



Perfect information games where each player acts only once

Kutay Cingiz¹ · János Flesch² · P. Jean-Jacques Herings³  · Arkadi Predtetchinski³

Received: 26 February 2018 / Accepted: 17 May 2019 / Published online: 28 May 2019
© The Author(s) 2019, corrected publication 2020

Abstract

We study perfect information games played by an infinite sequence of players, each acting only once in the course of the game. We introduce a class of frequency-based minority games and show that these games have no subgame perfect ϵ -equilibrium for any ϵ sufficiently small. Furthermore, we present a number of sufficient conditions to guarantee existence of subgame perfect ϵ -equilibrium.

Keywords Minority games · Subgame perfect ϵ -equilibria · Upper semicontinuous functions · Infinitely many players

JEL Classification C72 · C73 · D91

The authors would like to thank Nate Eldredge for an early version of the proof of Theorem 3.3, Natalia Chernova for suggesting the use of coupling, and Abraham Neyman and Jérôme Renault for their valuable comments. We are also grateful to the Editor and two anonymous referees for their helpful questions and suggestions.

✉ P. Jean-Jacques Herings
P.Herings@maastrichtuniversity.nl

Kutay Cingiz
kutay.cingiz@wur.nl

János Flesch
J.Flesch@maastrichtuniversity.nl

Arkadi Predtetchinski
A.Predtetchinski@maastrichtuniversity.nl

¹ Agricultural Economics and Rural Policy Group, Wageningen University, Hollandseweg 1, 6706 KN Wageningen, The Netherlands

² Department of Quantitative Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands

³ Department of Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands

1 Introduction

We study perfect information games with infinitely many players where each player is active only once. Our emphasis is on the question of existence of subgame perfect ϵ -equilibria.¹ The references below testify to the increasing attention that the concept has been receiving in the recent game-theoretic literature. In addition to the challenging existence issue, the paper reaches out to a large area of research on topics such as minority games, time-inconsistent preferences, and intergenerational games.

The contribution of the paper is twofold. The first contribution is to introduce a class of so-called frequency-based minority games. These are games where a player benefits from choosing an action that is less popular among the players. Minority games have been used to model a variety of different phenomena ranging from congestion, downloading in P2P networks, market entry, to speculative behavior in financial markets. For an extensive list of related works, we refer to Renault et al. (2007) and Renault et al. (2008).

Unlike much of the work on minority games that focuses on the setup with simultaneous moves, here we consider sequential games of perfect information. In more detail, we consider perfect information games where each player becomes active exactly once in the course of the game and has a choice between two actions. One is a risky action: It gives a high payoff if a minority of the players use it and a low payoff otherwise. The other is a safe action: Taking this action leads to an intermediate payoff for sure, irrespective of what the other players do. In minority games, payoffs depend on the set of players who take the action that the minority of the players choose. The subtlety of the construction lies in the exact meaning of the expression “a minority of the players.” We allow for several specifications, the most characteristic of which invokes the frequency of the appearance of the risky action in the population. If too many players choose the safe action or the risky action, then choosing the risky action or the safe action, respectively, provides a higher payoff.

Our main result on frequency-based minority games in Sect. 3 states that these games admit no subgame perfect ϵ -equilibria for small positive ϵ . The crucial tension arises from the fact that each player needs to predict the behavior of the future players. If, in the course of the game, it becomes likely that the majority of the future players will take the risky action, then the active player has an incentive to take the safe action with a large probability. But then, the probability for the majority of players to play risky is small.

Closely related to our work are two known counterexamples to existence of subgame perfect ϵ -equilibrium. One is a stopping game played by a sequence of players considered in Flesch, Kuipers, Mashiah-Yaakovi, Schoenmakers, Solan, and Vrieze (2010, page 753). The example shows that games of the type studied in this paper need not have subgame perfect ϵ -equilibria in pure strategies.

The other example is given in Flesch et al. (2014). This being a two-player game, it falls outside of the class of games studied here. All deviations in the current paper are automatically one-shot, since each player acts only once in the course of the

¹ For necessary and sufficient conditions for the existence of subgame perfect 0-equilibria in perfect information games, we refer the reader to Alós-Ferrer and Ritzberger (2017).

game. In contrast, in Flesch et al. (2014) deviations can be more complex and may involve changes by a player at infinitely many periods of time. Moreover, while the counterexample in the present paper admits an interpretation as an intergenerational game, the same cannot be said of the game in Flesch et al. (2014). Indeed, the agent form of the game in Flesch et al. (2014), that is, the game obtained by splitting each player into infinitely many “selves,” each self acting at most once, does have a subgame perfect 0-equilibrium in pure strategies.²

There are also important differences on a purely technical level. This paper relies on coupling. Coupling is used to compare any given strategy profile of the players to one particular stationary strategy profile whereby each player takes the risky action with a fixed positive probability. The main technique in Flesch et al. (2014) is Lévy’s zero–one law. The intuition behind the nonexistence of subgame perfect ϵ -equilibrium is that in order for player 2 to punish player 1, player 2 must, with a small but positive probability, repeatedly take a “bad” action. However, Lévy’s zero–one law then forces player 2 to take the bad action in deep subgames with probability close to one. Such behavior is incompatible with payoff maximization by player 2.

The second contribution of this paper is to provide several sufficient conditions for the existence of subgame perfect ϵ -equilibrium. Many of these conditions have appeared previously in the literature in various contexts. We now discuss each condition in turn.

[I] Games with qualitative objectives (also called Boolean objectives): These are games where each player’s payoff function takes only two values, 0 or 1. This class is motivated by a vast literature in computer science on qualitative games. The inventory of qualitative objectives considered in the computer science literature includes reachability, safety, Büchi, parity, and Müller objectives. For the precise definitions of these objectives, we refer to the reviews by Chatterjee and Henzinger (2012) and Bruyère (2017). We show that games with qualitative objectives admit subgame perfect 0-equilibria in pure strategies. The proof is based on the iterative procedure in Flesch et al. (2010).

[II] Games with upper semicontinuous payoffs. We show that games of this class admit a subgame perfect 0-equilibrium in pure strategies. This result complements the findings of Purves and Sudderth (2011), who consider games with finitely many players. A more general result, albeit with a more complicated proof, has been recently obtained by Flesch and Predtetchinski (2017).

[III] Games continuous outside a countable set. These are games in which the payoffs are continuous when restricted to a sufficiently large part of the domain. This somewhat technical condition is related to the work of Flesch and Predtetchinski (2016b). We show that this class of games admits a subgame perfect ϵ -equilibrium for every positive value of ϵ . The result generalizes the corresponding result of Cingiz et al. (2016).

[IV] Games played by a finite number of teams. A team is a group of players with identical payoff functions. In this class of games, the set of players can be partitioned into finitely many teams. The idea of teams consisting of individual players

² We are grateful to the editor for raising this question. An example of a subgame perfect 0-equilibrium of the agent form is: Each self of player 1 always takes action 2, and each self of player 2 always takes action 2.

has attracted considerable attention (see, e.g., von Stengel and Koller 1997; Solan 2000; Gossner and Tomala 2007; von Stengel and Zamir 2010; Gossner and Hörner 2010). We show that games of this class admit a subgame perfect 0-equilibrium in pure strategies.

Games where each player acts only once arise naturally in economics. One example is decision making with time-inconsistent preferences, as modeled by a sequence of multiple selves (Strotz 1956; Pollak 1968; Peleg and Yaari 1973). A special case of this is hyperbolic discounting (Phelps and Pollak 1968; Laibson 1994, 1997; Jaskiewicz and Nowak 2014). Another example relates to intergenerational games, where a player in the infinite sequence of players represents a generation (Phelps and Pollak 1968; Balbus et al. 2015). In a similar vein, Asheim (2010) considers a model with an infinite stream of generations to analyze intergenerational social welfare. A class of continuous games played by a sequence of players is studied, for example, in Hellwig et al. (1990). Callander and Hörner (2009) study social learning using games played by a sequence of players.

The paper is organized as follows. In Sect. 2, we provide the general model and formal definitions. In Sect. 3, we introduce frequency-based minority games and state our main result: Nonexistence of subgame perfect ϵ -equilibrium. In Sect. 4, we give several sufficient conditions for the existence of subgame perfect ϵ -equilibrium. Proofs are collected in “Appendix.”

2 The model

Preliminaries: Let \mathbb{N} denote the set of natural numbers $\{1, 2, \dots\}$. Let A be a non-empty finite set. For every $t \in \mathbb{N}$, let H_t denote the set of sequences of elements of A of length $t - 1$. Thus, in particular, H_1 is a singleton consisting of the empty sequence \emptyset . Let $H = \cup_{t \in \mathbb{N}} H_t$. Elements of H are called histories. Let $A^{\mathbb{N}}$ denote the set of infinite sequences of elements of A . Elements of $A^{\mathbb{N}}$ are called plays. A typical play is written as $p = (a_1, a_2, \dots)$. Finite sequences $\emptyset, (a_1), (a_1, a_2), \dots$ are said to be prefixes of the play (a_1, a_2, \dots) . A play $p \in A^{\mathbb{N}}$ is said to extend a history $h \in H$ if h is a prefix of p . In a similar way, we define the prefix of a history and a history extending another history.

Given a history $h \in H$, we define the cylinder set $C(h)$ as the set of plays $p \in A^{\mathbb{N}}$ such that h is a prefix of p . The collection of all cylinder sets is the basis of a topology on $A^{\mathbb{N}}$. In what follows, $A^{\mathbb{N}}$ is always assumed to be endowed with the topology generated by the cylinder sets. This topology is equal to the product topology, where each A is given the discrete topology, and is well known to be a Polish topology.

We let \mathcal{B} denote the Borel sigma-algebra on $A^{\mathbb{N}}$, the smallest sigma-algebra containing all the open sets. Since A is assumed to be finite, \mathcal{B} coincides with the sigma-algebra generated by the cylinder sets.

The game: We consider games where the set of players is equal to \mathbb{N} . Player $t \in \mathbb{N}$ is active in exactly one period, period t . In period t , player t chooses an action a_t from the set A . The chosen action is observed by all players. Depending on the infinite sequence of actions chosen $p = (a_1, a_2, \dots)$, player t receives a payoff equal to $u_t(p)$.

Throughout the paper, we assume that for every player $t \in \mathbb{N}$, u_t takes only finitely many values and is Borel measurable on $A^{\mathbb{N}}$.

Let some player $t \in \mathbb{N}$ be given. A strategy σ_t for player t is a mapping $\sigma_t : H_t \rightarrow \Delta(A)$. Recall that H_t is defined as the set of histories of length $t - 1$, so H_t corresponds to the set of histories where player t is active. The set $\Delta(A)$ contains all probability measures on A . A strategy profile is denoted by $\sigma = (\sigma_t)_{t \in \mathbb{N}}$. The set of strategies for player t is denoted by Σ_t and the set of strategy profiles by $\Sigma = \prod_{t \in \mathbb{N}} \Sigma_t$.

A strategy σ_t for player $t \in \mathbb{N}$ is called pure if for every $h \in H_t$, $\sigma_t(h)$ places all the probability on a single action. The set of pure strategies for player t is denoted by S_t . A pure strategy $s_t \in S_t$ is treated as a function $s_t : H_t \rightarrow A$. The set of pure strategy profiles is denoted by $S = \prod_{t \in \mathbb{N}} S_t$.

We define $\rho : S \times H \rightarrow A^{\mathbb{N}}$ by letting $\rho(s|h)$ denote the play induced by the pure strategy profile s conditional on reaching history h . Formally, for every $t \in \mathbb{N}$, for every $h \in H_t$, we let $\rho(s|h)$ be the play (h, a_t, a_{t+1}, \dots) where $a_t = s_t(h)$, $a_{t+1} = s_{t+1}(h, a_t)$, and so on. The payoff for a player $t \in \mathbb{N}$ induced by a pure strategy profile $s \in S$ at $h \in H_t$ is denoted by $v_t(s|h) = u_t(\rho(s|h))$.

Given a strategy profile $\sigma \in \Sigma$ and a history $h \in H$, we let $\mu_{\sigma,h}$ denote the Borel probability measure on $A^{\mathbb{N}}$ induced by σ in the subgame starting with history h . Thus, $\mu_{\sigma,h}$ places probability 1 on the cylinder set $C(h)$. The expected payoff for a player $t \in \mathbb{N}$, induced by σ at h , is equal to

$$v_t(\sigma|h) = \mathbb{E}_{\mu_{\sigma,h}}(u_t) = \int_{A^{\mathbb{N}}} u_t \, d\mu_{\sigma,h}. \tag{2.1}$$

We denote $v_t(\sigma|\emptyset)$ by $v_t(\sigma)$. This brings us to the following definitions.

ϵ -Equilibrium: For $\epsilon \geq 0$, a strategy profile $\sigma \in \Sigma$ is an ϵ -equilibrium if for every player $t \in \mathbb{N}$, for every strategy $\sigma'_t \in \Sigma_t$,

$$v_t(\sigma) \geq v_t(\sigma_{-t}, \sigma'_t) - \epsilon.$$

Mertens and Neyman (see Mertens 1987) have proven that a perfect information game with finitely many players and bounded Borel-measurable payoff functions admits a pure ϵ -equilibrium. Le Roux (2013) generalized their result in several directions: to infinitely many players (thus covering the games studied in this paper) and quasi-Borel preference relations rather than Borel payoff functions. Since we additionally assume that the payoff functions have finite range, the games studied in this paper admit a pure 0-equilibrium.

A strategy profile that constitutes an ϵ -equilibrium for a game does not necessarily induce an ϵ -equilibrium in every subgame, even though that would be highly desirable for a solution concept. This is the main motivation for considering the following refinement of ϵ -equilibrium.

Subgame perfect ϵ -equilibrium: For $\epsilon \geq 0$, a strategy profile $\sigma \in \Sigma$ is a subgame perfect ϵ -equilibrium if for every history $h \in H$, for every player $t \in \mathbb{N}$, for every strategy $\sigma'_t \in \Sigma_t$,

$$v_t(\sigma|h) \geq v_t(\sigma_{-t}, \sigma'_t|h) - \epsilon.$$

Thus, a strategy profile $\sigma \in \Sigma$ is a subgame perfect ϵ -equilibrium if and only if, for every $h \in H$, σ induces an ϵ -equilibrium in the subgame starting with history h if and only if for every $t \in \mathbb{N}$, for every $h \in H_t$, for every $a_t \in A$,

$$v_t(\sigma|h) \geq v_t(\sigma_{-t}, a_t|h) - \epsilon. \tag{2.2}$$

The goal of this paper is to study the existence of subgame perfect ϵ -equilibria.

3 Frequency-based minority games

This section introduces the class of frequency-based minority games. In Sect. 3.1, we describe our leading example of a frequency-based minority game, and in Sect. 3.2, we give a definition of the entire class. Section 3.3 provides some further discussion of the definition, and Sect. 3.4 contains the main result and an overview of its proof. The formal proof is relegated to ‘‘Appendix A.’’

3.1 The leading example

Example 3.1 Consider a game where the action set is given by $A = \{0, 1\}$. Define the function $f : A^{\mathbb{N}} \rightarrow [0, 1]$ by

$$f(a_1, a_2, \dots) = \limsup_{t \rightarrow \infty} \frac{1}{t}(a_1 + \dots + a_t). \tag{3.1}$$

The function f gives the frequency (interpreted as the limit superior of the empirical average) of the use of action 1. We define the set

$$F = \left\{ p \in A^{\mathbb{N}} \mid f(p) > \frac{1}{2} \right\}. \tag{3.2}$$

We interpret F as expressing the condition that ‘‘more than half’’ of the players use action 1. To illustrate, consider the play $p = (1, 1, 0, 1, 1, 0, \dots)$. Under this action sequence, only every third player takes the action 0, so $f(p) = 2/3$ and p is an element of F . Notice that p is an element of F even though 0 is played by infinitely many players.

For $t \in \mathbb{N}$, player t ’s payoff function is given by

$$u_t(a_1, a_2, \dots) = \begin{cases} 1, & \text{if } a_t = 1, \\ 2, & \text{if } a_t = 0 \text{ and } (a_1, a_2, \dots) \in F, \\ 0, & \text{if } a_t = 0 \text{ and } (a_1, a_2, \dots) \in A^{\mathbb{N}} \setminus F. \end{cases} \tag{3.3}$$

Action 1 is a safe action that leads to a payoff of 1 irrespective of the actions chosen by the other players. Action 0 is a risky action. It leads to a payoff of 2 if the play of the game belongs to the set F and to a payoff of 0 otherwise, so the risky action pays off if it is chosen by a minority of the players.

As noted earlier, the games studied in this paper have a 0-equilibrium. One particular 0-equilibrium of the game in Example 3.1 is as follows: All players take action 1 as long as no player has chosen action 0. As soon as some player takes action 0, all subsequent players take action 0 as well. This 0-equilibrium is not subgame perfect. At a history such that some player has chosen action 0, the 0-equilibrium prescribes to take action 0 leading to a payoff of 0. Switching to action 1 gives a payoff of 1 and is therefore a profitable deviation.

Surprisingly, as follows from the main result of the paper, the game admits no subgame perfect ϵ -equilibrium for any sufficiently small $\epsilon > 0$. It is clear that the payoff functions of the game are not continuous; otherwise, existence would follow, for example, from the positive result of Sect. 4.2.

3.2 The definition of frequency-based minority games

Let $A = \{0, 1\}$ and $F \subset A^{\mathbb{N}} = \{0, 1\}^{\mathbb{N}}$. A game is called *frequency-based minority game* if for every player $t \in \mathbb{N}$ the payoff function is given by equation (3.3) and the set F satisfies the following properties:

- [F1] F is a Borel subset of $A^{\mathbb{N}}$.
- [F2] F is a tail set in the following sense: If $(a_1, a_2, \dots) \in F$ and $(a'_1, a'_2, \dots) \in A^{\mathbb{N}}$ is such that $a'_n \neq a_n$ for at most finitely many $n \in \mathbb{N}$, then $(a'_1, a'_2, \dots) \in F$.
- [F3] F is closed under \leq : If $(a_1, a_2, \dots) \in F$ and $(a'_1, a'_2, \dots) \in A^{\mathbb{N}}$ is such that, for every $t \in \mathbb{N}$, $a_t \leq a'_t$, then $(a'_1, a'_2, \dots) \in F$.
- [F4] There exists a number $\nu^* \in (0, 1)$ with the following property: If β_1, β_2, \dots is a sequence of i.i.d. binary random variables, each taking the value 1 with probability ν^* and the value 0 with probability $1 - \nu^*$, then

$$\mathbb{P}((\beta_1, \beta_2, \dots) \in F) = 1.$$

- [F5] There exists a number $\nu_* \in (0, 1)$ with the following property: If $\beta'_1, \beta'_2, \dots$ is a sequence of i.i.d. binary random variables, each taking the value 1 with probability ν_* and the value 0 with probability $1 - \nu_*$, then

$$\mathbb{P}((\beta'_1, \beta'_2, \dots) \in F) = 0.$$

The leading example of a set satisfying [F1–F5] is the set F as defined in (3.2). It satisfies [F1] since f is a Borel-measurable function. It satisfies condition [F4] with for instance $\nu^* = 2/3$ and condition [F5] with $\nu_* = 1/3$ as a consequence of the strong law of large numbers.

The set F describes the plays such that action 0 is played by a minority of the players. As we remarked previously, a minority can consist of infinitely many players.

Condition [F1] is a technical measurability condition. Condition [F2] essentially means that each individual player is negligible: whether a play belongs to F or not depends only on the tail of the play. Condition [F3] says that starting from an action profile in F , if some of the players revise their action upward, then the resulting action profile is again in F . By Condition [F4], it holds that if all players choose action

1 independently with sufficiently high probability, then the play belongs to F with probability 1. Similarly, by Condition [F5], if all players choose action 1 independently with sufficiently low probability, then the play belongs to F with probability 0.

3.3 Further discussion of properties [F1–F5]

We have already seen that the set F as defined in (3.2) satisfies [F1–F5]. In fact, the choice of $1/2$ as a lower bound on the frequency of action 1 is inessential. Any number in the open interval $(0, 1)$ yields a set F satisfying [F1–F5]. Likewise, the set

$$\left\{ (a_1, a_2, \dots) \in A^{\mathbb{N}} \mid \liminf_{t \rightarrow \infty} \frac{1}{t} (a_1 + \dots + a_t) > \frac{1}{2} \right\}$$

satisfies [F1–F5]. A somewhat less obvious example of a set satisfying [F1–F5] is

$$\left\{ (a_1, a_2, \dots) \in A^{\mathbb{N}} \mid \ell(a_1, a_2, \dots) > \frac{1}{2} \right\},$$

where ℓ is a medial limit, a special type of Banach limit, as described in Mertens, Sorin and Zamir (2015, page 29). This medial limit ℓ is a positive linear functional that assigns a real number to every bounded sequence of real numbers and has various appealing properties. For example, for every $(a_1, a_2, \dots) \in A^{\mathbb{N}}$, it holds that

$$\begin{aligned} \ell(a_1, a_2, \dots) &= \ell(a_1, \frac{1}{2}(a_1 + a_2), \frac{1}{3}(a_1 + a_2 + a_3), \dots), \\ \ell(a_1, a_2, a_3, \dots) &= \ell(a_2, a_3, \dots), \\ \liminf_{t \rightarrow \infty} \frac{1}{t} (a_1 + \dots + a_t) &\leq \ell(a_1, a_2, \dots) \leq \limsup_{t \rightarrow \infty} \frac{1}{t} (a_1 + \dots + a_t). \end{aligned}$$

It is easy to make use of examples of sets with [F1–F5] to create new examples, since the collection of those sets is closed under finite unions and finite intersections.

It follows from the topological zero–one law (Kechris (1995), Theorem 8.47) that a set satisfying [F1–F5] is either meager or comeager. The Hewitt–Savage zero–one law (Shiryayev 1996, page 382) implies that if F satisfies [F1–F5] and β_1, β_2, \dots is an i.i.d. sequence of Bernoulli random variables, then the probability $\mathbb{P}((\beta_1, \beta_2, \dots) \in F)$ is either 0 or 1.

Next, we remark that any set satisfying [F1–F5] is sandwiched between the sets E and E' defined as follows:

$$\begin{aligned} E &= \{ (a_1, a_2, \dots) \in A^{\mathbb{N}} \mid \{t \in \mathbb{N} : a_t = 0\} \text{ is finite} \}, \\ E' &= \{ (a_1, a_2, \dots) \in A^{\mathbb{N}} \mid \{t \in \mathbb{N} : a_t = 1\} \text{ is infinite} \}. \end{aligned} \tag{3.4}$$

Claim 3.2 *If F satisfies [F1–F5], then $E \subset F \subset E'$.*

Proof The set F is non-empty by [F4]. By [F3], it must contain the play $(1, 1, \dots)$. It now follows by [F2] that it contains all plays in E .

Suppose that F is not a subset of E' . Then, F contains a play (a_1, a_2, \dots) such that only finitely many players choose action 1. Replacing the finitely many 1’s by 0’s

and making use of [F2], we conclude that F contains the play $(0, 0, \dots)$. But then, by [F3], we have $F = A^{\mathbb{N}}$, contradicting [F5]. \square

The set E satisfies [F1], [F2], [F3], and [F5], but not [F4]. Indeed, if β_1, β_2, \dots is a sequence of i.i.d. distributed random variables, each taking value 1 with probability $\nu > 0$, then with probability 1 the realization of infinitely many random variables β_t is equal to 0. The set E' satisfies [F1], [F2], [F3] [F4], but not [F5].

3.4 The main result and overview of the proof

The main result of this section is the following theorem.

Theorem 3.3 *Let G be a frequency-based minority game with a set F satisfying [F1–F5]. Let ν^* be as in [F4] and ν_* as in [F5]. Then, G has no subgame perfect ϵ -equilibrium whenever $\epsilon < \min\{\nu_*, (1 - \nu^*)/3\}$.*

Proof See “Appendix A.” \square

It is easy to see that G has no subgame perfect 0-equilibrium in pure strategies. Indeed, suppose that σ is a pure subgame perfect 0-equilibrium of G . We argue that, after each history, the play induced by σ lies in F . This is because each player can guarantee a payoff of 1 by taking the safe action 1. If the induced play would not be in F , then the players taking action 0 would be receiving a payoff of 0, contradicting the fact that σ is subgame perfect 0-equilibrium. But then, since the continuation play after any history is an element of F , every player prefers to choose action 0, leading to a contradiction.

The difficulty of Theorem 3.3 involves the incorporation of mixed strategies. The intuitive reason why mixing does not mend the nonexistence result can be explained as follows: Even if the strategy profile σ is mixed, in very deep subgames the probability of reaching the set F under σ is either very close to zero or very close to one as a consequence of Lévy’s zero–one law. Thus, in those deep subgames, the players can make very accurate predictions. Such accuracy, however, forces them to make nearly pure choices, leading us back to the situation of a pure strategy profile.

The formal proof relies on a technique known as coupling. Coupling is a technique from probability theory, which is useful whenever the distributions of two distinct random variables need to be compared to each other. The technique consists of creating a new, larger probability space and two new random variables, having the same distributions as the two original ones, with some additional property that affords a simpler comparison of the two. Coupling is used to compare any given strategy profile of the players to one particular stationary strategy profile whereby each player takes the risky action with a fixed positive probability.

4 Sufficient conditions for existence

For games with finitely many players, a number of sufficient conditions for the existence of a subgame perfect ϵ -equilibrium have been identified, see for instance Kuipers

et al. (2016) and Flesch and Predtetchinski (2016a, b). Much less is known about games with infinitely many players.

Below, we discuss four classes of games for which subgame perfect ϵ -equilibria exist: games with qualitative objectives, games with upper semicontinuous payoffs, games continuous outside a countable set, and games played by a finite number of teams. Many of the conditions that we discuss have previously appeared in the literature in various contexts.

As in Sect. 2, we continue to assume that for every player for every player $t \in \mathbb{N}$, u_t takes only finitely many values and is Borel measurable on $A^{\mathbb{N}}$.

4.1 Games with qualitative objectives

The game G is said to have *qualitative objectives* if for each player $t \in \mathbb{N}$ the payoff function u_t takes at most two values, 0 and 1. The literature on computer science provides many examples of qualitative objectives. Examples are the objectives of reachability and safety, and the criteria of Büchi, parity, and Müller, see the reviews by Chatterjee and Henzinger (2012) and Bruyère (2017).

Example 4.1, which is inspired by the work of Voorneveld (2010) on minority games, is a game with qualitative objectives. The payoff functions in this game are exactly those of Example 4.3 in Voorneveld (2010), an example which is used to show the nonexistence of a 0-equilibrium. What distinguishes our analysis is that we are looking at a perfect information version of the game, while Voorneveld (2010) analyzes the simultaneous-move version. Notice that the game in Example 4.1 is not a frequency-based minority game as defined in Sect. 3.

Example 4.1 The action set is equal to $A = \{0, 1\}$. For every $t \in \mathbb{N}$, player t 's payoff function is given by

$$u_t(a_1, a_2, \dots) = \begin{cases} 1, & \text{if } a_t = 1 \text{ and } f(a_1, a_2, \dots) \leq 1/2, \\ 1, & \text{if } a_t = 0 \text{ and } f(a_1, a_2, \dots) > 1/2, \\ 0, & \text{otherwise,} \end{cases}$$

where f is as in (3.1).

Player t receives a payoff of 1 if he chooses the minority action. Otherwise, player t receives a payoff of 0. More precisely, player t receives a payoff of 1 if he chooses action 1 and the frequency of players choosing action 1 is less than or equal to $1/2$. Player t also receives a payoff of 1 if he chooses action 0, and the frequency of players choosing action 1 is strictly above $1/2$.

The game has a pure strategy subgame perfect 0-equilibrium: Player 1 playing action 1, and each player $t > 1$ taking the same action as player $t - 1$, is an example of such an equilibrium. \square

Theorem 4.2 *If the game G has qualitative objectives, then it has a subgame perfect 0-equilibrium in pure strategies.*

We apply the iterative algorithm developed in Flesch et al. (2010) to prove Theorem 4.2. For each ordinal number ξ , we recursively define the collection $\{P_\xi(h) \mid h \in$

H of sets of plays and the collection $\{\alpha_\xi(h) \mid h \in H\}$ of real numbers augmented with $+\infty$. For ordinal number 0, for every $t \in \mathbb{N}$, for every history $h \in H_t$, we define

$$P_0(h) = \{p \in A^{\mathbb{N}} \mid h \text{ is a prefix of } p\}, \tag{4.1}$$

$$\alpha_0(h) = \min_{p \in P_0(h)} u_t(p). \tag{4.2}$$

For every successor ordinal $\xi + 1$, for every $t \in \mathbb{N}$, for every $h \in H_t$, we let

$$\alpha_{\xi+1}(h) = \max_{a \in A} \min_{p \in P_\xi(h,a)} u_t(p), \tag{4.3}$$

$$P_{\xi+1}(h) = \{p \in \bigcup_{a \in A} P_\xi(h, a) \mid u_t(p) \geq \alpha_{\xi+1}(h)\}, \tag{4.4}$$

with the convention that the minimum over the empty set is equal to $+\infty$. For every limit ordinal ξ , for every $t \in \mathbb{N}$, for every history $h \in H_t$, we let

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

$$\alpha_\xi(h) = \min_{p \in P_\xi(h)} u_t(p).$$

For $t \in \mathbb{N}$, we let W_t denote the set of plays $p \in A^{\mathbb{N}}$ for which $u_t(p) = 1$. We think of W_t as the winning set of player t .

Proof of Theorem 4.2: One can show, exactly as in Flesch et al. (2010), that for each $h \in H$ the sequence $P_\xi(h)$ is non-increasing by inclusion and that the sequence $\alpha_\xi(h)$ is non-decreasing. Moreover, the game admits a subgame perfect 0-equilibrium in pure strategies if and only if the sets $P_\xi(h)$ are non-empty for every ordinal ξ and every history h . Furthermore, the set $\bigcap_\xi P_\xi(\emptyset)$ is exactly the set of plays that could be induced by subgame perfect 0-equilibria of the game.

We show that $P_\xi(h)$ is a non-empty set for each history h by induction on ξ . The statement is clearly true for $\xi = 0$. Suppose the statement is true for the ordinal ξ . Consider a history $h \in H_t$, and let a be an action that reaches the maximum in (4.3). Then, $P_{\xi+1}(h) \supset P_\xi(h, a)$; hence, $P_{\xi+1}(h)$ is non-empty, as desired.

We now turn to the more challenging case of a limit ordinal ξ . Suppose that $P_\eta(h)$ is non-empty for each $h \in H$ and each ordinal $\eta < \xi$. Let $h \in H_t$ be given. We recursively define a play $p = (h, a_t, a_{t+1}, \dots)$. To define a_t , we distinguish two cases.

Case 1: If $\alpha_\eta(h) = 0$ for each $\eta < \xi$, let a_t be any element of A .

Case 2: Otherwise, let $\eta < \xi$ be the least successor ordinal such that $\alpha_\eta(h) = 1$. By (4.3), there is an action $a_t \in A$ such that

$$\min_{p \in P_{\eta-1}(h, a_t)} u_t(p) = 1.$$

Notice that $P_\eta(h, a_t) \subset W_t$.

To proceed with the definition of p , let $h_{t+1} = (h, a_t)$ and repeat the argument ad infinitum. The construction guarantees that $p \in P_\xi(h)$. \square

Grädel and Ummels (2008) consider games with qualitative objectives in a setup that is complementary to ours: There are finitely many players and each player can move infinitely many times. The authors show that the games of this class admit a subgame perfect 0-equilibrium. Unlike Grädel and Ummels (2008), who rely on Borel determinacy (Martin, 1975) to obtain their result, our technique does not require that the winning sets W_t be Borel measurable.

4.2 Games with upper semicontinuous payoffs

The payoff function u_t of player $t \in \mathbb{N}$ is *upper semicontinuous* if for every sequence $(p^n)_{n \in \mathbb{N}}$ of plays in $A^\mathbb{N}$ that converges to a limit $p \in A^\mathbb{N}$ it holds that $\limsup_{n \rightarrow \infty} u_t(p^n) \leq u_t(p)$. Equivalently, u_t is upper semicontinuous if for each real number $r \in \mathbb{R}$ the set $\{p \in A^\mathbb{N} \mid u_t(p) \geq r\}$ is closed. The game G is said to have upper semicontinuous payoffs if every player $t \in \mathbb{N}$ has an upper semicontinuous payoff function.

Theorem 4.3 *If the game G has upper semicontinuous payoffs, then it has a subgame perfect 0-equilibrium in pure strategies.*

Proof As in the proof of Theorem 4.2, the result follows once we show that the set $P_\xi(h)$ as defined in Sect. 4.1 is a non-empty compact set for each history h and each ordinal ξ . The proof is by induction on ξ . For $\xi = 0$, the statement is clearly true. Assume that for some ordinal number ξ the statement is true. Take $h \in H_t$, and let $a \in A$ be an action that attains the maximum in (4.3). Since $P_{\xi+1}(h) \supset P_\xi(h, a)$, the set $P_{\xi+1}(h)$ is non-empty. Since, for every $a \in A$, the set $P_\xi(h, a)$ is compact by assumption and since A is finite, the set $\cup_{a \in A} P_\xi(h, a)$ is also compact. The set $\{p \in A^\mathbb{N} \mid u_t(p) \geq \alpha_{\xi+1}(h)\}$ is compact because u_t is upper semicontinuous. We conclude that $P_{\xi+1}(h)$ is compact.

Let ξ be a limit ordinal, and assume the statement is true for each ordinal $\eta < \xi$. Then, the statement is true for ξ , because the intersection of a nested family of non-empty and compact sets is non-empty and compact. \square

This result complements the result in Purves and Sudderth (2011), which shows that perfect information games with a finite number of players and upper semicontinuous payoff functions with finite range admit a subgame perfect 0-equilibrium in pure strategies. Their method of proof relies on induction on the number of payoffs in the game and therefore seems difficult to extend to games with infinitely many players.

Theorem 4.3 is subsumed by a recent result of Flesch and Predtetchinski (2017), which, however, has a more complicated proof. Another related work is Le Roux and Pauly (2014), Theorem 22, which implies that games of our class with upper semicontinuous payoffs admit a 0-equilibrium in pure strategies.

An interesting open problem is existence of subgame perfect ϵ -equilibria in games with lower semicontinuous payoffs. Due to an example in Flesch et al. (2010), we do know that such games do not always admit a subgame perfect 0-equilibrium.

4.3 Games continuous outside a countable set

A subset C of $A^{\mathbb{N}}$ is said to be co-countable if the set $A^{\mathbb{N}} \setminus C$ is countable. The game G is said to be *continuous outside a countable set* if there exists a co-countable subset C of $A^{\mathbb{N}}$ such that for each player $t \in \mathbb{N}$ the restriction of the payoff function u_t to the set C , denoted by $u_t|_C$, is a continuous function.³ The game G is said to have uniformly bounded payoffs if there exists $B < \infty$ such that, for every $t \in \mathbb{N}$, for every $p \in A^{\mathbb{N}}$, $|u_t(p)| < B$.

An illustration of this condition is provided by Example 4.4. The payoff functions of this example are borrowed from Peleg (1969). In Peleg (1969), the players make simultaneous moves and it is shown that the game does not admit a 0-equilibrium. Moreover, Radner (1980) has shown that Peleg’s game does not admit an ϵ -equilibrium for arbitrarily small positive values of ϵ . The game of Example 4.4 is, like all games studied in this paper, a game of perfect information.

Example 4.4 The action set is equal to $A = \{0, 1\}$. For every $t \in \mathbb{N}$, player t ’s payoff function is given by⁴

$$u_t(a_1, a_2, \dots) = \begin{cases} 1, & \text{if } a_t = 1, \\ 2, & \text{if } a_t = 0 \text{ and } (a_1, a_2, \dots) \in E, \\ 0, & \text{if } a_t = 0 \text{ and } (a_1, a_2, \dots) \in A^{\mathbb{N}} \setminus E. \end{cases}$$

Recall that $E = \{(a_1, a_2, \dots) \in A^{\mathbb{N}} \mid \{t \in \mathbb{N} : a_t = 0\} \text{ is finite}\}$ has been defined in (3.4) as the set of plays such that only finitely many players take action 0.

The game is not a frequency-based minority game, since the set E violates condition [F4]; see also the discussion in Sect. 3.3. The game is continuous outside a countable set, with $C = A^{\mathbb{N}} \setminus E$. Indeed, E is countable. Moreover, the function $u_t|_C$ is a function of a_t only and is therefore continuous.

The game admits a subgame perfect ϵ -equilibrium for every positive ϵ , for instance the strategy profile where each player takes action 1 with probability $1 - \epsilon$ and action 0 with probability ϵ . This result is tight, as the game has no subgame perfect 0-equilibrium as is shown in “Appendix B.” □

Theorem 4.5 *If the game G has uniformly bounded payoffs and is continuous outside a countable set, then, for every $\epsilon > 0$, it has a subgame perfect ϵ -equilibrium.*

Proof Let $\delta > 0$ be such that $2BM\delta < \epsilon$, where M is the cardinality of the action set A . For each player $t \in \mathbb{N}$, consider a restricted strategy space Σ_t^δ consisting of strategies σ_t such that for each $h \in H_t$ the probability distribution $\sigma_t(h)$ places a probability of at least δ on each action in A . Let Σ^δ be the corresponding set of strategy profiles. It is clear that for each σ in Σ^δ the induced probability measure $\mu_{\sigma,h}$ places probability zero on each

³ The condition that a game is continuous outside a countable set is weaker than the condition used in Flesch and Predtetchinski (2016b).

⁴ We have added one to the payoffs in Peleg (1969) and have relabeled action 0 as action 1 and action 1 as action 0. Clearly, this is inconsequential for the analysis of the example.

singleton set $\{p\}$ and consequently also on each countable subset of $A^{\mathbb{N}}$. In particular, $\mu_{\sigma,h}$ assigns probability 0 to the complement of the set C . It follows that the payoff function $v_i(\sigma|h)$ is continuous on Σ^δ . By Peleg and Yaari (1973, Theorem 6.2), there is a 0-equilibrium in the restricted strategy space. Since each finite sequence of actions is reached with positive probability, this is also a subgame perfect 0-equilibrium of the restricted game. It follows by the choice of δ that this strategy profile is a subgame perfect ϵ -equilibrium of the original game. \square

Theorem 4.5 generalizes a result in Cingiz et al. (2016), who consider so-called centipede games played by an infinite sequence of players.

4.4 Games played by a finite number of teams

A team is a group of players with identical payoff functions. The game G is said to be played by a finite number of teams if there exists a finite partition $\{T_1, \dots, T_n\}$ of the player set \mathbb{N} such that, for every $k \in \{1, \dots, n\}$, for every $t, t' \in T_k$, $u_t = u_{t'}$.

Theorem 4.6 *If the game G is played by a finite number of teams, then it has a subgame perfect 0-equilibrium in pure strategies.*

Proof Consider a game where the set of players is $\{1, \dots, n\}$, player k 's payoff function is u_t for $t \in T_k$, and player k makes a move at each period $t \in T_k$. This results in a game with finitely many players. Flesch et al. (2010) now apply to find a strategy profile that is immune to one-shot deviations, i.e., a strategy profile with the property that no player can improve his payoff at any history by deviating only once. Such a strategy profile induces a subgame perfect 0-equilibrium of the game G . \square

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

Appendix A: The proof of Theorem 3.3

To prove the main result, we adopt the following notational convention. Given a player $t \in \mathbb{N}$, a history $h \in H_t$, and a strategy profile σ , we use $\sigma_t(h)$ to denote the probability that $\sigma_t(h)$ places on action 1. The probability that the resulting play belongs to F is denoted by $\tau(\sigma|h)$, so

$$\tau(\sigma|h) = \mu_{\sigma,h}(F).$$

We employ coupling twice, in Steps 1 and 3, but it is only essential in Step 3. The use of coupling in Step 1 could be replaced by an appeal to the much more general

Table 1 Value of $\lambda_t(a_1, b_1, \dots, a_{t-1}, b_{t-1})(a_t, b_t)$, where $h = (a_1, \dots, a_{t-1})$.

	$b_t = 0$	$b_t = 1$
$a_t = 0$	$1 - \sigma_t(h)$	0
$a_t = 1$	$\sigma_t(h) - v^*$	v^*

Theorem 5.8 in Lindvall (1992). We choose to use coupling in Step 1 because it is very similar to the construction employed in Step 3.

STEP 1: Let $\sigma \in \Sigma$ be a strategy profile. If, for every $t \in \mathbb{N}$, for every $h \in H_t$, $\sigma_t(h) \geq v^*$, then $\tau(\sigma|\emptyset) = 1$. If, for every $t \in \mathbb{N}$, for every $h \in H_t$, $\sigma_t(h) \leq v_*$, then $\tau(\sigma|\emptyset) = 0$.

Proof of Step 1: We prove the first statement. The proof of the second statement is similar.

We first construct a probability space $(\Omega, \mathcal{S}, \mathbb{P})$ and, for every $t \in \mathbb{N}$, random variables $\alpha_t : \Omega \rightarrow \{0, 1\}$ and $\beta_t : \Omega \rightarrow \{0, 1\}$ satisfying the following three properties:

1. For every $t \in \mathbb{N}$, for every $(a_1, \dots, a_{t-1}) \in H_t$,

$$\mathbb{P}(\alpha_t = 1 | \alpha_1 = a_1, \dots, \alpha_{t-1} = a_{t-1}) = \sigma_t(a_1, \dots, a_{t-1}).$$

2. The random variables β_1, β_2, \dots are i.i.d. with, for every $t \in \mathbb{N}$, $\mathbb{P}(\beta_t = 1) = v^*$.
3. For every $t \in \mathbb{N}$, $\alpha_t \geq \beta_t$.

Let $B = \{0, 1\}$ and consider a game as in Sect. 2 where the action set is equal to $A \times B = \{0, 1\} \times \{0, 1\}$. We define the strategy profile $\lambda : \prod_{t \in \mathbb{N}} (A \times B)^{t-1} \rightarrow \Delta(A \times B)$ as follows. For every player $t \in \mathbb{N}$, for every history $(a_1, b_1, \dots, a_{t-1}, b_{t-1}) \in (A \times B)^{t-1}$, for every action $(a_t, b_t) \in A \times B$, the value of $\lambda_t(a_1, b_1, \dots, a_{t-1}, b_{t-1})(a_t, b_t)$ is as given in Table 1.

Let Ω be the set $(A \times B)^{\mathbb{N}}$, \mathcal{S} the sigma-algebra of Borel sets of Ω , and \mathbb{P} the probability measure $\mu_{\lambda, \emptyset}$. For every $t \in \mathbb{N}$, let α_t and β_t be the random variables defined by $\alpha_t(a_1, b_1, a_2, b_2, \dots) = a_t$ and $\beta_t(a_1, b_1, a_2, b_2, \dots) = b_t$. It is easy to check that α_t and β_t have the three desired properties.

The sequence $(\alpha_1, \alpha_2, \dots)$ induces the probability measure $\mu_{\sigma, \emptyset}$ on the set of plays $A^{\mathbb{N}}$. It follows that

$$1 = \mathbb{P}((\beta_1, \beta_2, \dots) \in F) \leq \mathbb{P}((\alpha_1, \alpha_2, \dots) \in F) = \mu_{\sigma, \emptyset}(F) = \tau(\sigma|\emptyset),$$

where the first equality follows from [F4] and the inequality from [F3]. □

Now choose $\epsilon > 0$ such that $\epsilon < v_*$ and $1 - 3\epsilon > v^*$.

Suppose $\sigma \in \Sigma$ is a subgame perfect ϵ -equilibrium of G .

STEP 2: For every $t \in \mathbb{N}$, for every $h \in H_t$ such that $\tau(\sigma|h) < \epsilon$, it holds that $\sigma_t(h) \geq 1 - 3\epsilon$.

Proof of Step 2: Let some $t \in \mathbb{N}$ and some $h \in H_t$ be given. It holds that

$$\tau(\sigma|h) = \sigma_t(h) \cdot \tau(\sigma|h, 1) + (1 - \sigma_t(h)) \cdot \tau(\sigma|h, 0)$$

or, equivalently,

$$(1 - \sigma_t(h)) \cdot \tau(\sigma|h, 0) = \tau(\sigma|h) - \sigma_t(h) \cdot \tau(\sigma|h, 1). \tag{4.5}$$

Now player t 's expected payoff at history h is

$$v_t(\sigma|h) = \sigma_t(h) + 2 \cdot (1 - \sigma_t(h)) \cdot \tau(\sigma|h, 0). \tag{4.6}$$

Plugging Equation (4.5) into Equation (4.6), we get

$$v_t(\sigma|h) = \sigma_t(h) + 2 \cdot (\tau(\sigma|h) - \sigma_t(h) \cdot \tau(\sigma|h, 1)).$$

Since playing action 1 yields a payoff of 1 and σ is a subgame perfect ϵ -equilibrium, we have $v_t(\sigma|h) \geq 1 - \epsilon$. It follows that

$$1 - \epsilon \leq v_t(\sigma|h) = \sigma_t(h) + 2 \cdot (\tau(\sigma|h) - \sigma_t(h) \cdot \tau(\sigma|h, 1)) \leq \sigma_t(h) + 2 \cdot \tau(\sigma|h).$$

This implies the statement of Step 2. □

STEP 3: Definition of appropriate random variables $\alpha_1, \alpha_2, \dots$ and β_1, β_2, \dots

We construct a probability space $(\Omega, \mathcal{S}, \mathbb{P})$ and, for every $t \in \mathbb{N}$, random variables $\alpha_t : \Omega \rightarrow \{0, 1\}$ and $\beta_t : \Omega \rightarrow \{0, 1\}$ satisfying the following three properties:

P1. For every $t \in \mathbb{N}$, for every $(a_1, \dots, a_{t-1}) \in H_t$,

$$\mathbb{P}(\alpha_t = 1 | \alpha_1 = a_1, \dots, \alpha_{t-1} = a_{t-1}) = \sigma_t(a_1, \dots, a_{t-1}).$$

P2. The random variables β_1, β_2, \dots are i.i.d. with, for every $t \in \mathbb{N}$, $\mathbb{P}(\beta_t = 1) = 1 - 3\epsilon$.

P3. For every $t \in \mathbb{N}$, for every $(a_1, \dots, a_{t-1}) \in H_t$, $\tau(\sigma|a_1, \dots, a_{t-1}) < \epsilon$ implies $\alpha_t \geq \beta_t$.

Let $B = \{0, 1\}$ and consider a game as in Sect. 2 where the action set is equal to $A \times B = \{0, 1\} \times \{0, 1\}$. We define the strategy profile $\lambda : \prod_{t \in \mathbb{N}} (A \times B)^{t-1} \rightarrow \Delta(A \times B)$ as follows. For every player $t \in \mathbb{N}$, for every history $(a_1, b_1, \dots, a_{t-1}, b_{t-1}) \in (A \times B)^{t-1}$, for every action $(a_t, b_t) \in A \times B$, the value of $\lambda_t(a_1, b_1, \dots, a_{t-1}, b_{t-1})(a_t, b_t)$ is as given in Table 2.

Notice that the probability that $(a_t, b_t) = (1, 0)$ when $\tau(\sigma|h) < \epsilon$ is equal to $\sigma_t(h) - (1 - 3\epsilon)$, which is nonnegative by Step 2.

Let Ω be the set $(A \times B)^{\mathbb{N}}$, \mathcal{S} the sigma-algebra of Borel sets of Ω , and \mathbb{P} the probability measure $\mu_{\lambda, \phi}$. For every $t \in \mathbb{N}$, let α_t and β_t be the random variables defined by $\alpha_t(a_1, b_1, a_2, b_2, \dots) = a_t$ and $\beta_t(a_1, b_1, a_2, b_2, \dots) = b_t$. It is easy to

check that α_t and β_t have the three desired properties. This completes the definition of the sequences.

Property P1, the fact that $1 - 3\epsilon > v^*$, and Step 1 imply that

$$\mathbb{P}((\beta_1, \beta_2, \dots) \in F) = 1. \tag{4.7}$$

Furthermore, the sequence $(\alpha_1, \alpha_2, \dots)$ induces the probability measure $\mu_{\sigma, \theta}$ on the set of plays $A^{\mathbb{N}}$. It follows that

$$\mathbb{P}((\alpha_1, \alpha_2, \dots) \in F) = \mu_{\sigma, \theta}(F) = \tau(\sigma|\theta).$$

□

STEP 4: $\tau(\sigma|\theta) = 1$.

Proof of Step 4: For every $t \in \mathbb{N}$, define the random variable $\psi_t : \Omega \rightarrow \{0, 1\}$ by letting $\psi_t(a_1, b_1, a_2, b_2, \dots) = 1$ if $\tau(\sigma|a_1, \dots, a_{t-1}) < \epsilon$ and $\psi_t(a_1, b_1, a_2, b_2, \dots) = 0$, otherwise. Notice that $\psi_t = 1$ implies $\beta_t \leq \alpha_t$. Define the random variable $\psi : \Omega \rightarrow A$ by $\psi(a_1, b_1, a_2, b_2, \dots) = 1$ if $(a_1, a_2, \dots) \in F^c$ and $\psi(a_1, b_1, a_2, b_2, \dots) = 0$ if $(a_1, a_2, \dots) \in F$. Notice that by definition $\mathbb{P}((\alpha_1, \alpha_2, \dots) \in F^c) = \mathbb{P}(\psi = 1)$.

We argue that ψ_t converges \mathbb{P} -almost surely to ψ . To see this, let $1_F : A^{\mathbb{N}} \rightarrow \{0, 1\}$ be the indicator function of the set F and let C be the set of plays $(a_1, a_2, \dots) \in A^{\mathbb{N}}$ such that $\tau(\sigma|a_1, \dots, a_t)$ converges to $1_F(a_1, a_2, \dots)$ as t approaches infinity. Applying Lévy’s zero–one law to the set F , we conclude that $\mu_{\sigma, \theta}(C) = 1$.⁵ Now we have

$$\mathbb{P}(\psi_t \rightarrow \psi) \geq \mathbb{P}((\alpha_1, \alpha_2, \dots) \in C) = \mu_{\sigma, \theta}(C) = 1,$$

where the first equality follows because the sequence $(\alpha_1, \alpha_2, \dots)$ induces the probability measure $\mu_{\sigma, \theta}$ on $A^{\mathbb{N}}$. The inequality holds by set inclusion. Take a sequence $(a_1, b_1, a_2, b_2, \dots)$ such that $(a_1, a_2, \dots) \in C$. We distinguish two cases. First, suppose that $(a_1, a_2, \dots) \in F$. Then $\psi(a_1, b_1, a_2, b_2, \dots) = 0$. Thus, by definition of τ , as t goes to infinity, $\tau(\sigma|a_1, a_2, \dots, a_t)$ converges to 1 and $\psi_t(a_1, b_1, a_2, b_2, \dots)$ converges to 0. Hence, as t goes to infinity, ψ_t converges to ψ . Now suppose that

⁵ See “Appendix C” for a precise statement of Lévy’s zero–one law.

Table 2 Value of $\lambda_t(a_1, b_1, \dots, a_{t-1}, b_{t-1})(a_t, b_t)$, where $h = (a_1, \dots, a_{t-1})$

	$\tau(\sigma h) \geq \epsilon$	
	$b_t = 0$	$b_t = 1$
$a_t = 0$	$3\epsilon \cdot (1 - \sigma_t(h))$	$(1 - 3\epsilon) \cdot (1 - \sigma_t(h))$
$a_t = 1$	$3\epsilon \cdot \sigma_t(h)$	$(1 - 3\epsilon) \cdot \sigma_t(h)$
	$\tau(\sigma h) < \epsilon$	
	$b_t = 0$	$b_t = 1$
$a_t = 0$	$1 - \sigma_t(h)$	0
$a_t = 1$	$\sigma_t(h) - (1 - 3\epsilon)$	$1 - 3\epsilon$

$(a_1, a_2, \dots) \notin F$. Then, $\psi(a_1, b_1, a_2, b_2, \dots) = 1$. Thus, as t goes to infinity, $\tau(\sigma|a_1, a_2, \dots, a_t)$ converges to 0, $\psi_t(a_1, b_1, a_2, b_2, \dots)$ converges to 1, and ψ_t converges to ψ .

Combining these facts, we obtain the following chain of equalities and inequalities:

$$\begin{aligned} &\mathbb{P}((\alpha_1, \alpha_2, \dots) \in F^c) \\ &= \mathbb{P}(\psi = 1) \end{aligned} \tag{4.8}$$

$$= \mathbb{P}((\beta_1, \beta_2, \dots) \in F \text{ and } \psi = 1) \tag{4.9}$$

$$= \mathbb{P}(\exists m \in \mathbb{N} \text{ s.t. } (\beta_1, \beta_2, \dots) \in F \text{ and } \psi_m = \psi_{m+1} = \dots = 1) \tag{4.10}$$

$$\leq \mathbb{P}(\exists m \in \mathbb{N} \text{ s.t. } (\alpha_1, \alpha_2, \dots) \in F \text{ and } \psi_m = \psi_{m+1} = \dots = 1) \tag{4.11}$$

$$= \mathbb{P}((\alpha_1, \alpha_2, \dots) \in F \text{ and } \psi = 1) \tag{4.12}$$

$$\begin{aligned} &= \mathbb{P}(\psi = 0 \text{ and } \psi = 1) \\ &= 0, \end{aligned} \tag{4.13}$$

where (4.8) holds by definition of ψ , (4.9) is true because of (4.7), and equation (4.10) is true because ψ_t converges \mathbb{P} -almost surely to ψ . To prove inequality (4.11), suppose that $(\beta_1, \beta_2, \dots) \in F$ and $\psi_m = \psi_{m+1} = \dots = 1$. By construction, we have that $\beta_t \leq \alpha_t$ whenever $\psi_t = 1$. Since F is closed under \leq by [F3] we conclude that the vector $(\beta_1, \dots, \beta_{m-1}, \alpha_m, \alpha_{m+1}, \dots)$ is in F . Now using the tail property [F2], we can replace the finite prefix $(\beta_1, \dots, \beta_{m-1})$ by $(\alpha_1, \dots, \alpha_{m-1})$ and conclude that $(\alpha_1, \alpha_2, \dots)$ is in F . Inequality (4.11) follows. Finally, (4.12) holds because ψ_t converges \mathbb{P} -almost surely to ψ and (4.13) holds by definition of ψ . \square

STEP 5: For every $t \in \mathbb{N}$, for every $h \in H_t$, $\tau(\sigma|h) = 1$.

Proof of Step 5: Let some $t \in \mathbb{N}$ and some $h \in H_t$ be given. We define

$$F' = \{(a_1, a_2, \dots) \in A^{\mathbb{N}} \mid (h, a_1, a_2, \dots) \in F\}.$$

It is not difficult to show that F' satisfies [F1–F5]. In particular, it holds that $\mathbb{P}((\beta_1, \beta_2, \dots) \in F) = \mathbb{P}((\beta_1, \beta_2, \dots) \in F')$ for any sequence $(\beta_1, \beta_2, \dots)$ of i.i.d. Bernoulli random variables. This implies that F' satisfies [F4] and [F5] as we can take $v^{*'} = v^*$ and $v'_* = v_*$.

Let G' be the frequency-based minority game corresponding to F' . Define the strategy profile $\sigma' \in \Sigma$ by letting $\sigma_{t'}(h') = \sigma_{t+t'-1}(h, h')$ for every $t' \in \mathbb{N}$ and $h' \in H_{t'}$. Then, σ' is a subgame perfect ϵ -equilibrium of G' . Hence, Step 4 implies that $\mu_{\sigma', \emptyset}(F') = 1$. The result of Step 5 now follows since $\mu_{\sigma, h}(F) = \mu_{\sigma', \emptyset}(F')$. \square

STEP 6: The game G has no subgame perfect ϵ -equilibrium.

Proof of Step 6: Let some $t \in \mathbb{N}$ and some $h \in H_t$ be given. Since $\tau(\sigma|h, 0) = 1$ by Step 5, action 0 yields player t a payoff of 2 and action 1 a payoff of 1. It follows that $\sigma_t(h) \leq \epsilon < v_*$. By Step 1, we have $\tau(\sigma|\emptyset) = 0$, yielding a contradiction to the statement of Step 4. \square

Appendix B: The game of Example 4.4 has no subgame perfect 0-equilibrium

Suppose that the game has a subgame perfect 0-equilibrium, say $\gamma \in \Sigma$. Let E^c be the complement of E , so E^c denotes the set of plays (a_1, a_2, \dots) for which the set $\{t \in \mathbb{N} \mid a_t = 0\}$ is infinite. We proceed by considering two cases.

CASE 1: For every $h \in H$, $\mu_{\gamma,h}(E^c) = 0$.

Consider a player $t \in \mathbb{N}$ and a history $h \in H_t$. We have $v_t(\gamma|h, 1) = 1$ and $v_t(\gamma|h, 0) = 2$. Since γ is a subgame perfect 0-equilibrium, $\gamma_t(h)$ puts probability 1 on action 0. Since this is true for every player and every history, we obtain a contradiction to the assumption of case 1.

CASE 2: For some $h' \in H$, $\mu_{\gamma,h'}(E^c) > 0$.

By Lévy’s zero–one law (see “Appendix C”), there exists a history $h \in H$ extending h' such that

$$\mu_{\gamma,h}(E^c) > 1/2. \tag{4.14}$$

For every $k \in \mathbb{N}$, let $h^k = (h, 1^{k-1}, 0)$ be the history where after history h first $k - 1$ players take action 1 and next the last player takes action 0. For every $k \in \mathbb{N}$, we define $F_k = C(h^k)$, the set of plays extending history h^k , so the set of plays where it takes exactly k periods after history h to observe action 0. We define the singleton set of plays $F_\infty = \{(h, 1, 1, \dots)\}$. Observe that $\{F_1, F_2, \dots, F_\infty\}$ is a Borel partition of $C(h)$. We argue that

$$\mu_{\gamma,h}(E^c \cap F_k) \leq \frac{1}{2}\mu_{\gamma,h}(F_k), \quad k = 1, 2, \dots, \infty. \tag{4.15}$$

Intuitively, with respect to the measure $\mu_{\gamma,h}$, less than or equal to half of the plays in F_k are such that the action 0 is taken infinitely many times.

Take some $k \in \mathbb{N}$. If $\mu_{\gamma,h}(F_k) = 0$, then also $\mu_{\gamma,h}(E^c \cap F_k) = 0$ and inequality (4.15) holds. Let us therefore consider the case $\mu_{\gamma,h}(F_k) > 0$. Since F_k is equal to the cylinder set $C(h^k)$, we have

$$\mu_{\gamma,h}(E^c \cap F_k) = \mu_{\gamma,h^k}(E^c) \cdot \mu_{\gamma,h}(F_k). \tag{4.16}$$

Let t be the player who is active at history $g^k = (h, 1^{k-1})$. Then, player t chooses action 0 with positive probability at g^k . Since γ is a subgame perfect 0-equilibrium, it follows that $v_t(\gamma|g^k, 0) \geq v_t(\gamma|g^k, 1) = 1$. It holds that

$$v_t(\gamma|g^k, 0) = v_t(\gamma|h^k) = \mu_{\gamma,h^k}(E^c) \cdot 0 + (1 - \mu_{\gamma,h^k}(E^c)) \cdot 2,$$

so we obtain $\mu_{\gamma,h^k}(E^c) \leq 1/2$. Using equality (4.16), inequality (4.15) follows.

The set E^c is disjoint from F_∞ , hence $\mu_{\gamma,h}(E^c \cap F_\infty) = 0$, so inequality (4.15) holds.

Finally, since $\{F_1, F_2, \dots, F_\infty\}$ is a partition of $C(h)$, we have

$$\mu_{\gamma,h}(E^c) = \mu_{\gamma,h}(E^c \cap C(h)) = \sum_{k=1}^{\infty} \mu_{\gamma,h}(E^c \cap F_k) \leq \frac{1}{2} \cdot \sum_{k=1}^{\infty} \mu_{\gamma,h}(F_k) \leq \frac{1}{2},$$

obtaining a contradiction to (4.14).

Appendix C: Lévy's zero-one law

A general statement of Lévy's zero-one law can be found in, e.g., Bogachev (2007), Example 10.3.15. Here, we state a version that is sufficient for our purposes. Consider a Borel subset P of $A^{\mathbb{N}}$, and let $1_P : A^{\mathbb{N}} \rightarrow \{0, 1\}$ denote the indicator function of P , that is $1_P(p) = 1$ if $p \in P$ and $1_P(p) = 0$ if $p \in A^{\mathbb{N}} \setminus P$. Let $\sigma \in \Sigma$ be a strategy profile. For $t \in \mathbb{N}$, we define $\tau_t : A^{\mathbb{N}} \rightarrow [0, 1]$ by

$$\tau_t(a_1, a_2, \dots) = \mu_{\sigma, (a_1, \dots, a_t)}(P), \quad (a_1, a_2, \dots) \in A^{\mathbb{N}},$$

as the probability that the play belongs to P conditional on a history of length t .

Theorem A.1 *Let $\sigma \in \Sigma$ be a strategy profile and $h \in H$ be a history. Then, the sequence τ_1, τ_2, \dots converges to 1_P $\mu_{\sigma,h}$ -almost surely.*

According to Theorem A.1, the probability measure $\mu_{\sigma,h}$ assigns probability 1 to the set of plays p for which $\tau_t(p)$ converges to $1_P(p)$.

References

- Alós-Ferrer, C., Ritzberger, K.: Characterizing existence of equilibrium for large extensive form games: a necessity result. *Econ. Theory* **63**, 407–430 (2017). <https://doi.org/10.1007/s00199-015-0937-0>
- Asheim, G.B.: Intergenerational equity. *Ann. Rev. Econ.* **2**, 197–222 (2010)
- Balbus, L., Jaśkiewicz, A., Nowak, A.S.: Stochastic bequest games. *Games Econ. Behav.* **90**, 247–256 (2015)
- Bogachev, V.I.: *Measure Theory*, vol. II. Springer, Berlin (2007)
- Bruyère, V.: Computer aided synthesis: a game-theoretic approach. In: Charlier É., Leroy J., Rigo M. (eds) *Developments in Language Theory. DLT 2017. Lecture Notes in Computer Science*, vol 10396. Springer (2017)
- Callander, S., Hörner, J.: The wisdom of the minority. *J. Econ. Theory* **144**, 1421–1439 (2009)
- Chatterjee, K., Henzinger, T.A.: A survey of stochastic omega-regular games, *JCSS* 2012. (2012)
- Cingiz, K., Flesch, J., Herings, P.J.J., Predtetchinski, A.: Doing it now, later, or never. *Games Econ. Behav.* **97**, 174–185 (2016)
- Flesch, J., Predtetchinski, A.: Subgame-perfect epsilon-equilibria in perfect information games with common preferences at the limit. *Math. Oper. Res.* **41**, 1208–1221 (2016a)
- Flesch, J., Predtetchinski, A.: ϵ -equilibria in perfect information games with sigma-discrete discontinuities. *Econ. Theory* **61**, 479–495 (2016b). <https://doi.org/10.1007/s00199-015-0868-9>
- Flesch, J., Predtetchinski, A.: A characterization of subgame-perfect equilibrium plays in Borel games of perfect information. *Math. Oper. Res.* **42**, 1162–1179 (2017)
- Flesch, J., Kuipers, J., Mashiah-Yaakovi, A., Schoenmakers, G., Solan, E., Vrieze, K.: Perfect-information games with lower-semicontinuous payoffs. *Math. Oper. Res.* **35**, 742–755 (2010)

- Flesch, J., Kuipers, J., Mashiah-Yaakovi, A., Shmaya, E., Schoenmakers, G., Solan, E., Vrieze, K.: Non-existence of subgame-perfect epsilon-equilibrium in perfect information games with infinite horizon. *Int. J. Game Theory* **43**, 945–951 (2014)
- Gossner, O., Hörner, J.: When is the lowest equilibrium payoff in a repeated game equal to the payoff? *J. Econ. Theory* **145**, 63–84 (2010)
- Gossner, O., Tomala, T.: Secret correlation in repeated games with imperfect monitoring. *Math. Oper. Res.* **32**, 413–424 (2007)
- Grädel, E., Ummels, M.: Solution concepts and algorithms for infinite multiplayer games. In: *New Perspectives on Games and Interaction* (K. Apt and R. van Rooij, Eds.), vol. 4 of *Texts in Logic and Games*, pp. 151–178. Amsterdam University Press (2008)
- Hellwig, M., Leininger, W., Reny, P.J., Robson, A.J.: Subgame perfect equilibrium in continuous games of perfect information: an elementary approach to existence and approximation by discrete games. *J. Econ. Theory* **52**, 406–422 (1990)
- Jaśkiewicz, A., Nowak, A.S.: Stationary Markov perfect equilibria in risk sensitive stochastic overlapping generations models. *J. Econ. Theory* **151**, 411–447 (2014)
- Kechris, A.S.: *Classical Descriptive Set Theory*. Springer, Berlin (1995)
- Kuipers, J., Flesch, J., Schoenmakers, G., Vrieze, K.: Subgame-perfection in recursive perfect information games, where each player controls one state. *Int. J. Game Theory* **45**, 205–237 (2016)
- Laibson, D.: *Essays in Hyperbolic Discounting*. MIT Press, Cambridge (1994)
- Laibson, D.: Golden eggs and hyperbolic discounting. *Quart. J. Econ.* **112**, 443–477 (1997)
- Le Roux, S., Pauly, A.: Infinite sequential games with real-valued payoffs. *Proceedings of the Joint Meeting of the 23rd EACSL Annual Conference on Computer Science Logic and the 29th Annual ACM/IEEE Symposium on Logic in Computer Science Vol. 62*, pp. 1–10 (2014)
- Le Roux, S.: Infinite sequential Nash equilibrium. *Log. Methods Comput. Sci.* **9** (2013)
- Lindvall, T.: *Lectures on the Coupling Method*. Dover Publications, Mineola (1992)
- Mertens, J.-F.: Repeated games. *Proceedings of the International Congress of Mathematicians*, pp. 1528–1577 (1987)
- Mertens, J.-F., Sorin, S., Zamir, S.: *Repeated Games*. Cambridge University Press, Cambridge (2015)
- Peleg, B.: Equilibrium points for games with infinitely many players. *J. Lond. Math. Soc.* **1**, 292–294 (1969)
- Peleg, B., Yaari, M.E.: On the existence of a consistent course of action when tastes are changing. *Rev. Econ. Stud.* **40**, 391–401 (1973)
- Phelps, E.S., Pollak, R.A.: On second-best national saving and game-equilibrium growth. *Rev. Econ. Stud.* **35**, 185–199 (1968)
- Pollak, R.A.: Consistent planning. *Rev. Econ. Stud.* **35**, 201–208 (1968)
- Purves, R.A., Sudderth, W.D.: Perfect information games with upper semicontinuous payoffs. *Math. Oper. Res.* **36**, 468–473 (2011)
- Radner, R.: Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *J. Econ. Theory* **22**, 136–154 (1980)
- Renault, J., Scarsini, M., Tomala, T.: A minority game with bounded recall. *Math. Oper. Res.* **32**, 873–889 (2007)
- Renault, J., Scarlatti, S., Scarsini, M.: Discounted and finitely repeated minority games with public signals. *Math. Soc. Sci.* **56**, 44–74 (2008)
- Shiryayev, A.N.: *Probability*, 2nd edn. Springer, Berlin (1996)
- Solan, E.: Absorbing team games. *Games Econ. Behav.* **31**, 245–261 (2000)
- Strotz, R.H.: Myopia and inconsistency in dynamic utility maximization. *Rev. Econ. Stud.* **23**, 165–180 (1956)
- von Stengel, B., Koller, D.: Team-maxmin equilibria. *Games Econ. Behav.* **21**, 309–321 (1997)
- von Stengel, B., Zamir, S.: Leadership games with convex strategy sets. *Games Econ. Behav.* **69**, 446–457 (2010)
- Voorneveld, M.: The possibility of impossible stairways: tail events and countable player sets. *Games Econ. Behav.* **68**, 403–410 (2010)