : A^{\mathbb{N}} \times A^{\mathbb{N}} \to \mathbb{R}$ which is defined for each $p, q \in A^{\mathbb{N}}$ as follows: if $p = q$ then $d(p, q) = 0$, and otherwise $d(p, q) = 2^{-m(p,q)}$ where $m(p, q) \in \mathbb{N}$ is the first coordinate on which p and q differ.

A function $f : A^{\mathbb{N}} \to \mathbb{R}$ is said to be *lower semicontinuous*, lsc for short, at a play $p \in A^{\mathbb{N}}$ if, for every sequence of plays $(p_n)_{n \in \mathbb{N}}$ converging to p, we have $\liminf_{n \to \infty} f(p_n) \geq f(p)$. We say that f is lsc if it is lsc at each play in $A^{\mathbb{N}}$. Equivalently, f is lsc if the set $\{q \in A^{\mathbb{N}} : f(q) > \alpha\}$ is open in $A^{\mathbb{N}}$ for every $\alpha \in \mathbb{R}$.

A function $f : A^{\mathbb{N}} \to \mathbb{R}$ is said to be *upper semicontinuous*, usc for short, at a play $p \in A^{\mathbb{N}}$ if, for every sequence of plays $(p_n)_{n \in \mathbb{N}}$ converging to p, we have $\limsup_{n \to \infty} f(p_n) \leq f(p)$. We say that f is usc if it is usc at each play in $A^{\mathbb{N}}$. Equivalently, f is usc if the set $\{q \in A^{\mathbb{N}} : f(q) < \alpha\}$ is open in $A^{\mathbb{N}}$ for every $\alpha \in \mathbb{R}$.

2.5 Existence results

Mertens and Neyman (cf. Mertens 1987) showed, by using the result of Martin (1975), that if each player's payoff function is bounded and Borel measurable then an ϵ-equilibrium exists for every $\epsilon > 0$. An ϵ-SPE does not necessarily exist under these conditions, which was demonstrated by an example in Flesch et al. (2014). Nevertheless, semicontinuity proved to be a useful condition. Flesch et al. (2010) proved that an ϵ-SPE exists for every $\epsilon > 0$ when the payoffs are bounded and lsc, whereas Purves and Sudderth (2011) proved the same when the payoffs are bounded and usc. Very general topological conditions for existence of SPE are given in Alós-Ferrer and Ritzberger (2013).

3 Continuity ϵ-SPE

In this section, we examine a refinement of ϵ-SPE, which requires that, in every subgame, the payoff functions are continuous at the induced play. We start with the following illustrative example.

Example 3.1 Consider a 1-player game where $A = \{0, 1\}$ and the payoff $u_1(p)$ is 1 if $p = 1^\infty = (1, 1, 1, \dots)$ and is 0 otherwise. Any SPE (and in fact any equilibrium) induces the play 1^∞ from the root. Such an SPE however does not guarantee within finite time that the payoff will be 1, or at least close to 1, and the player has to play action 1 consistently at every period. To be more precise, for any error-term $\delta \in [0, 1)$, the player is not guaranteed a payoff within a distance of δ from 1 by following the strategy for a finite number of periods. Consequently, any mistake, even at far periods, will be costly. The exact cause of this problem lies in the fact that the payoff function is discontinuous at the play p (even though the payoff function is usc).

Now consider the same game, but now with a different payoff function. At every period of the game, the player receives a payoff of 1 if he chooses action 1 and receives a payoff of 0 if he chooses action 0. The payoff for a play $q = (a_0, a_1, \dots)$ is then the total discounted payoff with discount factor $\beta \in (0, 1)$, i.e.

$$u_1(q) = (1 - \beta) \sum_{t=0}^{\infty} \beta^t a_t. \tag{3.1}$$

Clearly, the unique SPE is to choose action 1 at every period, which again induces the play $p = 1^\infty$ from the root with payoff $u_1(p) = 1$. But now, in contrast with the

previous payoff function, u_1 is continuous at p (and in fact continuous everywhere). And indeed, for any error-term $\delta > 0$, the player can guarantee a payoff of at least $1 - \delta$ in finite time, by simply playing action 1 at the first N_δ periods, for a large enough N_δ. A similar property holds in every subgame.

We call a play p of the game a continuity play if the payoff function of every player is continuous at p. Otherwise, the play is called a discontinuity play.

Definition 3.2 Let $\epsilon \geq 0$. A strategy profile σ is called a continuity strategy profile if σ induces a continuity play in every subgame, i.e. $\pi(\sigma; h)$ is a continuity play for every history $h \in H$. A continuity strategy profile σ that is an ϵ-SPE is called a continuity ϵ-SPE.

Suppose that σ is a continuity ϵ-SPE in a game G. Then σ enjoys the following property: for any $\delta > 0$ and any history $h \in H$ there is a period $N_{\delta,h}$ such that if, starting at h, the players follow σ until period $N_{\delta,h}$, then regardless how they play after period $N_{\delta,h}$, their payoffs will be within a distance of δ from $u(\pi(\sigma; h))$. This means that, in every subgame, the payoffs are essentially fixed after a finite number of periods (up to payoff 2δ), and hence the game effectively ends in finite time. In contrast, when the induced play of the game fails to be a continuity play of the payoff functions, the game never settles, and the players may have to play the entire infinite sequence of actions with full precision to be sure about the payoffs that they receive.

Now we will examine the existence of continuity ϵ-SPE when the payoff functions are semicontinuous. As Example 3.1 demonstrated, bounded and usc payoffs are not sufficient to guarantee the existence. It turns out however that bounded and lsc payoffs are sufficient, and we will prove it below.

Example 3.3 We construct a game with bounded and lsc payoffs having an SPE such that *all* payoff functions are discontinuous at the equilibrium play. There are two players, and the action set is $A = \{L, R\}$. The players move alternatingly, starting from player 1. Player 1's payoff function only depends on actions of player 2. Player 1's payoff is 0 if player 2 plays L all the time, and it is 1 otherwise. Symmetrically, player 2's payoff only depends on player 1's actions:

$$u_1(a_0, a_1, a_2, \dots) = \begin{cases} 0 & \forall t \in \mathbb{N}, \quad a_{2t+1} = L \\ 1 & \text{otherwise.} \end{cases}$$

$$u_2(a_0, a_1, a_2, \dots) = \begin{cases} 0 & \forall t \in \mathbb{N}, \quad a_{2t} = L \\ 1 & \text{otherwise.} \end{cases}$$

Both payoff functions are lsc. Consider the strategy profile σ such that $\sigma(h) = L$ for each $h \in H$. It is an SPE, and it induces the play $(L, L, L, \dots)$, which is a discontinuity point of both payoff functions. Note that there is another SPE in which both players play R at every history. This is clearly a continuity SPE.

The result below strengthens the main theorem in Flesch et al. (2010). A set $D \subset A^{\mathbb{N}}$ is dense if it intersects every non-empty open subset of $A^{\mathbb{N}}$, or equivalently, if it intersects every cylinder set $O(h)$.

Theorem 3.4 *Suppose that all payoff functions are bounded and lsc. Let D be a subset of $A^{\mathbb{N}}$ and let $\epsilon > 0$. The game admits an ϵ-SPE σ for which $\pi(\sigma; h) \in D$ for each $h \in H$ if and only if D is dense in $A^{\mathbb{N}}$.*

Proof $\Longrightarrow$: Suppose that σ is an ϵ-SPE (or in fact any strategy profile) such that $\pi(\sigma; h) \in D$ for each $h \in H$. Take the cylinder set $O(h) = \{p \in A^{\mathbb{N}} : h < p\}$ corresponding to an arbitrary history $h \in H$. Since $\pi(\sigma; h) \in D \cap O(h)$, we have $D \cap O(h) \neq \emptyset$. Because the cylinder sets form a basis of the topology of $A^{\mathbb{N}}$, the set D is dense in $A^{\mathbb{N}}$ indeed.

$\Longleftarrow$: Suppose that D is dense in $A^{\mathbb{N}}$. The proof is nearly identical to the proof of Theorem 2.3 in Flesch et al. (2010). The only alteration we introduce is in the definition of $P_0(h)$. We replace formula (2) in Flesch et al. (2010) by the following:

$$P_0(h) = \{p \in D : h < p\}.$$

The rest of the proof remains the same, word for word. $\square$

If player i's payoff function u_i is bounded and lsc, then the set of continuity points D_i of u_i is a dense subset of $A^{\mathbb{N}}$ (see for example Fort 1951). By taking $D = \cap_{i \in N} D_i$, Theorem 3.4 yields the following stronger result.

Corollary 3.5 *Suppose that all payoff functions are bounded and lsc. Then, for every $\epsilon > 0$, the game admits a continuity ϵ-SPE.*

4 Loss tolerance equilibrium

An ϵ-SPE allows a player to play a suboptimal strategy in any subgame, provided that the payoff loss the player incurs does not exceed a fixed margin of ϵ. We could say that a payoff loss of ϵ is tolerated in every subgame.

In games with discounted payoffs, ϵ-SPE places no restrictions on behavior at remote decision nodes. As an illustration, consider the following example.

Example 4.1 Consider a 1-player game where $A = [0, 1)$ and the payoff function is given for a play $q = (a_0, a_1, \ldots)$ by the total discounted payoff with discount factor $\beta \in (0, 1)$, i.e.

$$u_1(q) = (1 - \beta) \sum_{t=0}^{\infty} \beta^t a_t. \tag{4.1}$$

The game admits no SPE but does admit an ϵ-SPE for each $0 < \epsilon < 1$, for example playing $1 - \epsilon$ in each period. However, there is also an arguably unnatural ϵ-SPE where the player plays $1 - \frac{1}{2}\epsilon$ in the first t periods and plays 0 thereafter, with t being large enough to satisfy $\beta^t < \frac{\epsilon}{2}$. To see that such a strategy is indeed an ϵ-SPE simply notice

that after period t the distance between any two feasible payoffs is at most ϵ, and so any course of action is compatible with the concept of ϵ-SPE.

Same observations apply to all games with continuous payoffs. As the play of the game proceeds, the set of feasible payoffs shrinks, and so a period is eventually reached after which all feasible payoffs differ by no more than ϵ. In any such subgame, ϵ-SPE places no restrictions.

These considerations motivate us to abandon constant loss tolerance and introduce a loss tolerance function assigning a tolerance level to each decision node. Clearly we are most interested in the case where tolerance level declines as the game proceeds.

Definition 4.2 A *tolerance function* is any function $\phi : H \rightarrow (0, +\infty)$. The value $\phi(h)$ of the tolerance function is called the *tolerance level* at h. Given a loss tolerance function ϕ, the strategy profile σ is said to be a ϕ-*tolerance equilibrium* if for each player $i \in N$, each strategy η_i of player i and each history $h \in H$ it holds that

$$u_i(\pi(\sigma/\eta_i; h)) - u_i(\pi(\sigma; h)) \leq \phi(h). \tag{4.2}$$

Equivalently, σ is a ϕ-tolerance equilibrium if for each $h \in H$ it induces a $\phi(h)$-equilibrium in the subgame that starts at h. The most interesting case is the one where ϕ is decreasing, in the sense of assigning smaller tolerance levels to longer histories. If $\phi(h) \leq \epsilon$ for each h then clearly each ϕ-tolerance equilibrium is an ϵ-SPE. With $\phi(h) = \beta^{|h|}$, where β is a discount factor, our ϕ-tolerance equilibrium is closely related to the concept of contemporaneous equilibrium of Mailath et al. (2005).

The differences between three concepts: SPE, ϵ-SPE, and loss tolerance equilibrium are illustrated by the following example.

Example 4.3 This is a 3-period game. Formally, the game is infinite, but only the actions a_0, a_1, and a_2 matter for the payoffs. In period 0 player 1 can choose to play *in* or *out*. If player 1 chooses *out* the game ends and both players receive payoff zero. If player 1 chooses *in*, the game continues and in period 1 player 1 chooses a positive integer a_1. In period 2 player 2 can play *yes* or *no*. If player 2 plays *yes* the payoffs are $(1, 0)$, if *no* the payoffs are $(-\frac{1}{a_1}, \frac{1}{a_1})$. Thus formally $\iota(\emptyset) = \iota(in) = 1$, and $\iota(in, a_1) = 2$, and the payoffs are given by

$$(u_1(a_0, a_1, a_2), u_2(a_0, a_1, a_2)) = \begin{cases} (0, 0) & a_0 = out \\ (1, 0) & a_0 = in, \quad a_2 = yes \\ \left(-\frac{1}{a_1}, \frac{1}{a_1}\right) & a_0 = in, \quad a_2 = no, \end{cases}$$

where $a_1 \in \{1, 2, \dots\}$.

Intuitively we should expect that in this game player 2 always plays *no* and player 1 chooses *out*. Turning to the formal analysis, notice that the game has no SPE. Indeed, if player 2 plays *no* at history (in, a_1) for each a_1, then player 1 at history (in) has no best response.

Fix an $n \geq 1$. The game has a $\frac{1}{n}$-SPE. For example the strategy profile

$$\sigma(\emptyset) = out, \quad \sigma(in) = n, \quad \sigma(in, a_1) = no \quad \text{for each } a_1 \in \{1, 2, \dots\}$$

is a $\frac{1}{n}$–SPE. However, in a $\frac{1}{n}$–SPE it is not necessarily the case that player 2 always chooses *no*. For example, the strategy profile

$$\sigma'(\emptyset) = in, \quad \sigma'(in) = n, \quad \sigma'(in, a_1) = \begin{cases} no & a_1 < n \\ yes & a_1 \geq n \end{cases}$$

is also a $\frac{1}{n}$–SPE. The reason is that player 2, even though he has an exact best response at each of his decision nodes, is allowed to play suboptimally provided that he incurs a payoff loss of not more than $\frac{1}{n}$.

Now consider the loss tolerance function ϕ given by $\phi(\emptyset) = \phi(in) = \frac{1}{n}$ and $\phi(in, a_1) = \frac{1}{2a_1}$. It is clear that σ is a ϕ-tolerance equilibrium while σ' is not. In fact, in each ϕ-tolerance equilibrium with ϕ specified above, player 2 always plays *no* and player 1 chooses *out*.

The main result of this section is the following theorem.

Theorem 4.4 *Suppose that all payoff functions are bounded and continuous. Then, for each loss tolerance function ϕ, the game admits a ϕ-tolerance equilibrium.*

The novelty of Theorem 4.4 is that it covers games where the action set A is infinite. For the special case when A is finite, the conclusion of the theorem follows from the fact that for such games an SPE exist, by the results of Fudenberg and Levine (1983) and Harris (1985). From the technical point of view, the main challenge is that the space of plays $A^{\mathbb{N}}$ is not compact unless A is finite. It is for this reason that the truncation technique of Fudenberg and Levine (1983) does not suffice to prove Theorem 4.4. Instead, we carry out an iterative procedure similar to that in Flesch et al. (2010). It differs from the one in Harris (1985) in that the latter terminates in countably many steps, whereas ours is transfinite. For more discussion on our proof see Sect. 5.

Thus Theorem 4.4 is particularly useful in games where the payoffs are continuous at infinity but where the action set is not compact, as is the case in Example 4.3. Another important class of games that is covered by the theorem are discounted games where the instantaneous reward function is discontinuous in actions.

To conclude this section we argue, by means of two examples, that loss tolerance equilibrium may fail to exist if the payoff functions are only semicontinuous.

*Example 4.5* Consider the following 1-player game with lsc payoffs. The action set is $A = \{s, c\}$ where s stands for *stop* and c for *continue*. If player 1 plays s in period t the game stops[4] with payoff $\frac{t}{t+1}$. If the player never plays action s his payoff is 0.

---
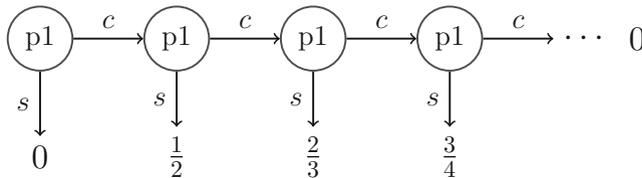
[4] Strictly speaking the game goes on, but once action s has been played, subsequent actions do not affect the payoff.

Define the loss tolerance function ϕ by letting $\phi(t) = \frac{1}{t+2}$ where we write t to denote the history reached if action c is played t times. We show that this game admits no ϕ-tolerance equilibrium. Suppose on the contrary that σ is a ϕ-tolerance equilibrium. First assume that $\sigma(t) = s$ for some time t. Then the payoff for the strategy σ at t is $\frac{1}{t+1}$. However, because $1 - \frac{t}{t+1} > \phi(t)$, the player could improve this payoff at t by more than $\phi(t)$ by playing c sufficiently many times and then playing s. Hence we conclude that $\phi(t) = c$ for each t. But then the payoff for σ is 0 and it could be improved by more than $\phi(0)$ at the beginning of the game by playing c twice and then s.

Example 4.6 The following is a game with usc payoffs having no ϕ-tolerance equilibrium:
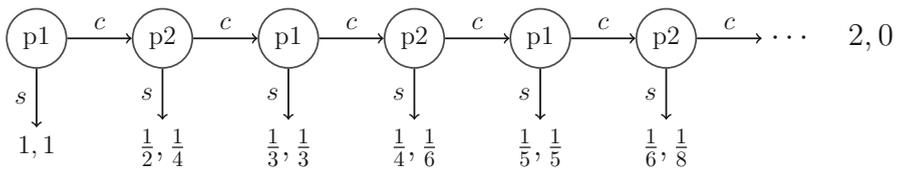
In this game, there are two players, who move in turn. The action set is $A = \{s, c\}$ where s stands for *stop* and c stands for *continue*. There are three possible scenarios. If player 1 stops the game at period $t \in \{0, 2, 4, \ldots\}$, then the payoffs are $(\frac{1}{t+1}, \frac{1}{t+1})$. If player 2 stops the game at period $t \in \{1, 3, 5, \ldots\}$, then the payoffs are $(\frac{1}{t+1}, \frac{1}{t+3})$. Finally, if neither player stops, then the payoffs are $(2, 0)$. Note that the payoff function of player 1 is usc, while the payoff function of player 2 is even continuous.

Define the tolerance level by

$$\phi(t) = \tfrac{1}{2} \min \left\{ \tfrac{1}{t+1} - \tfrac{1}{t+2}, \tfrac{1}{t+2} - \tfrac{1}{t+3} \right\}$$

where we write t to denote the history reached once c has been played t times. We argue that there is no ϕ-tolerance equilibrium. Suppose by way of contradiction that σ is a ϕ-tolerance equilibrium. We distinguish two cases.

Suppose first that there is some period t_0 after which σ_2 never quits. Since the tolerance level is always smaller than 1, σ_1 prescribes player 1 to always continue after t_0, and this yields payoff $(2, 0)$. Then, however, at any odd period $t > t_0$, player 2 could improve his payoff with $\frac{1}{t+3}$ by quitting, and this gain is greater than $\phi(t)$.

Now suppose that σ_2 quits at infinitely many odd periods, say at $t_1 < t_2 < t_3 < \cdots$ etc. Then our choice of the loss tolerance function forces player 1 to quit at $t_2 - 1$. But then it forces player 2 to continue at $t_2 - 2$. At position $t_2 - 3$ player 1 has to stop, and at position $t_2 - 4$ player 2 must continue. Working backwards we find that player 1 quits at each position $t < t_2$ and that player 2 continues at each position $t < t_2$. This contradicts our supposition that player 2 quits at t_1. This completes the argument.

5 The proof of Theorem 4.4

The proof of the theorem relies on Lemma 5.1 which warrants some attention on its own right. It is in essence an ϕ-version of a well-known one-shot deviation principle.

We say that the strategy σ is ϕ-robust to one-shot deviations if by deviating from σ at any single history h the player $\iota(h)$ cannot improve his payoff by more than $\phi(h)$. More precisely, σ is *ϕ-robust to one-shot deviations* if for every $h \in H$

$$u_{\iota(h)}(\pi(\sigma; h)) \geq u_{\iota(h)}(\pi(\sigma; ha)) - \phi(h) \quad \text{for each } a \in A. \tag{5.1}$$

Clearly each ϕ-tolerance equilibrium is ϕ-robust to one-shot deviations. The converse is not true, but we have the following result.

Lemma 5.1 *Suppose the functions u_i are bounded and lsc. Then the following conditions are equivalent:*

[1] *For each loss tolerance function ϕ there exists a ϕ-tolerance equilibrium.*
[2] *For each loss tolerance function ϕ there exists a strategy profile that is ϕ-robust to one-shot deviations.*

Proof That [1] implies [2] is immediate since each ϕ-tolerance equilibrium is ϕ-robust to one-shot deviations. We show the converse implication. Let the loss tolerance function ϕ be given. Without loss of generality we can assume that it is a decreasing function in the sense that $\phi(h) > \phi(h')$ whenever $h < h'$. Define the loss tolerance function ϕ_0 by $\phi_0(h) = 2^{-|h|-1}\phi(h)$. Let the strategy profile σ be ϕ_0-robust to one-shot deviations. We show that σ is a ϕ-tolerance equilibrium.

Let h be any history and η_i a pure strategy for player i. We want to prove inequality (4.2). Write $\eta = \sigma/\eta_i$. Let $h_0, h_1, h_2, \ldots$ be the successive enumeration of the initial segments of the play $\pi(\eta; h)$ starting from h, that is $h_0 = h$ and $h_{k+1} = h_k\eta(h_k)$ for each k. We have

$$u_i(\pi(\sigma; h_k)) \geq u_i(\pi(\sigma; h_k\eta(h_k))) - \phi_0(h_k). \tag{5.2}$$

Indeed, at histories h_k that belong to player i this inequality is an instance of (5.1) with $a = \eta(h_k)$. At histories h_k that do not belong to player i the inequality holds as we have $\eta(h_k) = \sigma(h_k)$ and so $\pi(\sigma; h_k) = \pi(\sigma; h_k\eta(h_k))$. Now we have

$$\phi_0(h_k) = 2^{-|h_k|-1}\phi(h_k) \leq 2^{-|h_k|-1}\phi(h_0) \leq 2^{-k-1}\phi(h_0), \tag{5.3}$$

where we used the fact that $|h_k| = |h_0| + k$. Combining (5.2) and (5.3) and using the fact that $h_{k+1} = h_k\eta(h_k)$ we obtain

$$u_i(\pi(\sigma; h_k)) \geq u_i(\pi(\sigma; h_{k+1})) - 2^{-k-1}\phi(h_0).$$

Unraveling this recursive relation we obtain

$$u_i(\pi(\sigma; h_0)) \geq u_i(\pi(\sigma; h_{m+1})) - \sum_{k=0}^{m} 2^{-k-1}\phi(h_0).$$

The sequence $\{\pi(\sigma; h_m)\}_{m \in \mathbb{N}}$ of plays converges to the play $\pi(\eta; h_0)$ as m approaches infinity. To see this, simply notice that h_m is the initial segment of both plays $\pi(\eta; h_0)$

and $\pi(\sigma; h_m)$. Hence passing to the limit and using the lower semicontinuity of the payoff function, we obtain

$$u_i(\pi(\sigma; h_0)) \geq \liminf_{m \to \infty} u_i(\pi(\sigma; h_{m+1})) - \sum_{k=0}^{\infty} 2^{-k-1} \phi(h_0) \geq u_i(\pi(\eta; h_0)) - \phi(h_0),$$

as desired.   □

We can now complete the proof of the theorem. We follow the approach of Flesch et al. (2010) with two crucial modifications. The first one is concerned with the discretization of the payoff function. In both cases discretizations are introduced as a player is allowed to ignore a small difference in the payoff. However, while Flesch et al. (2010) use one fixed discretization as determined by a fixed level of loss tolerance ϵ, we are forced to use different discretizations for each of the players' decision nodes h, as determined by the respective tolerance levels $\phi(h)$.

The second modification is in the argument that the procedure always returns a non-empty set of plays. By making use of the continuity of the payoff functions we are able to avoid some of the more difficult steps of the proof of Flesch et al. (2010). For more on this point see the remarks at the end of the proof.

Proof of Theorem 4.4 In view of Lemma 5.1 it is enough to show that for each loss tolerance function ϕ there exists a strategy profile that is ϕ-robust to one-shot deviations. For each $h \in H$ define the function v_h by letting

$$v_h(p) = \phi(h) \left\lceil \frac{u_{\iota(h)}(p)}{\phi(h)} \right\rceil$$

where $\lceil r \rceil$ is the largest integer that is not greater than r. Since the functions u_i are bounded, the function v_h only takes finitely many values. Moreover $\lceil r \rceil$ is usc as a function of r, and so v_h is usc as a function of p.

We define by transfinite recursion the following sequences. For ordinal 0 let

$$P_0(h) = \{p \in A^{\mathbb{N}} : h < p\}$$
$$\alpha_0(h) = \min_{p \in P_0(h)} v_h(p).$$

For each successor ordinal $\xi + 1$ let

$$\alpha_{\xi+1}(h) = \max_{a \in A} \min_{p \in P_\xi(h,a)} v_h(p) \tag{5.4}$$

$$P_{\xi+1}(h) = \left\{ p \in \bigcup_{a \in A} P_\xi(h,a) : v_h(p) \geq \alpha_{\xi+1}(h) \right\}. \tag{5.5}$$

For each limit ordinal ξ we set

$$
\begin{aligned}
P_\xi(h) &= \bigcap_{\lambda < \xi} P_\lambda(h) \\
\alpha_\xi(h) &= \min_{p \in P_\xi(h)} v_h(p).
\end{aligned}
\tag{5.6}
$$

By convention, a minimum over the empty set if $+\infty$. These definitions are precisely the same as in Flesch et al. (2010) except that we use the function v_h in place of the function $u_{\iota(h)}$. We state the following basic properties of these sequences without proof. The proof (which is straightforward) is exactly the same as the proof of Lemma 3.1, and of Properties Q2 and Q3 in Flesch et al. (2010).

Lemma 5.2 *Let ξ be an ordinal and $h \in H$. [1] Let a be an element of A that attains the maximum in (5.4). Then $P_\xi(h, a) \subseteq P_{\xi+1}(h)$. [2] If ξ is a limit ordinal or zero, $p \in P_\xi(h)$ and $h < h' < p$, then $p \in P_\xi(h')$.*

Lemma 5.3 (Monotonicity of the sequences) *Let η and λ be ordinals and $h \in H$. If $\eta < \lambda$ then $P_\eta(h) \supseteq P_\lambda(h)$ and $\alpha_\eta(h) \leq \alpha_\lambda(h)$.*

We now turn to the key step of the proof. It is here at this step that we are able to use the continuity of the payoff functions to simplify the argument.

Lemma 5.4 *The set $P_\xi(h)$ is non-empty for each ordinal ξ and each $h \in H$.*

Proof The proof is by transfinite induction on ξ. Thus fix an ordinal ξ and suppose that $P_\nu(h)$ is non-empty for each ordinal $\nu < \xi$ and each history $h \in H$. We have to prove that $P_\xi(h) \neq \varnothing$ for each $h \in H$.

If $\xi = 0$ the claim is clearly true. If ξ is a successor ordinal it follows from Lemma 5.2. Suppose ξ is a limit ordinal.

For each $h \in H$ define

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

Note that the maximum is well defined: v_h and hence also $\alpha_\lambda(h)$ only take finitely many values. Moreover since the sequence $\alpha_\lambda(h)$ is non-decreasing in λ we have $\tilde{\alpha}(h) = \alpha_\lambda(h)$ for all sufficiently large ordinals $\lambda < \xi$.

Now fix a history $h \in H$. We recursively define a sequence of histories and ordinals as follows. Set $h_0 = h$. Suppose for some $k \in \mathbb{N}$ the history h_k has been defined. Let ξ_k be a successor ordinal satisfying

$$
\alpha_{\xi_k}(h_k) = \tilde{\alpha}(h_k).
$$

If $k \geq 1$ we moreover require that $\xi_{k-1} - 1 \leq \xi_k$. An ordinal ξ_k with these properties exists by the remarks above. By (5.4) there is an $a_k \in A$ such that
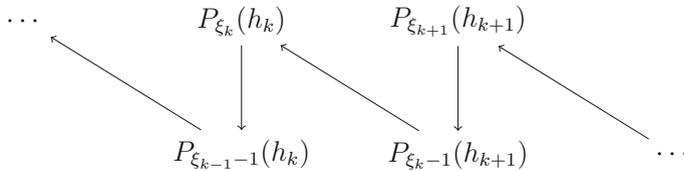
$$
\alpha_{\xi_k}(h_k) = \min_{p \in P_{\xi_k-1}(h_k, a_k)} v_{h_k}(p).
$$

Set $h_{k+1} = (h_k, a_k)$. This completes the induction step.

We have the inclusions

$$P_{\xi_k}(h_k) \supseteq P_{\xi_k-1}(h_{k+1}) \supseteq P_{\xi_{k+1}}(h_{k+1}),$$

where the first inclusion holds by the choice of the action a_k and by part [1] of Lemma 5.2, and the second inclusion holds because $\xi_k - 1 \leq \xi_{k+1}$ and by Lemma 5.3. Combining these we obtain an infinite chain of inclusions

For each $\ell \in \mathbb{N}$ let p_ℓ be an arbitrary element of the set $P_{\xi_\ell}(h_\ell)$. Let $p = (h_0, a_0, a_1, \ldots)$. Since h_ℓ is an initial segment of p_ℓ, the sequence $\{p_\ell\}_{\ell \in \mathbb{N}}$ of plays converges to the play p. If $k \leq \ell$ then $P_{\xi_k}(h_k) \supseteq P_{\xi_\ell}(h_\ell)$ and so p_ℓ is an element of $P_{\xi_k}(h_k)$. Thus we have

$$v_{h_k}(p_\ell) \geq \alpha_{\xi_k}(h_k) = \tilde{\alpha}(h_k)$$

whenever $k \leq \ell$, where the inequality follows from (5.5). Taking the limit as ℓ approaches infinity and making use of the upper semicontinuity of v_{h_k} we obtain

$$v_{h_k}(p) \geq \limsup_{\ell \to \infty} v_{h_k}(p_\ell) \geq \tilde{\alpha}(h_k).$$

Consequently for each $\lambda < \xi$ it holds that

$$v_{h_k}(p) \geq \alpha_\lambda(h_k). \tag{5.7}$$

Now let Φ_λ be the assertion that $p \in P_\lambda(h_k)$ for each $k \in \mathbb{N}$. We prove that Φ_λ holds for each $\lambda < \xi$. The assertion Φ_0 is true since h_k is an initial segment of p. Suppose Φ_ν is true for each $\nu < \lambda$. If λ is a limit ordinal then Φ_λ is true by definition (5.6). Suppose that λ is a successor ordinal. Take a $k \in \mathbb{N}$. Then $p \in P_{\lambda-1}(h_{k+1})$ by $\Phi_{\lambda-1}$. Together with (5.7) and definition (5.5) this implies that $p \in P_\lambda(h_k)$. This completes the induction step.

In particular we have shown that $p \in P_\lambda(h)$ for each $\lambda < \xi$. Hence $p \in P_\xi(h)$. □

A reader will notice that the construction above is similar to that of the odd iteration in Flesch et al. (2010). The difference is that our version of the odd iteration is infinite. The continuity of the payoff functions helps us to avoid having to switch between odd and even iterations as is done in Flesch et al. (2010). We remark that the proof of Lemma 3.3 in Flesch et al. (2010) asserting the finiteness of the odd iteration cannot be carried over to our setup, because of our use of different dicretizations at different histories.

To complete the proof one constructs a strategy profile σ. This is done just as in the proof of Theorem 2.3 in Flesch et al. (2010) using the function v_h in place of $u_{\iota(h)}$. By Lemma 5.3 there exists an ordinal ξ^* such that $P_{\xi^*}(h) = P_{\xi^*+1}(h)$ for each $h \in H$ and consequently $\alpha_{\xi^*}(h) = \alpha_{\xi^*+1}(h)$ for each $h \in H$. By Lemma 5.4 the sets $P_{\xi^*}(h)$ are non-empty for each h. We can assume without loss of generality that ξ^* is a limit ordinal. Let $p(\emptyset)$ be an arbitrary element of $P_{\xi^*}(\emptyset)$, and for each history of the form ha let $p(ha) \in P_{\xi^*}(ha)$ be a play that attains the minimum in the following problem

$$\min_{p \in P_{\xi^*}(ha)} v_h(p).$$

The strategy profile σ is defined as follows: Follow $p(\emptyset)$, as long as no one deviates from it. Suppose the first deviation from $p(\emptyset)$ occurs at history h_0 where player i_0 plays action a_0. Then switch to the play $p(h_0 a_0)$ and follow it as long as no one deviates. Suppose the first deviation from $p(h_0 a_0)$ occurs at history h_1 where player i_1 plays action a_1. Then switch to the play $p(h_1 a_1)$. And so on.

We next argue that for every history $h \in H$ and action $a \in A$

$$v_h(\pi(\sigma; h)) \geq v_h(\pi(\sigma; ha)).$$

If $\sigma(h) = a$ then $\pi(\sigma; h) = \pi(\sigma; ha)$ so there is nothing to prove. Suppose otherwise. By construction there exists a history $h' \leq h$ such that $\pi(\sigma; h) = p(h')$. Moreover, $\pi(\sigma; ha) = p(ha)$. Hence it is sufficient to show that $v_h(p(h')) \geq v_h(p(ha))$. Now $p(h') \in P_{\xi^*}(h')$. By claim [2] of Lemma 5.2 we have $p(h') \in P_{\xi^*}(h)$ and so $p(h') \in P_{\xi^*+1}(h)$. Hence

$$v_h(p(h')) \geq \alpha_{\xi^*+1}(h) \geq \min_{p \in P_{\xi^*}(ha)} v_h(p) = v_h(p(ha)),$$

as desired.

Finally we prove that σ is ϕ-robust to one-shot deviations. We have

$$u_{\iota(h)}(\pi(\sigma; h)) \geq v_h(\pi(\sigma; h)) \geq v_h(\pi(\sigma; ha)) \geq u_{\iota(h)}(\pi(\sigma; ha)) - \phi(h).$$

This completes the proof of Theorem 4.4.   □

Remark 5.5 For certain games the number of iterations in the proof of Theorem 4.4 may need to be transfinite. Indeed, let τ be the smallest ordinal for which $P_\tau(h) = P_{\tau+1}(h)$ for each $h \in H$. The following example shows that τ can be arbitrarily large.

Let ξ be any ordinal. Slightly modifying Example 4.1.1 in Flesch et al. (2010) we obtain a game where $\tau = \xi$. In this game players 1 and 2 choose a decreasing sequence $a_0 > a_1 > \cdots$ of ordinals, with $a_0 = \xi$. The rules are as follows: If the current ordinal a_t is a successor ordinal, then player 1 plays $a_{t+1} = a_t - 1$ or he quits. If the current ordinal a_t is a limit ordinal, then player 2 chooses any ordinal a_{t+1} smaller than a_t. Finally, if $a_t = 0$, then the play terminates. The payoff for player 1 equals 1 if the sequence (a_t) eventually reaches 0 and equals 0 otherwise. The payoff for player 2 equals 0 for every play. The payoffs are continuous and the proof that $\tau = \xi$ is similar to that in Flesch et al. (2010).

6 Vanishing loss

In this section we discuss a property of vanishing loss, which is closely related to the two refinements introduced earlier. We deliberately keep our discussion informal. The point here is not to prove new results but rather to emphasize the connections between continuity ϵ-SPE and loss tolerance equilibrium.

We are interested in strategy profiles such that the maximal improvement any player can gain by unilaterally deviating from the given strategy approaches zero as the play proceeds. To fix the ideas we define a loss $\lambda(\sigma; h)$ associated to a strategy profile σ at history h as the maximal improvement that a player could gain by unilaterally deviating from his strategy in the subgame that starts at h. Formally $\lambda(\sigma; h)$ is the supremum of the left-hand side of Eq. (4.2) over all strategies η_i and all players i. In particular σ is an ϵ-SPE if $\lambda(h; \sigma) \leq \epsilon$ for each h and it is a ϕ-tolerance equilibrium if $\lambda(\sigma; h) \leq \phi(h)$ for each h.

Definition 6.1 A strategy profile σ exhibits *vanishing loss along the play $p = (a_0, a_1, \dots) \in A^\mathbb{N}$* if

$$\lim_{t \to \infty} \lambda(\sigma; (a_0, \dots, a_t)) = 0.$$

A strategy profile σ has vanishing loss if σ exhibits vanishing loss along the play $\pi(\sigma; h)$ for every history $h \in H$.

The requirement of vanishing losses is intuitive: it might reflect the fact that the players become more experienced in the course of the game. It might also manifest the fact that a player thinks harder about a choice at a specific decision node once that node is reached, than he does in the beginning of the game. Finally, it might simply reflect the fact that the possibilities for deviating profitably shrink as time goes by.

Vanishing loss is a property shared by the two refinements we introduced. More precisely:

[1] Each continuity ϵ-SPE has vanishing loss.
[2] Any ϕ-tolerance equilibrium has vanishing loss provided that the tolerance level $\phi(h)$ converges to zero as the length of the history increases.

Item [2] follows immediately from the definitions. Item [1] is only slightly less trivial, and can be verified by using the fact that the induced plays are continuity points of all payoff functions. In fact [1] follows from the following somewhat more general observation:

[3] For each strategy profile σ if $\pi(\sigma; h)$ is a point of upper semicontinuity of all payoff functions, then σ exhibits vanishing loss along the path $\pi(\sigma; h)$.

Hence in games with usc payoff functions, any strategy profile σ has vanishing loss, and thus Definition 6.1 has no bite in such games. If the payoff functions are lsc rather than usc, it is no longer true that every strategy profile has vanishing loss. However, by item [1] above and Corollary 3.5, a game with bounded and lsc payoffs does have a strategy profile with vanishing loss.

Beyond games where the payoff functions are semicontinuous, the existence of strategy profiles with vanishing loss is difficult to guarantee. In the following example, for small ϵ no ϵ-SPE has vanishing loss along the induced play.

Example 6.2 In this game, there is only one player, who can choose at every period between playing down (d) or moving to the right (r). If in total he plays k times action right then his payoff is $\frac{k}{k+1}$, whereas if he moves to the right infinitely many times then his payoff is 0. This game can be schematically represented as follows, where it is understood that the payoffs corresponding to the arc are only obtained if action r is not played any more:
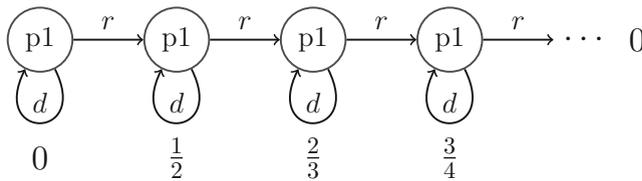
Formally $A = \{d, r\}$, and the payoff $u_1(a_0, a_1, \dots)$ is zero if the set

$$R(a_0, a_1, \dots) = \{t \in \mathbb{N} : a_t = r\}$$

is infinite, and it is $\frac{k}{k+1}$ if it has cardinality $k < \infty$. It is easy to see that u_1 is discontinuous at each play in $A^{\mathbb{N}}$ and is consequently neither lsc nor usc.

For each $\epsilon > 0$ there exists an ϵ-SPE, for example: Let k be such that $\frac{k}{k+1} > 1 - \epsilon$. Play r if in the past action r has been played less than k times, play d otherwise.

As we argue below, there is no ϵ-SPE σ, with $\epsilon < 1$, that exhibits vanishing loss along the induced play. Let $p = \pi(\sigma)$. Then $R(p)$ is finite, say consisting of k elements, so that $u_1(p) = \frac{k}{k+1}$. As the action r is played exactly k times along p, the play p is of the form $(h, d, d, d, \dots)$ for some finite history h. Now take any history $h' < p$ that is longer than h and consider the play $p' = (h', r, d, d, d, \dots)$. The play p' results in the payoff of $\frac{k+1}{k+2}$. Hence for every h' such that $h < h' < p$ it holds that

$$\lambda(\sigma; h') \geq \tfrac{k+1}{k+2} - \tfrac{k}{k+1}.$$

This provides a lower bound on the loss along the play p.

Interestingly, for $\epsilon \in (0, 1)$, each ϵ-SPE exhibits vanishing loss along the play $q = (r, r, r, \dots)$, even though q is not the induced equilibrium play.

References

Apt KR, Grädel E (eds) (2011) Lectures in game theory for computer scientists. Cambridge University Press

Alós-Ferrer C, Ritzberger K (2013) Characterizing existence of equilibrium for large extensive form games. Working paper

Binmore K, Osborne MJ, Rubinstein A (1992) Noncooperative models of bargaining. In: Aumann RJ, Hart S (eds) Handbook of game theory with economic applications. Elsevier, Amsterdam

Brihaye T, Bruyére V, De Pril J, Gimbert H (2013) On (subgame perfect) secure equilibrium in quantitative reachability games. Log Methods Comput Sci 9:1–32

Carmona G (2005) On games of perfect information: equilibria, ϵ-equilibria and approximation by simple games. Int Game Theory Rev 7:491–499

Flesch J, Kuipers J, Schoenmakers G, Vrieze K (2013) Subgame-perfection in free transition games. Eur J Oper Res 228:201–207

Flesch J, Kuipers J, Mashiah-Yaakovi A, Schoenmakers G, Solan E, Vrieze K (2010) Perfect-information games with lower semicontinuous payoffs. Math Oper Res 35:742–755

Flesch J, Kuipers J, Mashiah-Yaakovi A, Shmaya E, Schoenmakers G, Solan E, Vrieze K (2014) Non-existence of subgame-perfect ϵ-equilibrium in perfect information games with infinite horizon. Int J Game Theory 43:945–951

Fort MK (1951) Points of continuity of semicontinuous functions. Publ Math Debr 2:100–102

Fudenberg D, Levine D (1983) Subgame-perfect equilibria of finite and infinite-horizon games. J Econ Theory 31:251–268

Harris C (1985) Existence and characterization of perfect equilibrium in games of perfect information. Econometrica 53:613–628

Himmelberg CJ, Parthasarathy T, Raghavan TES, Van Vleck F (1976) Existence of p-equilibrium and optimal stationary strategies in stochastic games. Proc Am Math Soc 60:245–251

Lehrer E, Sorin S (1998) ϵ-Consistency in repeated games. Int J Game Theory 27:231–244

Laraki R, Maitra AP, Sudderth WD (2013) Two-person zero-sum stochastic games with semicontinuous payoff. Dyn Games Appl 3:162–171

Mailath GJ, Postlewaite A, Samuelson L (2005) Contemporaneuos perfect epsilon-equilibria. Games Econ Behav 53:126–140

Maitra AP, Sudderth WD (2007) Subgame-perfect equilibria for stochastic games. Math Oper Res 32:711–722

Martin DA (1975) Borel determinacy. Ann Math 2:363–370

Mashiah-Yaakovi A (2009) Periodic stopping games. Int J Game Theory 38:169–181

Mashiah-Yaakovi A (2014) Subgame perfect equilibria in stopping games. Int J Game Theory 43:89–135

Mertens J-F (1987) Repeated games. In: Proceedings of the international congress of mathematicians, (Berkeley, California, 986), 0528–0577. American Mathematical Society, Providence, RI

Purves RA, Sudderth WD (2011) Perfect information games with upper semicontinuous payoffs. Math Oper Res 36:468–473

Ray D (2007) A game-theoretic perspective on coalition formation. Oxford University Press, Oxford

Selten R (1965) Spieltheoretische Behandlung eines Oligopolmodells mit Nachfragentragheit. Zeitschrift fur die gesamte Staatswissenschaft 12:201–324

Solan E, Vieille N (2001) Quitting games. Math Oper Res 26:265–285

Solan E, Vieille N (2003) Deterministic multi-player Dynkin games. J Math Econ 39:911–929

Solan E (2005) Subgame-perfection in quitting games with perfect information and differential equations. Math Oper Res 30:51–72