

Perfect Information Games where Each Player Acts Only Once

Citation for published version (APA):

Cingiz, K., Flesch, J., Herings, P. J.-J., & Predtetchinski, A. (2016). *Perfect Information Games where Each Player Acts Only Once*. Maastricht University, Graduate School of Business and Economics. GSBE Research Memoranda No. 036 <https://doi.org/10.26481/umagsb.2016036>

Document status and date:

Published: 01/01/2016

DOI:

[10.26481/umagsb.2016036](https://doi.org/10.26481/umagsb.2016036)

Document Version:

Publisher's PDF, also known as Version of record

Please check the document version of this publication:

- A submitted manuscript is the version of the article upon submission and before peer-review. There can be important differences between the submitted version and the official published version of record. People interested in the research are advised to contact the author for the final version of the publication, or visit the DOI to the publisher's website.
- The final author version and the galley proof are versions of the publication after peer review.
- The final published version features the final layout of the paper including the volume, issue and page numbers.

[Link to publication](#)

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal.

If the publication is distributed under the terms of Article 25fa of the Dutch Copyright Act, indicated by the "Taverne" license above, please follow below link for the End User Agreement:

www.umlib.nl/taverne-license

Take down policy

If you believe that this document breaches copyright please contact us at:

repository@maastrichtuniversity.nl

providing details and we will investigate your claim.

Kutay Cingiz, János Flesch,
P. Jean-Jacques Herings,
Arkadi Predtetchinski

**Perfect Information Games
Where Each Player Acts Only
Once**

RM/16/036

GSBE

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

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

Perfect Information Games Where Each Player Acts Only Once*

Kutay Cingiz[†] János Flesch[‡] P. Jean-Jacques Herings[§]
Arkadi Predtetchinski[¶]

November 14, 2016

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 admit no subgame perfect ϵ -equilibrium for small positive values of ϵ . Furthermore we derive 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 Numbers: C72, C73, D91.

*The authors would like to thank Nate Eldredge and Natalia Chernova for their contribution to the proof of Theorem 4.4 and Abraham Neyman and Jérôme Renault for their valuable comments.

[†]Department of Economics, Maastricht University, P.O. Box 616, 6200 MD, Maastricht, The Netherlands. K.Cingiz@maastrichtuniversity.nl

[‡]Department of Quantitative Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands. J.Flesch@maastrichtuniversity.nl

[§]Department of Economics, Maastricht University, P.O. Box 616, 6200 MD, Maastricht, The Netherlands. P.Herings@maastrichtuniversity.nl

[¶]Department of Economics, Maastricht University, P.O. Box 616, 6200 MD, Maastricht, The Netherlands. A.Predtetchinski@maastrichtuniversity.nl

1 Introduction

We study perfect information games where each player is active only once. Our emphasis is on the question of existence of subgame perfect ϵ -equilibria. 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 arena 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 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 safe action: taking this action leads to an intermediate payoff irrespective of what the other players do. The other is a risky action: it gives a high payoff if a minority of the players use it, and a low payoff otherwise. 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. Our main result on frequency-based minority games states that these games admit no subgame perfect ϵ -equilibria for small positive ϵ .

To motivate our construction, we consider two examples introduced in Peleg (1969) and Voorneveld (2010) that were used to highlight equilibrium existence issues in normal-form games with infinitely many players. We adapt them to our setting and turn them into perfect information games where each player is active only once. Under this formulation, they are very similar in spirit to frequency-based minority games. Nevertheless, we show that these examples do admit subgame perfect ϵ -equilibrium, underlining the subtlety of the definition of a frequency-based minority game.

Closely related to our work are two known counter-examples 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). The example shows that games of the type studied in this paper need not admit subgame perfect ϵ -equilibria in pure strategies. The other example is given in Flesch, Kuipers, Mashiah-Yaakovi, Shmaya, Schoenmakers, Solan, and Vrieze (2014). Being a game with two players, this example falls outside the class of games considered here. The authors show that the game admits no subgame perfect ϵ -equilibrium, whether in pure or mixed strategies. Frequency-based minority games thus demonstrate that even in games where each player only acts once, the existence of subgame perfect ϵ -equilibrium is not guaranteed.

Our second contribution is providing several conditions for existence of subgame perfect ϵ -equilibrium. Many of these conditions have previously appeared in the literature in

various contexts. We discuss each condition in turn.

[I] Games with qualitative objectives: These are games where each players' payoff function takes only two values, 0 or 1. This class is motivated by a vast computer science literature on qualitative games. The inventory of qualitative objectives considered in computer science includes reachability, safety, Büchi, parity, and Müller objectives. We show that games with qualitative objectives admit subgame perfect 0–equilibria in pure strategies. The proof is based on an iterative procedure from 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.

[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, Flesch, Herings, and Predtetchinski (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 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), Jaśkiewicz and Nowak (2014)). Another example is intergenerational games, where a player in the infinite sequence of players represents an entire generation (Phelps and Polak (1968), Balbus, Jaskiewicz, Nowak (2015)). In each period, a player decides upon the amount of savings, which affects available income in the next period. 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, Leininger, Reny, and Robson (1990). Callander and Hörner (2009) study social learning using games played by a sequence of players.

The paper is organized as follows. In Section 2, we provide the general model and formal definitions. In Section 3, we study the two examples as adapted from Peleg (1969) and Voorneveld (2010). In Section 4 we introduce frequency–based minority games and show

that they admit no subgame perfect ϵ -equilibria for small positive values of ϵ . In Section 5 we give several sufficient conditions for the existence of subgame perfect ϵ -equilibrium.

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 is said to extend a history $h \in H$ if h is a prefix of p . In a similar vein, 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 , 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 $t \in \mathbb{N}$ is denoted by Σ_t and the set of strategy profiles by $\Sigma = \times_{t \in \mathbb{N}} \Sigma_t$.

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 $S = \times_{t \in \mathbb{N}} S_t$.

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

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

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

Definition 2.1 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 shown that, for any game satisfying our assumptions, a pure strategy 0-equilibrium exists.¹

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.

Definition 2.2 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 $a_t \in A$,

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

The main emphasis of the paper is to study the existence of subgame perfect ϵ -equilibrium in the class of games satisfying our assumptions.

¹The argument is spelled out for games with finitely many players. It is clear that it remains valid in the present setting where the number of players is infinite.

3 Perfect information minority games

In this and the next section we discuss three particular games in detail. These three games share the following feature: each player chooses one out of two actions and the action that is chosen by a minority of the players yields the better payoff. The precise meaning of minority is what distinguishes the three games from each other.

Minority games with simultaneous moves have been considered in Peleg (1969) and Voorneveld (2010). Both authors present minority games that do 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 .

In this section we will revisit the minority games of Peleg (1969) and Voorneveld (2010), but within a setup with sequential moves and perfect information. Perfect information versions of both minority games have a 0–equilibrium in pure strategies, a fact that follows from the general result by Mertens and Neyman alluded to above.

Our focus is primarily on subgame perfect ϵ –equilibria. We show that the perfect information version of the minority game as presented in Voorneveld (2010) has a subgame perfect 0–equilibrium in pure strategies and that the perfect information version of the minority game as presented in Peleg (1969) has a subgame perfect ϵ –equilibrium for every positive value of ϵ .

We start with the perfect information version of the minority game as presented in Voorneveld (2010).

3.1 Perfect information version of Voorneveld (2010)

Consider the action set $A = \{0, 1\}$. For every $T \in \mathbb{N}$, we define the function $f_T : A^{\mathbb{N}} \rightarrow [0, 1]$ by

$$f_T(a_1, a_2, \dots) = \frac{1}{T} \sum_{t=1}^T a_t, \quad (a_1, a_2, \dots) \in A^{\mathbb{N}}. \quad (3.1)$$

Moreover, we define the function $f : A^{\mathbb{N}} \rightarrow [0, 1]$ by

$$f(a_1, a_2, \dots) = \limsup_{T \rightarrow \infty} f_T(a_1, a_2, \dots), \quad (a_1, a_2, \dots) \in A^{\mathbb{N}}.$$

The number $f(a_1, a_2, \dots)$ measures the frequency of action 1 in the sequence (a_1, a_2, \dots) . The function f is Borel–measurable.

Example 3.1 Perfect information version of Voorneveld (2010).

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

by

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

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 $1/2$.

Consider the game of Example 3.1. The pure strategy profile $s^V = (s_t^V)_{t \in \mathbb{N}} \in S$ is defined by

$$\begin{aligned} s_1^V(\emptyset) &= 1, \\ s_t^V(a_1, a_2, \dots, a_{t-1}) &= a_{t-1}, \quad t \in \mathbb{N} \setminus \{1\}, (a_1, a_2, \dots, a_{t-1}) \in H_t. \end{aligned}$$

According to s^V , player 1 chooses action 1 and subsequent players choose the same action as their predecessor. On the induced path, every player chooses action 1. The next result establishes that s^V is a subgame perfect 0-equilibrium in pure strategies.

Claim 3.2 *The strategy profile s^V is a subgame perfect 0-equilibrium in pure strategies of the game in Example 3.1.*

Proof: Take an arbitrary player $t \in \mathbb{N}$ and a history $h \in H_t$. If player t chooses any action a_t at history h and the other players follow s_{-t}^V , the induced play becomes

$$\rho(s_{-t}^V, a_t | h) = (h, a_t, a_t, a_t, \dots),$$

and hence

$$v_t(s_{-t}^V, a_t | h) = 0.$$

In view of (2.2), the pure strategy profile s^V is a subgame perfect 0-equilibrium indeed. \square

Example 3.1 is an example of a game with qualitative objectives considered in more detail in Section 5. There we argue that all such games admit a pure strategy subgame perfect 0-equilibrium.

We now turn to our second example, the perfect information version of the minority game as presented in Peleg (1969).

3.2 Perfect information version of Peleg (1969)

Example 3.3 Perfect information version of Peleg (1969).

The action set is equal to $A = \{0, 1\}$. Let E be the set of plays such that only finitely many players take action 0, so

$$E = \{(a_1, a_2, \dots) \in A^{\mathbb{N}} \mid \{t \in \mathbb{N} : a_t = 0\} \text{ is finite}\}. \quad (3.2)$$

For every $t \in \mathbb{N}$, player t 's payoff function is given by:²

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

Player t has the option to take the safe action 1 and receive a payoff of 1 irrespective of the actions chosen by the other players. Player t can also take the risky action 0 and receive a payoff of 2 when only finitely many other players choose action 0. Otherwise, the payoff when taking the risky action is equal to 0.

As noted earlier, the result of Mertens and Neyman implies that the game of Example 3.3 has a 0–equilibrium. One particular 0–equilibrium of the game 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. We show next that the game of Example 3.3 does not have a subgame perfect 0–equilibrium.

Claim 3.4 *The game of Example 3.3 has no subgame perfect 0–equilibrium.*

Proof: 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 $\{t \in \mathbb{N} : a_t = 0\}$ 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.

²To make the transition to the game studied in Section 4 more natural, 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.

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

By Lévy's zero-one law,³ there exists a history $h \in H$ extending h' such that

$$\mu_{\gamma, h}(E^c) > 1/2. \quad (3.3)$$

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 partition of $C(h)$. We argue that

$$\mu_{\gamma, h}(E^c \cap F_k) \leq \frac{1}{2} \mu_{\gamma, h}(F_k), \quad k = 1, 2, \dots, \infty. \quad (3.4)$$

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 (3.4) 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

$$\mu_{\gamma, h}(E^c \cap F_k) = \mu_{\gamma, h^k}(E^c) \cdot \mu_{\gamma, h}(F_k). \quad (3.5)$$

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 (3.5), inequality (3.4) follows.

The set E^c is disjoint from F_∞ , hence $\mu_{\gamma, h}(E^c \cap F_\infty) = 0$, so inequality (3.4) 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 (3.3). \square

Consider the game of Example 3.3 and let some $\epsilon \in (0, 1)$ be given. We define the strategy profile $\sigma^{P, \epsilon} \in \Sigma$ by specifying, for every $t \in \mathbb{N}$, for every $h \in H_t$,

$$\begin{aligned} \sigma_t^{P, \epsilon}(h)(1) &= 1 - \epsilon, \\ \sigma_t^{P, \epsilon}(h)(0) &= \epsilon. \end{aligned}$$

³For a formal statement of Lévy's zero-one law, see Appendix A.

Under strategy $\sigma_t^{P,\epsilon}$, player t puts probability ϵ on action 0, irrespective of the prevailing history. The strategy $\sigma_t^{P,\epsilon}$ is therefore stationary. We show next that $\sigma^{P,\epsilon} = (\sigma_t^{P,\epsilon})_{t \in \mathbb{N}}$ is a subgame perfect ϵ -equilibrium in the game of Example 3.3.

Claim 3.5 *Let some $\epsilon \in (0, 1)$ be given. The strategy profile $\sigma^{P,\epsilon}$ is a subgame perfect ϵ -equilibrium in the game of Example 3.3.*

Proof: Take some player $t \in \mathbb{N}$ and a history $h \in H_t$. Under $\sigma^{P,\epsilon}$, the probability that infinitely many players play action 0 after history $(h, 0)$ is equal to 1. It follows that $v_t(\sigma_{-t}^{P,\epsilon}, 0|h) = 0$. Since $v_t(\sigma_{-t}^{P,\epsilon}, 1|h) = 1$, we have

$$v_t(\sigma^{P,\epsilon}|h) = \epsilon \cdot 0 + (1 - \epsilon) \cdot 1 = 1 - \epsilon.$$

Thus condition (2.2) holds for every $a_t \in \{0, 1\}$. We conclude that $\sigma^{P,\epsilon}$ is a subgame perfect ϵ -equilibrium in the game of Example 3.3. \square

Example 3.3 falls within the class of games continuous outside a countable set, as defined below in Section 5. We argue that all such games admit a subgame perfect ϵ -equilibrium for each positive ϵ .

We remark that minority games have been studied in a variety of very different contexts. For an extensive list of related work we refer to Renault, Scarlatti, and Scarsini (2008) and Renault, Scarsini, and Tomala (2007).

4 Frequency-based minority games

In the preceding section we have discussed two examples of minority games and shown that both games have a subgame perfect ϵ -equilibrium for every $\epsilon > 0$. This section is devoted to a class of frequency-based minority games that, unlike the two examples of the previous section, do not have a subgame perfect ϵ -equilibrium for small positive ϵ .

This section is divided into three subsections.

4.1 The definition

Before we give a definition of the class of frequency-based minority games, we start with an example. Recall the definition of the function f_T in (3.1).

Example 4.1 Consider a game where the action set is given by $A = \{0, 1\}$. For each $t \in \mathbb{N}$, player t 's payoff function is given by

$$u_t(a_1, a_2, \dots) = \begin{cases} 1, & \text{if } a_t = 1, \\ 2, & \text{if } a_t = 0 \text{ and } (a_1, a_2, \dots) \in F, \\ 0, & \text{if } a_t = 0 \text{ and } (a_1, a_2, \dots) \in A^{\mathbb{N}} \setminus F, \end{cases} \quad (4.1)$$

where

$$F = \{(a_1, a_2, \dots) \in A^{\mathbb{N}} : \limsup_{T \rightarrow \infty} f_T(a_1, a_2, \dots) > \frac{1}{2}\}. \quad (4.2)$$

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. The set F expresses the condition that only a minority of the players use action 0. More precisely, F consists of the plays such that the limsup frequency of action 1 exceeds $1/2$.

Thus the main difference between Examples 4.1 and 3.1 is that in 4.1 the two actions are not symmetric. The main difference to Example 3.3 lies in the difference between the sets F and E , a point to which we come back in the next subsection.

As follows from our main result, the game in Example 4.1 has no subgame perfect ϵ -equilibrium for any ϵ sufficiently small. Rather than proving this result specifically for the game of Example 4.1, we identify a whole class of games where the same result applies. These are the so-called frequency-based minority games.

We say that a set $F \subset A^{\mathbb{N}}$ is *admissible* if

[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] Let β_1, β_2, \dots be a sequence of i.i.d. random variables taking values in $\{0, 1\}$. There exists $\nu^* \in (0, 1)$ such that if each β_t is equal to 1 with probability ν^* , then

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

[F5] Let $\beta'_1, \beta'_2, \dots$ be a sequence of i.i.d. random variables taking values in $\{0, 1\}$. There exists $\nu_* \in (0, 1)$ such that if each β'_t is equal to 1 with probability ν_* , then

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

Definition 4.2 *A game as in Section 2 is a frequency-based minority game if $A = \{0, 1\}$, and for each $t \in \mathbb{N}$, player t 's payoff function is given by equation (4.1), where the set F is an admissible set.*

The set F describes the plays being such that action 0 is played by a minority of the players. Condition [F1] is a technical measurability condition. By Condition [F2], whether a play belongs to F or not depends only on the tail of the play. Condition [F3] expresses that if some play belongs to F and the set of players choosing action 1 is a subset of the set of players choosing action 1 at another play, then the other play belongs to F as well. 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.

4.2 Further examples of admissible sets

The set F in (4.2) is an example of an admissible set. In particular it satisfies condition [F4] with for instance $\nu^* = 2/3$ and condition [F5] with $\nu_* = 1/3$ as a consequence of the law of large numbers. The choice of $1/2$ as a lower bound on the frequency of action 1 in (4.2) is inessential. Any number in the open interval $(0, 1)$ yields an admissible set. Likewise the set

$$\{(a_1, a_2, \dots) \in A^{\mathbb{N}} : \liminf_{T \rightarrow \infty} f_T(a_1, a_2, \dots) > \frac{1}{2}\}$$

is admissible. A somewhat less obvious example of an admissible set is the following one:

$$\{(a_1, a_2, \dots) \in A^{\mathbb{N}} : \ell(f_1(a_1, a_2, \dots), f_2(a_1, a_2, \dots), \dots) > \frac{1}{2}\},$$

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

- $\ell(a_1, a_2, \dots) = \ell(f_1(a_1, a_2, \dots), f_2(a_1, a_2, \dots), \dots)$
- $\ell(a_1, a_2, a_3, \dots) = \ell(a_2, a_3, \dots)$,
- $\liminf_{T \rightarrow \infty} f_T(a_1, a_2, \dots) \leq \ell(a_1, a_2, \dots) \leq \limsup_{T \rightarrow \infty} f_T(a_1, a_2, \dots)$.

It is easy to use examples of admissible sets to create new examples, since the collection of admissible sets is closed under finite unions and finite intersections.

It follows from the topological zero–one law (Kechris [1995], Theorem 8.47) that an admissible set is either meagre or comeagre. The Hewitt–Savage zero–one law (Shiryaev [1996], p. 382) implies that if F is admissible, 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 admissible set is sandwiched between the sets E and E' , where E is given by (3.2) and

$$E' = \{(a_1, a_2, \dots) \in A^{\mathbb{N}} \mid \{t \in \mathbb{N} : a_t = 1\} \text{ is infinite}\}.$$

Claim 4.3 *If F is an admissible set, 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 sets E and E' are not admissible. More precisely, 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. When taking F equal to E in the game of Example 4.1, we obtain the game of Example 3.3. It follows from Claim 3.5 that, for every $\epsilon > 0$, this game admits a subgame perfect ϵ -equilibrium.

The set E' is not admissible either. Indeed E' satisfies [F1], [F2], [F3] [F4], but not [F5]. Consider the game G' obtained by taking F to be equal to E' in Example 4.1. The resulting game has subgame perfect ϵ -equilibrium in which each player plays 1 with probability ϵ .

4.3 The main result

The main result of this section is the following theorem.

Theorem 4.4 *Let G be a frequency-based minority game. Then G has no subgame perfect ϵ -equilibrium whenever $\epsilon < \min\{\nu_*, (1 - \nu^*)/3\}$.*

The proof of 4.4 makes use of a technique known in probability theory as coupling.

Before we embark on the proof of the theorem in its full generality, we discuss one special case of the result, the non-existence of a subgame perfect 0-equilibrium in pure strategies. This consequence of Theorem 4.4 can be proved in an elementary way.

Claim 4.5 *Let G be a frequency-based minority game. Then G has no subgame perfect 0-equilibrium in pure strategies.*

Proof: Suppose $s \in S$ is a subgame perfect 0-equilibrium in pure strategies of G . We argue first that, for every $h \in H$, $\rho(s|h) \in F$.

Take some player $t \in \mathbb{N}$ and a history $h \in H_t$. Denote the play $\rho(s|h)$ by $p = (h, a_t, a_{t+1}, \dots)$.

Suppose that p is not in F , so by Claim 4.3 it holds that $p \notin E$. Then there is a player $t' \geq t$ and a prefix $g = (h, a_t, \dots, a_{t'-1})$ of p in $H_{t'}$ such that $a_{t'} = s_{t'}(g) = 0$. Since $p \in A^{\mathbb{N}} \setminus F$, it holds that $v_{t'}(s|g) = u_{t'}(p) = 0$. But this contradicts the fact that s is a subgame perfect 0–equilibrium, since playing action 1 yields a payoff of 1. Consequently, it holds that $p \in F$.

We have shown that, for every player $t \in \mathbb{N}$, for every history $h \in H_t$, it holds that $\rho(s|h) \in F$. It follows that, for every player $t \in \mathbb{N}$, for every history $h \in H_t$, player t can get a payoff of 2 by taking action 0, so $s_t(h) = 0$. It follows that $\rho(s|\emptyset) = (0, 0, \dots)$. Since $(0, 0, \dots) \notin E'$, it holds by Claim 4.3 that $(0, 0, \dots) \notin F$, so we have obtained a contradiction to $\rho(s|\emptyset) = (0, 0, \dots) \in F$. \square

The intuition behind Claim 4.5 is as follows. It is argued first that after each history, the subgame perfect 0–equilibrium play is such that a minority plays the risky action 0. Otherwise, playing action 0 would lead to a payoff of 0 and every player would prefer to choose action 1, leading to a contradiction. But if the subgame perfect 0–equilibrium play is such that a minority plays action 0, then every player prefers to choose action 0, leading to a contradiction with the fact that a minority plays action 0.

4.4 The proof of Theorem 4.4

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

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

Our main tool to prove Theorem 4.4 is the technique of coupling from probability theory. It is used whenever the distributions of two distinct random variables defined on different probability spaces need to be compared to each other. The technique of coupling consists of creating a new, larger probability space, and two new random variables, having the same distributions as the two original ones.

To illustrate the technique, consider Step 1 of the proof. It asserts that if at each history the active player chooses action 1 with probability at least ν^* , then the resulting play belongs to the set F with probability 1. Intuitively, such a strategy profile generates at least the same number of 1's as a sequence of i.i.d. Bernoulli random variables each

taking value 1 with probability ν^* . The reverse statement is made for a strategy profile such that at each history the active player chooses action 1 with probability at most ν_* .

Let $\sigma \in \Sigma$ be a strategy profile such that $\sigma_t(h) \geq \nu^*$ for every $t \in \mathbb{N}$ and for every $h \in H_t$. In the proof of Step 1 we construct two sequences of random variables, β_1, β_2, \dots and $\alpha_1, \alpha_2, \dots$, defined on a single state space. The first sequence is as in condition [F4]. The second sequence induces the same conditional probabilities on actions as the strategy profile σ and satisfies $\alpha_t \geq \beta_t$ for every $t \in \mathbb{N}$.

Let G be a frequency-based minority game with the set F . Let ν^* be as in [F4] and ν_* as in [F5].

STEP 1: *Let $\sigma \in \Sigma$ be a strategy profile. If, for every $t \in \mathbb{N}$, for every $h \in H_t$, $\sigma_t(h) \geq \nu^*$, then $\tau(\sigma|\emptyset) = 1$. If $\sigma_t(h) \leq \nu_*$ for every $t \in \mathbb{N}$ and every $h \in H_t$, 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. β_1, β_2, \dots are i.i.d. with $\mathbb{P}(\beta_t = 1) = \nu^*$.
2. For every $(a_1, \dots, a_{t-1}) \in H_t$,

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

3. $\alpha_t \geq \beta_t$.

We define the set $B = \{0, 1\}$. Consider a game where the action set is equal to $A \times B = \{0, 1\} \times \{0, 1\}$. We define the strategy profile $\lambda : \times_{t \in \mathbb{N}} (A \times B)^{t-1} \rightarrow \Delta(A \times B)$ 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 given in Table 1, where $h = (a_1, \dots, a_{t-1})$.

| | $b_t = 0$ | $b_t = 1$ |
|-----------|-----------------------|-----------|
| $a_t = 0$ | $1 - \sigma_t(h)$ | 0 |
| $a_t = 1$ | $\sigma_t(h) - \nu^*$ | ν^* |

Table 1: The value of $\lambda_t(a_1, b_1, \dots, a_{t-1}, b_{t-1})(a_t, b_t)$.

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]. \square

Now choose an $\epsilon > 0$ such that $\epsilon < \nu_*$ and $1 - 3\epsilon > \nu^*$ and 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,

$$(1 - \sigma_t(h)) \cdot \tau(\sigma|h, 0) = \tau(\sigma|h) - \sigma_t(h) \cdot \tau(\sigma|h, 1). \quad (4.3)$$

Now player t 's expected payoff at history h is

$$v_t(\sigma|h) = \sigma_t(h) + 2 \cdot (1 - \sigma_t(h)) \cdot \tau(\sigma|h, 0). \quad (4.4)$$

Plugging equation (4.3) into equation (4.4), 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 since 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. \square

STEP 3: Defining 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.

1. β_1, β_2, \dots are i.i.d. with $\mathbb{P}(\beta_t = 1) = 1 - 3\epsilon$.
2. For every $(a_1, \dots, a_{t-1}) \in H_t$,

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

3. For every $(a_1, \dots, a_{t-1}) \in H_t$, if $\tau(\sigma|a_1, \dots, a_{t-1}) < \epsilon$ then $\alpha_t \geq \beta_t$.

We define the set $B = \{0, 1\}$. Consider a game where the action set is equal to $A \times B = \{0, 1\} \times \{0, 1\}$. We define the strategy profile $\lambda : \times_{t \in \mathbb{N}} (A \times B)^{t-1} \rightarrow \Delta(A \times B)$ 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 given in Table 2, where $h = (a_1, \dots, a_{t-1})$.

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

Table 2: The value of $\lambda_t(a_1, b_1, \dots, a_{t-1}, b_{t-1})(a_t, b_t)$.

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 non-negative 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 1, the fact that $1 - 3\epsilon > \nu^*$, and Step 1 imply that

$$\mathbb{P}((\beta_1, \beta_2, \dots) \in F) = 1. \quad (4.5)$$

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

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

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

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. 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 $(a_1, a_2, \dots) \in F$.

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 $\mu_{\sigma, \emptyset}(C) = 1$.⁴ Now we have

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

where the first equality follows because the sequence $(\alpha_1, \alpha_2, \dots)$ induces the probability measure $\mu_{\sigma, \emptyset}$ on $A^{\mathbb{N}}$.

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

$$\begin{aligned} & \mathbb{P}((\alpha_1, \alpha_2, \dots) \in F^c) \\ &= \mathbb{P}(\psi = 1) \end{aligned} \tag{4.6}$$

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

$$= \mathbb{P}(\exists m \in \mathbb{N} \text{ s.t. } (\beta_1, \beta_2, \dots) \in F \text{ and } \psi_m = \psi_{m+1} = \dots = 1) \tag{4.8}$$

$$\leq \mathbb{P}(\exists m \in \mathbb{N} \text{ s.t. } (\alpha_1, \alpha_2, \dots) \in F \text{ and } \psi_m = \psi_{m+1} = \dots = 1) \tag{4.9}$$

$$= \mathbb{P}((\alpha_1, \alpha_2, \dots) \in F \text{ and } \psi = 1) \tag{4.10}$$

$$= \mathbb{P}(\psi = 0 \text{ and } \psi = 1) \tag{4.11}$$

$$= 0,$$

where (4.6) holds by definition of ψ , (4.7) is true because of (4.5), and equation (4.8) is true because ψ_t converges \mathbb{P} -almost surely to ψ . To prove inequality (4.9), 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 (4.9) follows.

Finally (4.10) holds because ψ_t converges \mathbb{P} -almost surely to ψ and (4.11) 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

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

It is not difficult to show that F' is an admissible set. 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 $\nu^{*'} = \nu^*$ and $\nu_*' = \nu_*$.

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

⁴See Appendix A for a precise statement of Lévy's zero-one law.

subgame perfect ϵ -equilibrium of G' . Hence Step 4 implies that $\mu_{\sigma',\emptyset}(F') = 0$. 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 $\sigma_t(h) \leq \epsilon < \nu_*$. By Step 1 we have $\tau(\sigma|\emptyset) = 0$, yielding a contradiction to the statement of Step 4. \square

5 Sufficient conditions for existence

For games with finitely many players, a number of sufficient conditions for the existence of subgame perfect ϵ -equilibrium have been identified, see for instance Kuipers, Flesch, Schoenmakers and Vrieze (2016), Flesch and Predtetchinski (2016a, 2016b). Much less is known about games with infinitely many players.

Below we discuss four classes of games that admit a subgame perfect ϵ -equilibrium: 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. Example 3.1 falls into the class of games with qualitative objectives and Example 3.3 into the class of games continuous outside a countable set.

Throughout Section 5, we fix a game G satisfying the assumptions made in Section 2.

5.1 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 game of Example 3.1 is a game with qualitative payoffs. The literature on computer science provides many examples of qualitative objectives: the reachability objective, the safety objective, Büchi, parity, and Müller criteria are all examples of qualitative objectives, see the review by Chatterjee and Henzinger (2012). 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 a winning set of player t .

Theorem 5.1 *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 5.1. For each ordinal number ξ , we recursively define the collection $\{P_\xi(h) : h \in H\}$ of sets of

plays and the collection $\{\alpha_\xi(h) : h \in H\}$ 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

$$P_0(h) = \{p \in A^{\mathbb{N}} : h \text{ is a prefix of } p\}, \quad (5.1)$$

$$\alpha_0(h) = \min_{p \in P_0(h)} u_t(p). \quad (5.2)$$

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

$$\alpha_{\xi+1}(h) = \max_{a \in A} \min_{p \in P_\xi(h, a)} u_t(p), \quad (5.3)$$

$$P_{\xi+1}(h) = \{p \in \bigcup_{a \in A} P_\xi(h, a) : u_t(p) \geq \alpha_{\xi+1}(h)\}, \quad (5.4)$$

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

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

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

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.

Proof of Theorem 5.1: 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 (5.3). 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 $\alpha_\eta(h) = 1$. Let $a_t \in A$ be an action that reaches the maximum in (5.3) with $\xi + 1 = \eta$. Notice that $P_\eta(h, a_t) \subset W_{\iota(h)}$.

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 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.

5.2 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 $\{p \in A^{\mathbb{N}} : u_t(p) \geq r\}$ is closed. The game G is said to have upper semicontinuous payoffs if every player $t \in \mathbb{N}$ has upper semicontinuous payoffs.

Theorem 5.2 *If the game G has upper semicontinuous payoffs, then it admits a subgame perfect 0–equilibrium in pure strategies.*

Proof: We argue that the set $P_\xi(h)$ as defined in the previous subsection 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 (5.3). 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 $\{p \in A^{\mathbb{N}} : u_t(p) \geq \alpha_{\xi+1}(h)\}$ is compact because u_t is upper semicontinuous. We conclude that $P_{\xi+1}(h)$ is compact.

Let ξ be a limit ordinal and the statement is true for each $\lambda < \xi$. 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 for games with infinitely many players.

A related work is Le Roux and Pauly (2014), which considers games with bounded upper semicontinuous payoffs but do not assume that the payoff functions only take finitely many payoffs. The authors show that the game admits 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.

5.3 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 a countable. The game G is said to be *continuous outside a countable set*⁵ 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 $u_t|_C$, is a continuous function. The game of Example 3.3 is continuous outside a countable set. Let $C = A^{\mathbb{N}} \setminus E$ and note that the set E is countable. For every player $t \in \mathbb{N}$, the function $u_t|_C$ only depends on a_t and is hence continuous. The game G is said to have uniformly bounded payoff functions if there exists $B < \infty$ such that, for every $t \in \mathbb{N}$, for every $p \in A^{\mathbb{N}}$, $|u_t(p)| < B$.

Theorem 5.3 *If the game G has uniformly bounded payoff functions and is continuous outside a countable set, then it admits a subgame perfect ϵ -equilibrium for each $\epsilon > 0$.*

Proof: Take a positive δ so small that $2BM\delta < \epsilon$ where M is the cardinality of the action set A . For each player t 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 . Now 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 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 . Consequently the payoff function $v_i(\sigma|h)$ is continuous on Σ^δ . A standard fixed point argument could now be invoked to prove the existence of a subgame perfect 0-equilibrium of the game with restricted strategy spaces. Any such strategy profile is a subgame perfect ϵ -equilibrium of the original game. \square

Theorem 5.3 generalizes a result in Cingiz et al (2016) who consider so-called centipede games played by a sequence of players.

5.4 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 $u_t = u_{t'}$ whenever $t, t' \in T_k$ for some $k \in \{1, \dots, n\}$. The idea

⁵The condition that a game is continuous outside a countable set is weaker than the condition used in Flesch and Predtetchinski (2016b).

of teams consisting of individual players is not new, see for instance the references in the introduction.

Theorem 5.4 *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 applies 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

Appendix A: 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 5.5 *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 5.5, 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

- [1] Asheim G. B. (2010). Intergenerational equity. *Annual Review of Economics*, 2, 197–222.
- [2] Balbus L., Jaśkiewicz A., Nowak A. S. (2015). Stochastic bequest games. *Games and Economic Behavior*, 90, 247–256.

- [3] Bogachev V. I. (2007). Measure theory, Volume II. Springer-Verlag, Berlin.
- [4] Cingiz K., Flesch J., Herings P. J. J., Predtetchinski A. (2016). Doing it now, later, or never. *Games and Economic Behavior*, 97, 174–185.
- [5] Callander S., and Hörner J. (2009). The wisdom of the minority. *Journal of Economic Theory*, 144, 1421–1439.
- [6] Chatterjee K., and Henzinger T. A. (2012). A survey of stochastic omega-regular games. *JCSS 2012*.
- [7] Flesch J., Kuipers J., Mashiah-Yaakovi A., Schoenmakers G., Solan E., Vrieze K. (2010). Perfect-information games with lower-semicontinuous payoffs. *Mathematics of Operations Research*, 35, 742–755.
- [8] Flesch, J., Kuipers, J., Mashiah-Yaakovi, A., Shmaya, E., Schoenmakers, G., Solan, E. and Vrieze, K. (2014). Non-existence of subgame-perfect epsilon-equilibrium in perfect information games with infinite horizon. *International Journal of Game Theory*, 43, 945–951.
- [9] Flesch J., Predtetchinski A. (2016a). Subgame-perfect epsilon-equilibria in perfect information games with common preferences at the limit. *Mathematics of Operations Research*, 41, 1208–1221.
- [10] Flesch J., Predtetchinski A. (2016b). Subgame-perfect ϵ -equilibria in perfect information games with sigma-discrete discontinuities. *Economic Theory*, 61, 479–495.
- [11] Gimbert H. and Kelmendi E. (2015). Two-player perfect-information shift-invariant submixing stochastic games are half-positional. Preprint.
- [12] Gossner O. and Hörner J. (2010). When is the lowest equilibrium payoff in a repeated game equal to the payoff? *Journal of Economic Theory*, 145, 63–84.
- [13] Gossner O. and Tomala T. (2007). Secret correlation in repeated games with imperfect monitoring. *Mathematics of Operations Research*, 32, 413–424.
- [14] Grädel E., and Ummels. M. (2008). 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.
- [15] Hellwig M., Leininger W., Reny P. J., and Robson A. J. (1990). Subgame perfect equilibrium in continuous games of perfect information: an elementary approach to

- existence and approximation by discrete games. *Journal of Economic Theory*, 52, 406–422.
- [16] Jaśkiewicz A., Nowak, A. S. (2014). Stationary Markov perfect equilibria in risk sensitive stochastic overlapping generations models. *Journal of Economic Theory*, 151, 411–447.
- [17] Kechris A.S. (1995). *Classical Descriptive Set Theory*. Springer-Verlag, Berlin.
- [18] Kuipers J., Flesch J., Schoenmakers G., Vrieze K. (2016). Subgame-perfection in recursive perfect information games, where each player controls one state. *International Journal of Game Theory*, 45, 205–237.
- [19] Laibson D. (1994). *Essays in Hyperbolic Discounting*. MIT Press, Cambridge, Massachusetts.
- [20] Laibson D. (1997). Golden eggs and hyperbolic discounting. *Quarterly Journal of Economics*, 112, 443–477.
- [21] Le Roux S., Pauly A. (2014). 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*, 62, 1–10.
- [22] Mertens, J.-F. (1987): Repeated games. *Proceedings of the International Congress of Mathematicians*, 1528–1577.
- [23] Mertens, J.-F., Sorin S., Zamir S. (2015): *Repeated Games*. Cambridge University Press.
- [24] Peleg B. (1969). Equilibrium points for games with infinitely many players. *Journal of the London Mathematical Society*, 1, 292–294.
- [25] Peleg B., Yaari M. E. (1973). On the existence of a consistent course of action when tastes are changing. *Review of Economic Studies*, 40, 391–401.
- [26] Phelps E. S., Pollak R. A. (1968). On second-best national saving and game-equilibrium growth. *Review of Economic Studies*, 35, 185–199.
- [27] Pollak R. A. (1968). Consistent planning. *Review of Economic Studies*, 35, 201–208.
- [28] Purves R. A., Sudderth W. D. (2011). Perfect information games with upper semi-continuous payoffs. *Mathematics of Operations Research*, 36, 468–473.

- [29] Radner R. (1980). Collusive behavior in noncooperative epsilon-equilibria of oligopolies with long but finite lives. *Journal of Economic Theory*, 22, 136–154.
- [30] Renault J., Scarlatti S., Scarsini. M. (2008). Discounted and finitely repeated minority games with public signals. *Mathematical Social Sciences*, 56, 44–74.
- [31] Renault J., Scarsini M., Tomala. T. (2007). A minority game with bounded recall. *Mathematics of Operations Research*, 32, 873–889.
- [32] Shiryaev A.N. (1996). *Probability*. Second Edition, Springer-Verlag, Berlin.
- [33] Solan E. (2000). Absorbing team games. *Games and Economic Behavior*, 31, 245–261.
- [34] Solan E., Vieille N. (2003). Deterministic multi-player Dynkin games. *Journal of Mathematical Economics*, 39, 911–929.
- [35] von Stengel B., Koller D. (1997). Team-maxmin equilibria. *Games and Economic Behavior*, 21, 309–321.
- [36] von Stengel B., Zamir S. (2010). Leadership games with convex strategy sets. *Games and Economic Behavior*, 69, 446–457.
- [37] Strotz R. H. (1956). Myopia and inconsistency in dynamic utility maximization. *Review of Economic Studies*, 23, 165–180.
- [38] Voorneveld M. (2010). The possibility of impossible stairways: tail events and countable player sets. *Games and Economic Behavior*, 68, 403–410.