



Existence of equilibria in repeated games with long-run payoffs

Galit Ashkenazi-Golan^a, János Flesch^b, Arkadi Predtetchinski^c, and Eilon Solan^d

Edited by Robert Aumann, Hebrew University of Jerusalem, Jerusalem, Israel; received March 26, 2021; accepted January 13, 2022

We consider repeated games with tail-measurable payoffs, i.e., when the payoffs depend only on what happens in the long run. We show that every repeated game with tail-measurable payoffs admits an ε -equilibrium, for every $\varepsilon > 0$, provided that the set of players is finite or countably infinite and the action sets are finite. The proof relies on techniques from stochastic games and from alternating-move games with Borel-measurable payoffs.

repeated games | Nash equilibrium | countably many players | tail-measurable payoffs

1. Introduction

Nash equilibrium, of central importance in strategic game theory and its applications [Holt and Roth (1)], exists in all finite games [Nash (2)]. It exists also in some infinitely repeated (IR) games [Shapley (3) and Sorin (4)]; in others, only ε -equilibria (see end of section 2.1) exist [Vieille (5, 6)]; and in others still, even they do not [Gale and Stewart (7)]. Here we prove that ε -equilibria exist in all IR games in which the payoff depends only on what happens in the long run, more precisely, whenever action streams that differ in only finitely many stages have the same payoff (technically, we require Borel measurability and tail measurability [section 2.2]). Moreover, we allow countably infinitely many players [Peleg (8) and Voorneveld (9)]. Thus, interestingly, even though one-shot games with countably infinitely many players and finite action sets may fail to have an ε -equilibrium, IR games with countably infinitely many players, finite action sets, and tail-measurable payoffs do have an ε -equilibrium. We remark that the existence of ε -equilibrium is not known even when the number of players in the game is finite.

The world is finite; infinite models capture those aspects of the finite world that are difficult to capture with finite models. For example, IR games capture the idea of a long repetition when it is not clear or too costly to compute how long the repetition is. Models with countably infinitely many players are used to capture interactions with many players, for example, models in which long-run players coexist with a sequence of short-run ones [Mailath and Samuelson (10)], overlapping generations models [Kandori and Obayashi (11)], large networks [Jackson (12)], and models with time-inconsistent decision makers [O'Donoghue and Rabin (13)]. IR games with countably infinitely many players model situations in which these two phenomena (long duration and many players) interact, for example, models in which a certain phenomenon occurring in a local area of a network propagates over time to the distant reaches of that network.

To prove the result we introduce a technique to the study of IR games with general payoff functions, which combines Martin's method (14) for Borel games and Blackwell games, with a method developed by Solan and Vieille (15) for stochastic games.

The structure of the paper is as follows. In section 2, we introduce the model, state the existence result, and discuss an illustrative example. In section 3, we prove the main result. In section 4 we make some concluding remarks.

2. Model and Main Result

2.1. The Model. *Definition:* An IR game with countably many players is a triplet $\Gamma = (I, (A_i)_{i \in I}, (f_i)_{i \in I})$, where

- $I = \{0, 1, 2, \dots\}$ is a countable set of players;
- A_i is a nonempty and finite action set for player i , for each $i \in I$ (let $A = \prod_{i \in I} A_i$ denote the set of action profiles, and $A^{\mathbb{N}}$ denote the set of plays, i.e., infinite streams of action profiles); and
- $f_i : A^{\mathbb{N}} \rightarrow \mathbb{R}$ is player i 's payoff function, for each $i \in I$.

The game is played in discrete time as follows. In each stage $t \in \mathbb{N} = \{0, 1, 2, \dots\}$, each player $i \in I$ selects an action $a_i^t \in A_i$, simultaneously with the other players. This

Significance

Nash equilibrium, of central importance in strategic game theory, exists in all finite games. Here we prove that it exists also in all infinitely repeated games, with a finite or countably infinite set of players, in which the payoff function is bounded and measurable and the payoff depends only on what is played in the long run, i.e., not on what is played in any fixed finite number of stages. To this end we combine techniques from stochastic games with techniques from alternating-move games with Borel-measurable payoffs.

Author affiliations: ^aDepartment of Mathematics, London School of Economics and Political Science, London WC2A 2AE, United Kingdom; ^bDepartment of Quantitative Economics, Maastricht University, 6200 MD Maastricht, The Netherlands; ^cDepartment of Economics, Maastricht University, 6200 MD Maastricht, The Netherlands; and ^dSchool of Mathematical Sciences, Tel-Aviv University, Tel-Aviv, Israel 6997800

Author contributions: G.A.-G., J.F., A.P., and E.S. wrote the paper.

The authors declare no competing interest.

This article is a PNAS Direct Submission.

Copyright © 2022 the Author(s). Published by PNAS. This article is distributed under Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 (CC BY-NC-ND).

¹To whom correspondence may be addressed. Email: eilons@tauex.tau.ac.il.

Published March 8, 2022.

induces an action profile $a^t = (a_i^t)_{i \in I}$ in stage t , which is observed by all players. Given the realized play $p = (a^0, a^1, \dots)$, the payoff to each player $i \in I$ is $f_i(p)$.

Histories. A history is a finite sequence of elements of A , including the empty sequence \emptyset . The set of histories is denoted by $H = \bigcup_{t=0}^{\infty} A^t$, where A^t is the t -fold product of A . The current stage (or length) of a history $h \in A^t$ is denoted by $\text{stage}(h) = t$. For a play $p \in A^{\mathbb{N}}$, write $p^{\leq t}$ to denote the prefix of p of length $t + 1$.

For two histories $h, h' \in H$ and a play $p \in A^{\mathbb{N}}$, we denote the concatenation of h and h' by hh' and that of h and p by hp .

Each history induces a subgame of Γ . Given $h \in H$, the subgame that starts at h is the game $\Gamma^h = (I, (A_i)_{i \in I}, (f_{i,h})_{i \in I})$, where $f_{i,h} = f_i \circ s_h$ and $s_h : A^{\mathbb{N}} \rightarrow A^{\mathbb{N}}$ is given by $p \mapsto hp$.

Measure theoretic setup. We endow the set A of action profiles, the sets A^t for $t \geq 2$, and the set $A^{\mathbb{N}}$ of plays with the product topology, where the action set A_i of each player $i \in I$ has its natural discrete topology. We denote the corresponding Borel sigma-algebras by $\mathcal{F}(A)$, $\mathcal{F}(A^t)$, and $\mathcal{F}(A^{\mathbb{N}})$. A function that is measurable with respect to one of these sigma-algebras is called product-measurable.

Strategies. A mixed action for player i is a probability measure x_i on A_i . The set of mixed actions for player i is denoted by X_i . A mixed action profile is a vector of mixed actions $x = (x_i)_{i \in I}$, one for each player. The set of mixed action profiles is denoted by X .

A (behavior) strategy of player i is a function $\sigma_i : H \rightarrow X_i$ that satisfies the following measurability condition: for each stage $t \in \mathbb{N}$ and each action $a_i \in A_i$, the mapping from A^t to $[0, 1]$ given by $h \in A^t \mapsto \sigma_i(h)(a_i)$ is product-measurable. We denote by Σ_i the set of strategies of player i .

A strategy profile is a vector of strategies $\sigma = (\sigma_i)_{i \in I}$, one for each player. We denote by Σ the set of strategy profiles.

Strategy profiles with finite support. A mixed action profile x is said to have finite support if only finitely many players randomize in x , i.e., all players except finitely many play a pure action with probability 1. We denote by X^F the set of mixed action profiles with finite support.

A strategy profile σ is said to have finite support (at each stage) if $\sigma(h) \in X^F$ for every history $h \in H$. The set of players who randomize may be history dependent, and its size may not be bounded. We denote by Σ^F the set of strategy profiles with finite support.

Definitions for the opponents of a player. For $i \in I$, let $-i$ denote the set of its opponents, $I \setminus \{i\}$. When restricting attention to players in $-i$, we can similarly define mixed action profiles $x_{-i} = (x_j)_{j \neq i}$, strategy profiles $\sigma_{-i} = (\sigma_j)_{j \neq i}$, mixed action profiles with finite support, and strategy profiles with finite support. The corresponding sets are denoted by X_{-i} , Σ_{-i} , X_{-i}^F , and Σ_{-i}^F , respectively.

Expected payoffs. For each history $h \in H$ the mixed action profile $\sigma(h) = (\sigma_i(h))_{i \in I}$ induces a unique probability measure $\mathbb{P}_{\sigma(h)}$ on $(A, \mathcal{F}(A))$ by Kolmogorov's extension theorem, and then σ induces a unique probability measure \mathbb{P}_{σ} on $(A^{\mathbb{N}}, \mathcal{F}(A^{\mathbb{N}}))$ by the Ionescu-Tulcea theorem. If player i 's payoff function f_i is bounded and product-measurable, then player i 's expected payoff under σ is

$$\mathbb{E}_{\sigma}[f_i] = \int_{A^{\mathbb{N}}} f_i(p) \mathbb{P}_{\sigma}(dp).$$

In that case, player i 's expected payoff under σ in the subgame starting at a history h is

$$\mathbb{E}_{\sigma}[f_i | h] = \int_{A^{\mathbb{N}}} f_{i,h}(p) \mathbb{P}_{\sigma_h}(dp),$$

where σ_h is the strategy profile defined as $\sigma_h(h') = \sigma(hh')$ for each history h' .

Equilibrium. Consider a game with countably many players and bounded and product-measurable payoffs. Let $\vec{\varepsilon} = (\varepsilon_i)_{i \in I}$, where $\varepsilon_i \geq 0$ for each player $i \in I$. A strategy profile σ^* is called an $\vec{\varepsilon}$ -equilibrium, if for every player $i \in I$ and every strategy $\sigma_i \in \Sigma_i$, we have $\mathbb{E}_{\sigma^*}[f_i] \geq \mathbb{E}_{(\sigma_i, \sigma_{-i}^*)}[f_i] - \varepsilon_i$.

2.2. Tail Measurability. A set $B \subseteq A^{\mathbb{N}}$ of plays is called tail if whenever $p = (a^0, a^1, a^2, \dots) \in B$ and $p' = (a'^0, a'^1, a'^2, \dots)$ satisfy $a'^t = a^t$ for every t sufficiently large, we have $p' \in B$. Intuitively, a set $B \subseteq A^{\mathbb{N}}$ is tail if changing finitely many coordinates does not change the membership relation for B . The tail sets of $A^{\mathbb{N}}$ form a sigma-algebra, denoted by $\mathcal{F}^{\text{Tail}}$. The sigma-algebras $\mathcal{F}^{\text{Tail}}$ and $\mathcal{F}(A^{\mathbb{N}})$ are generally not related: neither of them includes the other [Rosenthal (16) and Blackwell and Diaconis (17)]. A function from $A^{\mathbb{N}}$ to \mathbb{R} is called tail-measurable if it is measurable with respect to $\mathcal{F}^{\text{Tail}}$.

Various important evaluation functions in the literature of IR games are tail-measurable, for example, the limsup and liminf of the average stage payoffs [e.g., Sorin (4), p. 73]. Various classical winning conditions in the computer science literature, such as the Büchi, co-Büchi, parity, Streett, and Müller [e.g., Horn and Gimbert (18) and Chatterjee and Henzinger (19)], are also tail-measurable. The discounted payoff [e.g., Shapley (3)] is not tail-measurable.

2.3. The Main Result. We now present the main result of the paper.

Theorem 2.1. Consider an IR game with countably many players. If each player's payoff function is bounded, product-measurable, and tail-measurable, then the game admits an $\vec{\varepsilon}$ -equilibrium for each $\vec{\varepsilon} = (\varepsilon_i)_{i \in I}$, where $\varepsilon_i > 0$ for each $i \in I$.

We show the existence of an $\vec{\varepsilon}$ -equilibrium with a simple structure: The players are supposed to follow a play; that is, each player is supposed to play a fixed sequence of actions. If a player deviates, her opponents switch to a strategy profile with finite support to punish the deviator. Thus, it never happens that, on or off the equilibrium play, infinitely many players randomize at the same stage.

A 0-equilibrium does not always exist under the conditions of *Theorem 2.1*, not even in games with only one player (the other players can be treated as dummies). Indeed, suppose that this player has two actions, 0 and 1. For each play $p = (a^0, a^1, \dots)$, let $\zeta(p)$ denote the limsup frequency of action 1: $\zeta(p) = \limsup_{t \rightarrow \infty} (a^0 + \dots + a^t)/(t + 1)$. Define the payoff to be $\zeta(p)$ if $\zeta(p) < 1$ and 0 if $\zeta(p) = 1$. This payoff function is tail-measurable, the player can obtain a payoff arbitrarily close to 1, but the player cannot obtain 1.

The example in section 2.4 illustrates the main ideas of the construction of $\vec{\varepsilon}$ -equilibria in the proof of *Theorem 2.1*.

2.4. An Illustrative Example. Voorneveld (9) studied the following one-shot game with countably infinitely many players, which is a variation of the example in Peleg (8) and in which each player would like to choose the action that is chosen by the minority of the players.

- The set of players is $I = \mathbb{N}$.
- The action set of each player $i \in I$ is $A_i = \{0, 1\}$.
- The payoff of each player $i \in I$ is

$$r_i(a) = \begin{cases} 1, & \text{if } a_i = 0 \text{ and } \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{j < n} a_j > \frac{1}{2}, \\ 1, & \text{if } a_i = 1 \text{ and } \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{j < n} a_j \leq \frac{1}{2}, \\ 0, & \text{otherwise.} \end{cases} \quad [1]$$

Voorneveld (9) showed that this one-shot game has no $\bar{\varepsilon}$ -equilibrium when the ε_i values are not too large, for example, when $\varepsilon_i \in [0, 1/3]$ for every i . Indeed, suppose by way of contradiction that x^* is an $\bar{\varepsilon}$ -equilibrium. The set of action profiles

$$B = \left\{ a \in A : \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{j < n} a_j > \frac{1}{2} \right\} \subseteq \prod_{i \in I} A_i$$

is tail. Hence, by Kolmogorov's 0–1 law, either $\mathbb{P}_{x^*}(B) = 0$ or $\mathbb{P}_{x^*}(B) = 1$. In the former case, choosing action 0 gives payoff 0, while choosing action 1 gives payoff 1, and therefore the strategy x_i^* of each player i chooses action 1 with probability at least $1 - \varepsilon_i \geq 2/3$, which contradicts $\mathbb{P}_{x^*}(B) = 0$. The argument is analogous for the latter case. Hence, x^* cannot be an $\bar{\varepsilon}$ -equilibrium.

Consider now the following repeated version Γ of Voorneveld's game:

- The set of players is $I = \mathbb{N}$.
- The action set of each player $i \in I$ is $A_i = \{0, 1\}$.
- The payoff of each player $i \in I$ is equal to 1 if $r_i(a^t) = 1$ for all sufficiently large stages t , and it is equal to 0 otherwise, where r_i is defined in Eq. 1.

Thus, player i wins in Γ if she wins the one-shot Voorneveld game in all but finitely many stages. The payoff function of this game is product-measurable and tail-measurable, and hence, the game satisfies the conditions of *Theorem 2.1*.

Even though the one-shot Voorneveld game has no $\bar{\varepsilon}$ -equilibrium when $\varepsilon_i \in [0, 1/3]$ for every i , according to *Theorem 2.1* the IR game Γ does have an $\bar{\varepsilon}$ -equilibrium when $\varepsilon_i > 0$ for every i . In fact, the following strategy profile is a pure $\bar{0}$ -equilibrium in the game Γ : fix an arbitrary default action $\hat{a}_i \in \{0, 1\}$ for each player i ; let each player i play her default action \hat{a}_i in all stages $t < i$, and in each stage $t \geq i$, let her play action 0 if $(\hat{a}_{t+1}, \hat{a}_{t+2}, \dots)$ is an element of B and action 1 otherwise. Under this strategy profile, each player i wins the one-shot Voorneveld game in all stages $t \geq i$, i.e., $r_i(a^t) = 1$ for each $t \geq i$. Hence, in the game Γ each player's payoff is 1.

In the above $\bar{0}$ -equilibrium, the players are supposed to follow a specific play of the game. Since each player's payoff is maximal under this play, no player has an incentive to deviate from it.

The general construction shares one crucial feature with the example above: the $\bar{\varepsilon}$ -equilibrium play, denoted p , is pure. The key property of the play p is that it yields each player i at least her minmax value, possibly up to ε_i . The minmax value of a player is the highest expected payoff that the player can defend against any strategy profile of her opponents; see section 3.1 for a formal definition. If a player deviates from the play p , the other players start punishing the deviator at her minmax value. Showing that a play p with the desired property exists is the main challenge of the construction. It is to tackle this problem that Martin's (14) techniques and tools from stochastic games are called for.

To apply the results of Martin (14), we cannot allow countably infinitely many players to randomize at the same stage. We therefore do not use the classical minmax value of player i , where all opponents may randomize their actions. Rather, we will define

the concept of the finitistic minmax value of player i , restricting the opponents of player i to strategy profiles with finite support.

3. The Proof

In section 3.1 we introduce the notion of finitistic minmax value, designed specifically for games with countably infinitely many players. In section 3.2 we establish a relation between the finitistic minmax value of the IR game and the finitistic minmax values of a collection of one-shot games, where each one-shot game is associated with a history in the IR game. In section 3.3 we use these one-shot games to derive the existence of a play in which each player's payoff is at least her finitistic minmax value, up to some error term, and complete the proof of *Theorem 2.1*.

3.1. The Finitistic Minmax Value. In IR games it is customary to use threats of punishment at the minmax value to discipline players. In the presence of countably infinitely many players, when the payoff function is not product-measurable, the expected payoff and therefore the minmax value are not always well defined. In our proof we will use one-shot games where the payoff functions are not necessarily product-measurable; hence, we introduce a suitable related notion, the so-called finitistic minmax value. This notion imposes that all but finitely many of a player's opponents play a pure strategy. As a consequence, it is well defined regardless of whether the payoff function is product-measurable or not.

The finitistic minmax value in one-shot games. Consider a one-shot game $G = (I, (A_i)_{i \in I}, (r_i)_{i \in I})$, where I is a countable set of players, and for each $i \in I$, A_i is a nonempty and finite action set of player i and $r_i : A \rightarrow \mathbb{R}$ is her payoff function.

Consider a player $i \in I$. If r_i is bounded and $\mathcal{F}(A)$ -measurable, then each mixed action profile $x = (x_j)_{j \in I}$ induces an expected payoff $r_i(x)$ for player i . In the game G , the (classical) minmax value of player i is defined as

$$\text{val}_i(r_i) = \inf_{x_{-i} \in X_{-i}} \sup_{x_i \in X_i} r_i(x_i, x_{-i}). \quad [2]$$

The minmax value can be interpreted as the highest payoff that player i can defend against any mixed action profile of her opponents.

Recall that X_{-i}^F denotes the set of mixed action profiles $x_{-i} = (x_j)_{j \neq i} \in X_{-i}$ such that x_j is pure for all but finitely many $j \neq i$. Each mixed action profile $x = (x_i, x_{-i}) \in X_i \times X_{-i}^F$ induces a well-defined expectation $r_i(x)$ for any function $r_i : A \rightarrow \mathbb{R}$. The finitistic minmax value of player i in G is

$$\text{val}_i^F(r_i) = \inf_{x_{-i} \in X_{-i}^F} \sup_{x_i \in X_i} r_i(x_i, x_{-i}). \quad [3]$$

Since A_i is finite, we can replace $\sup_{x_i \in X_i}$ by $\max_{x_i \in X_i}$. The classical and the finitistic versions of the minmax value coincide for any function r_i that is bounded and $\mathcal{F}(A)$ -measurable.

We note that in Eqs. 2 and 3 we cannot replace the infimum by minimum. Indeed, suppose that $i = 0$, and player 0 has a single action, while all other players have two actions: a and b . Define r_0 to be $1/n$ if player $n > 0$ is the only opponent who plays action b and to be 1 otherwise. Then player 0's minmax value and finitistic minmax value are equal to 0, but player 0's payoff is never 0.

The finitistic minmax value in IR games. Let Γ be an IR game with countably many players. Assume that player i 's payoff function f_i is bounded and product-measurable. Define the finitistic minmax value of player i in the IR game Γ as

$$\text{Val}_i^F(\Gamma) = \inf_{\sigma_{-i} \in \Sigma_{-i}^F} \sup_{\sigma_i \in \Sigma_i} \mathbb{E}_{(\sigma_i, \sigma_{-i})} [f_i]. \quad [4]$$

Note the distinction between val_i^F and Val_i^F ; the former is used for the finitistic minmax value in one-shot games, and the latter is used for the finitistic minmax value in IR games.

Player i 's finitistic minmax value in the subgame starting at a history h is

$$\text{Val}_i^F(\Gamma^h) = \inf_{\sigma_{-i} \in \Sigma_{-i}^F} \sup_{\sigma_i \in \Sigma_i} \mathbb{E}_{(\sigma_i, \sigma_{-i})}[f_i | h].$$

When a player's payoff function is tail-measurable, her payoff is independent of the initial segment of the play. As a result, her finitistic minmax value is independent of the history.

Theorem 3.1. *Let Γ be an IR game with countably many players. Assume that some player i 's payoff function f_i is bounded, product-measurable, and tail-measurable. Then, $\text{Val}_i^F(\Gamma^h) = \text{Val}_i^F(\Gamma)$ for every history $h \in H$.*

3.2. Defending the Finitistic Minmax Value. We start with an auxiliary result. Let $\Gamma = (I, (A_i)_{i \in I}, (f_i)_{i \in I})$ be an IR game with countably many players. Fix a player $i \in I$. For a function $d : H \rightarrow \mathbb{R}$ and a history $h \in H$, we write $d^h : A \rightarrow \mathbb{R}$ to denote the function $a \mapsto d(ha)$. This way, given a function $d : H \rightarrow \mathbb{R}$, we can associate the following one-shot game $G(d^h)$ with each history $h \in H$: the action set of each player $j \in I$ is A_j , and the payoff of player i under each action profile $a \in A$ is $d^h(a) = d(ha)$. The payoff functions of the other players are irrelevant.

Lemma 3.1. *Let Γ be an IR game with countably many players. Consider a player $i \in I$ and assume that her payoff function f_i takes values in $[0, 1]$ and is product-measurable. Suppose moreover that $0 < \text{Val}_i^F(\Gamma)$ and let $0 < w < \text{Val}_i^F(\Gamma)$. There exists a function $d : H \rightarrow [0, 1]$ satisfying the following conditions:*

- Condition a:** At the empty history: $d(\emptyset) = w$.
- Condition b:** $d(h) \leq \text{val}_i^F(d^h)$ for every $h \in H$.
- Condition c:** $\limsup_{t \rightarrow \infty} d(p^{\leq t}) \leq f_i(p)$ for every play $p \in A^{\mathbb{N}}$.

Intuitively, the function d has the following properties. *Condition a* is the initialization. *Condition b* requires that at any history h , if the opponents of player i use a mixed strategy profile x_{-i} with finite support, then player i has a response x_i such that the expectation of the function d^h under the mixed action profile (x_i, x_{-i}) is at least $d(h)$. In other words, player i is able to defend the payoff of $d(h)$ in the one-shot game $G(d^h)$ against finitistic mixed action profiles of the other players. *Condition c* states that the actual payoff is at least the limsup of the sequence of values under d .

Before proving *Lemma 3.1*, we provide two examples.

Example 3.1: Consider the IR game in section 2.4. Define d recursively: let $w \in (0, 1)$ and set $d(\emptyset) = w$. Let $h \in H$ be such that $d(h)$ has been defined. If $d(h) = 0$, let $d(ha) = 0$ for each $a \in A$. If $d(h) > 0$, let $d(ha) = r_i(a)$ for each $a \in A$. The function d satisfies *Condition a* of *Lemma 3.1*, and it satisfies *Condition b* since $\text{val}_i^F(r_i) = 1$. To see that it satisfies *Condition c* of *Lemma 3.1*, suppose that $f_i(p) = 0$ for a play $p = (a^0, a^1, \dots)$. Then $r_i(a^k) = 0$ for some stage $k \in \mathbb{N}$, and therefore, $d(p^{\leq t}) = 0$ for all stages $t \geq k$. \diamond

Example 3.2: Let Γ be the following two-player zero-sum IR game: The action set is $A_1 = \{T, B\}$ for player 1 and $A_2 = \{L, R\}$ for player 2 (and all other players are dummies). The payoff function f_1 assigns payoff 1 to each play $p \in (A_1 \times A_2)^{\mathbb{N}}$ such that (T, L) or (B, R) are played infinitely many times and assigns payoff 0 to all other plays. One can think of Γ as a repeated matching pennies game

$$\begin{matrix} & L & R \\ T & \begin{pmatrix} 1 & 0 \end{pmatrix} \\ B & \begin{pmatrix} 0 & 1 \end{pmatrix} \end{matrix},$$

where player 1 wins if she receives payoff 1 infinitely often. The function f_1 is tail-measurable, and $\text{Val}_1^F(\Gamma^h) = 1$ for each history h . Moreover, one strategy of player 1 that guarantees the payoff 1 is playing the mixed action $(\frac{1}{2}, \frac{1}{2})$ at each stage.

Let $\psi_-, \psi_+ : [0, 1] \rightarrow [0, 1]$ be the functions $\psi_-(x) = x - \frac{1}{2} \min\{x, 1-x\}$ and $\psi_+(x) = x + \frac{1}{2} \min\{x, 1-x\}$. Define d recursively. Fix $w \in (0, 1)$ and let $d(\emptyset) = w$. If $d(h)$ has been defined, let $d^h : A \rightarrow [0, 1]$ be given as follows:

$$\begin{matrix} & L & R \\ T & \begin{pmatrix} \psi_+(d(h)) & \psi_-(d(h)) \end{pmatrix} \\ B & \begin{pmatrix} \psi_-(d(h)) & \psi_+(d(h)) \end{pmatrix} \end{matrix}. \quad [5]$$

Note that the value of d^h is exactly $d(h)$. Also, if p is a play where both (T, L) and (B, R) are played only finitely many times, then $d(p^{\leq t})$ is monotonically decreasing after some stage and is approaching 0. Thus, the function d satisfies the conditions of *Lemma 3.1* for $i = 1$. \diamond

The proof of *Lemma 3.1* is largely based on the arguments in Martin (14) and Maitra and Sudderth (20). These arguments involve an auxiliary zero-sum perfect information game in which the winning strategy of the first player produces the desired function d . We provide a sketch of the argument, noting what modifications to Martin's and Maitra and Sudderth's proofs are required, but omit the details.

Proof of Lemma 3.1. Given a game Γ , a player i satisfying the conditions of *Lemma 3.1*, and a number $w \in (0, 1]$, consider the following perfect information game, denoted $\mathcal{M}_i(\Gamma, w)$:

- Player I chooses a one-shot payoff function $r^0 : A \rightarrow [0, 1]$ such that $\text{val}_i^F(r^0) \geq w$.
- Player II chooses an action profile $a^0 \in A$ such that $r^0(a^0) > 0$.
- Player I chooses a one-shot payoff function $r^1 : A \rightarrow [0, 1]$ such that $\text{val}_i^F(r^1) \geq r^0(a^0)$.
- Player II chooses an action profile $a^1 \in A$ such that $r^1(a^1) > 0$.
- Player I chooses a one-shot payoff function $r^2 : A \rightarrow [0, 1]$ such that $\text{val}_i^F(r^2) \geq r^1(a^1)$, and so on.

To distinguish between histories and plays in Γ and in $\mathcal{M}_i(\Gamma, w)$, we call the latter positions and runs. The outcome of $\mathcal{M}_i(\Gamma, w)$ is a run $(r^0, a^0, r^1, a^1, \dots)$. Player I wins $\mathcal{M}_i(\Gamma, w)$ if

$$\limsup_{t \rightarrow \infty} r^t(a^t) \leq f_i(a^0, a^1, \dots).$$

Note that player I has a legal move at any stage of the game, for instance, the function $r : A \rightarrow \mathbb{R}$ such that $r(a) = 1$ for all $a \in A$. By induction one can show that player II always has a legal move as well. Let T be the set of all positions in the game $\mathcal{M}_i(\Gamma, w)$. Then T is a pruned tree on the set $R \cup A$, where R denotes the set of functions $r : A \rightarrow [0, 1]$. One can check that player I's winning set in $\mathcal{M}_i(\Gamma, w)$ is a Borel subset of the set of runs, where the set of runs is endowed with the product topology and $R \cup A$ is endowed with the discrete topology. By Martin (21), $\mathcal{M}_i(\Gamma, w)$ is determined: either player I has a winning strategy in the game or player II does. We next establish the following fact.

Claim 3.1: Player I has a winning strategy in $\mathcal{M}_i(\Gamma, w)$ whenever $0 < w < \text{Val}_i^F(\Gamma)$.

Sketch of the proof of Claim 3.1. Fix w such that $0 < w < \text{Val}_i^F(\Gamma)$ and suppose that player I has no winning strategy in $\mathcal{M}_i(\Gamma, w)$. Then player II does. Fix player II's winning strategy, say σ_{II} , and $0 < \delta < \text{Val}_i^F(\Gamma) - w$. We spell out the key definitions and give a sketch of the steps of the proof, which follows closely the arguments in Maitra and Sudderth (ref. 20, lemmata 4.3 and 4.4).

Step 1: One defines a notion of consistency (with respect to σ_{II} and δ) of a history of the IR game Γ and simultaneously a function ψ mapping each consistent history of Γ into a player I's position in $\mathcal{M}_i(\Gamma, w)$ of the same length.

The definition is recursive in the length of the history. The empty history \emptyset of Γ is consistent; we define $\psi(\emptyset)$ to be the empty position of $\mathcal{M}_i(\Gamma, w)$. Take a $t \in \mathbb{N}$ and a consistent history $h^t \in A^t$ for which player I's position $\psi(h^t)$ has been defined; let $a^t \in A$ and consider the history (h^t, a^t) in Γ . Let $R(h^t, a^t)$ be the set of player I's legal moves r^t at position $\psi(h^t)$ in $\mathcal{M}_i(\Gamma, w)$ to which σ_{II} responds with a^t :

$$R(h^t, a^t) = \{r^t \in R : (\psi(h^t), r^t) \in T \text{ and } \sigma_{II}(\psi(h^t), r^t) = a^t\}.$$

If $R(h^t, a^t)$ is not empty, (h^t, a^t) is defined to be consistent. In this case, define

$$u_{h^t}(a^t) = \inf\{r^t(a^t) : r^t \in R(h^t, a^t)\}. \quad [6]$$

To define $\psi(h^t, a^t)$, choose $r^t \in R(h^t, a^t)$ so that

$$r^t(a^t) \leq u_{h^t}(a^t) + \delta \cdot 3^{-t-2}, \quad [7]$$

and let

$$\psi(h^t, a^t) = (\psi(h^t), r^t, a^t).$$

Step 2: It holds that

$$\text{val}_i^F(u_\emptyset) \leq w \text{ and } \text{val}_i^F(u_{h^t}) \leq r^{t-1}(a^{t-1}) \quad [8]$$

whenever $h^t = (h^{t-1}, a^{t-1})$ is a consistent history and $\psi(h^t) = (\psi(h^{t-1}), r^{t-1}, a^{t-1})$.

The proof of Step 2 is by contradiction. Assuming that $\text{val}_i^F(u_{h^t}) > \varepsilon + r^{t-1}(a^{t-1})$ for some $\varepsilon > 0$, we define player I's move at position $\psi(h^t)$ by $r^t = \max\{u_{h^t} - \varepsilon, 0\}$, argue that the move is legal, and consider player II's reaction a^t as prescribed by σ_{II} . Then $r^t \in R(h^t, a^t)$ and we obtain a contradiction to Eq. 6.

Step 3: Obtain a contradiction by showing that $\text{Val}_i^F(\Gamma) \leq w + \delta$.

Step 3.1: Define a process W^0, W^1, \dots on $A^\mathbb{N}$.

Define first $\overline{W}^t : A^t \rightarrow \mathbb{R}$ for each $t \in \mathbb{N}$ as follows: let $\overline{W}^0(\emptyset) = w$; for each $t \in \mathbb{N}$, let $\overline{W}^{t+1}(h^t, a^t) = u_{h^t}(a^t)$ whenever (h^t, a^t) is consistent and $\overline{W}^{t+1}(h^t, a^t) = 1$ otherwise. Now define $W^t : A^t \rightarrow \mathbb{R}$ by $W^t = \overline{W}^t + \delta \cdot 3^{-t}$ for each $t \in \mathbb{N}$. This defines a process on $A^\mathbb{N}$ adapted to the filtration $\{\mathcal{F}(A^t)\}_{t \in \mathbb{N}}$. From Eq. 7 and the fact that σ_{II} is winning, one deduces that $f_i \leq \limsup_{t \rightarrow \infty} W^t$ everywhere on $A^\mathbb{N}$.

Step 3.2: Define $\sigma_{-i} \in \Sigma_{-i}^F$, a strategy profile with finite support for i 's opponents in Γ .

For each history h^t of Γ , choose a mixed action profile with finite support $\sigma_{-i}(h) \in X_{-i}^F$ so that for each $x_i \in X_i$,

$$u_{h^t}(x_i, \sigma_{-i}(h^t)) \leq \text{val}_i^F(u_{h^t}) + \delta \cdot 3^{-t-1}. \quad [9]$$

At this point, a delicate issue is the measurability property of strategies. Recall that for each stage t , player j , and action a_j ,

the mapping $h \mapsto \sigma_j(h)(a_j)$ should be measurable. Since we have chosen $\sigma_{-i}(h)$ to be finitely supported, the number of histories that occur with positive probability at each stage under σ_{-i} is finite. Since the definition of the strategy in histories that occur with probability 0 is irrelevant, we can set the action chosen at these histories to be an arbitrary fixed action, thus ensuring that the measurability property is satisfied.

Step 3.3: The process W^t is a supermartingale with respect to the measure $\mathbb{P}_{(\sigma_i, \sigma_{-i})}$ for each $\sigma_i \in \Sigma_i$.

This follows by combining Eqs. 7–9.

Step 3.4: It holds that $\mathbb{E}_{(\sigma_i, \sigma_{-i})}(f_i) \leq w + \delta$ for each $\sigma_i \in \Sigma_i$.

This follows by the martingale convergence theorem and Step 3.1. This inequality contradicts the choice of δ . \square

Now, using Claim 3.1, we complete the proof of Lemma 3.1.

Take $0 < w < \text{Val}_i^F(\Gamma)$ and let σ_I be a winning strategy of player I in $\mathcal{M}_i(\Gamma, w)$. The strategy σ_I induces the function d with the desired properties, as follows. Define $d(\emptyset) = w$. Let $T_1 \subseteq T$ denote the set of positions in the game $\mathcal{M}_i(\Gamma, w)$ of even length (i.e., player I's positions) that are consistent with σ_I . Let $H_1 \subseteq H$ denote the set of histories of Γ consisting of the empty history \emptyset , and nonempty histories (a^0, \dots, a^{t-1}) for which there is a sequence r^0, \dots, r^{t-1} of elements of R such that $h^* = (r^0, a^0, \dots, r^{t-1}, a^{t-1})$ is an element of T_1 . If $h \in H \setminus H_1$, define $d(ha) = 0$ for each $a \in A$. If $h \in H_1$, let $r^t = \sigma_I(h^*)$ be player I's move at position h^* , and let $d(ha) = r^t(a)$ for each $a \in A$. Since r^t is determined by σ_I , h^* is uniquely determined by $h \in H_1$. The reader can verify that d satisfies Conditions a–c. \square

By using Lemma 3.1, we derive the existence of a function $d : H \rightarrow \mathbb{R}$ with certain desirable properties.

Theorem 3.2. Let Γ be an IR game with countably many players. Consider a player $i \in I$ and assume that her payoff function f_i takes values in $[0, 1]$ and is product-measurable. Then, for every $\delta > 0$, there exists a function $d : H \rightarrow [0, 1]$ satisfying the following conditions:

Condition M.1: $d(h) \leq \text{val}_i^F(d^h)$ for every $h \in H$.

Condition M.2: $\text{Val}_i^F(\Gamma^h) - \delta \leq \text{val}_i^F(d^h)$ for every $h \in H$.

Condition M.3: For every strategy profile with finite support $\sigma \in \Sigma^F$ such that

$$\text{val}_i^F(d^h) \leq \mathbb{E}_{\sigma(h)}[d^h] \quad \forall h \in H, \quad [10]$$

we have

$$\text{val}_i^F(d^h) \leq \mathbb{E}_\sigma[f_i | h] \quad \forall h \in H.$$

Condition M.1 is identical to Condition b of Lemma 3.1. Condition M.2 requires that at each history h , the finitistic minmax value $\text{val}_i^F(d^h)$ of the one-shot game $G(d^h)$ is not too low: it should be at least $\text{Val}_i^F(\Gamma^h) - \delta$. Condition M.3 links local conditions on the one-shot games $G(d^h)$, $h \in H$, with a global condition on the IR game: under any strategy profile with finite support, if in each one-shot game $G(d^h)$, player i 's expected payoff under the mixed actions played at h is at least $\text{val}_i^F(d^h)$, then in the IR game, for each history h , player i 's expected payoff in the subgame at h is also at least $\text{val}_i^F(d^h)$.

Example 3.3: The function d defined in Example 3.1 violates Condition M.2 since there are histories h such that the one-shot payoff function d^h is identically zero, whereas $\text{Val}_i^F(\Gamma^h) = 1$ for each history h .

Consider instead the function d defined by $d(\emptyset) = 1$ and $d(ha) = r_i(a)$ for each $h \in H$ and $a \in A$. Thus, $d^h = r_i$, and so $\text{val}_i^F(d^h) = 1$ for each h . Conditions M.1, M.2, and M.3 are all satisfied. To verify Condition M.3, consider a strategy profile $\sigma \in \Sigma^F$ satisfying Eq. 10. Then for each $h \in H$ the mixed action profile $\sigma(h)$ is supported by the set of action profiles $a \in A$ such

that $r_i(a) = 1$. This implies that $\mathbb{P}_\sigma(\cdot | h)$ is supported by the set of plays $p \in A^\mathbb{N}$ with $f_i(p) = 1$.

The function d does not satisfy *Condition c* of *Lemma 3.1*. Consider, for instance, the play p where at each stage all the players in $-i$ play 0, while player i plays 0 at even stages and 1 at odd stages. Then i 's reward is 0 at even stages and 1 at odd stages, implying that $f_i(p) = 0$, while $\limsup_{t \rightarrow \infty} d(p^{\leq t}) = 1$. In fact, one can show that in this example, there is no function d that simultaneously satisfies *Conditions M.2* and *M.3* of *Theorem 3.2* and *Condition c* of *Lemma 3.1*. \diamond

Example 3.4: The function d constructed in *Example 3.2* fails to satisfy *Condition M.2*: we have seen that $\text{val}_i^F(d^h)$ can be arbitrarily close to zero.

Fix $\delta \in (0, \frac{2}{3})$, and consider instead the function d defined recursively as follows. Let $d(\emptyset) = 1 - \delta$. Take $h \in H$ such that $d(h)$ has been defined. If $1 - \delta \leq d(h)$, define $d^h : A \rightarrow [0, 1]$ to be given by Eq. 5; if $d(h) < 1 - \delta$, define d^h to be the function

$$\begin{array}{cc} & \begin{array}{cc} L & R \end{array} \\ \begin{array}{c} T \\ B \end{array} & \begin{pmatrix} 1 - \frac{1}{2}\delta & 1 - \frac{3}{2}\delta \\ 1 - \frac{3}{2}\delta & 1 - \frac{1}{2}\delta \end{pmatrix}. \end{array}$$

Note that if $1 - \delta \leq d(h)$, then $d(h) = \text{val}_i^F(d^h)$, and if $d(h) < 1 - \delta$, then $1 - \delta = \text{val}_i^F(d^h)$. This shows that both *Conditions M.1* and *M.2* are satisfied. To verify *Condition M.3*, let σ satisfy Eq. 10. Then, for each $h \in H$, the measure $\sigma(h)$ places probability of at least $\frac{1}{2}$ on the set $\{(T, L), (B, R)\}$. This implies that $a^t \in \{(T, L), (B, R)\}$ for infinitely many $t \in \mathbb{N}$ almost surely with respect to the measure $\mathbb{P}_\sigma(\cdot | h)$. Thus, $\mathbb{E}_\sigma[f_i | h] = 1$. \diamond

Proof of Theorem 3.2. Fix $\delta > 0$.

The function \widehat{D}_h . Consider an arbitrary history $h \in H$. Let $H_h \subseteq H$ be the set of all histories that extend h (including h itself). We define a function $\widehat{D}_h : H_h \rightarrow [0, 1]$ as follows.

Suppose first that $\text{Val}_i^F(\Gamma^h) > \delta/2$. Apply *Lemma 3.1* to the subgame Γ^h and $w = \text{Val}_i^F(\Gamma^h) - \delta/2$. This yields a function $\widehat{D}_h : H_h \rightarrow [0, 1]$ with the following properties:

- At the history h ,

$$\widehat{D}_h(h) = \text{Val}_i^F(\Gamma^h) - \frac{\delta}{2}. \quad [11]$$

- For every $g \in H_h$,

$$\widehat{D}_h(g) \leq \text{val}_i^F(\widehat{D}_h^g), \quad [12]$$

where the symbol \widehat{D}_h^g stands for the function d^g as defined in section 3.2, with $d = \widehat{D}_h$.

- For every play p that extends h we have

$$\limsup_{t \rightarrow \infty} \widehat{D}_h(p^{\leq t}) \leq f_i(p). \quad [13]$$

If $\text{Val}_i^F(\Gamma^h) \leq \delta/2$, then we let \widehat{D}_h be the constant zero function. This definition satisfies Eqs. 12 and 13, whereas Eq. 11 has to be replaced by

$$\text{Val}_i^F(\Gamma^h) - \frac{\delta}{2} \leq \widehat{D}_h(h) = 0. \quad [14]$$

The function d . We recursively define a function $d : H \rightarrow [0, 1]$ and, simultaneously, an auxiliary function $\alpha : H \rightarrow H$, which assigns to each history h a prefix $\alpha(h)$ of h .

- At the empty history, set $\alpha(\emptyset) = \emptyset$ and $d(\emptyset) = \widehat{D}_{\alpha(\emptyset)}(\emptyset) = \widehat{D}_\emptyset(\emptyset) \geq \text{Val}_i^F(\Gamma) - \delta/2$.
- Consider a history $h \neq \emptyset$, and suppose that $\alpha(g)$ and $d(g)$ are already defined for each strict prefix g of h . Let h^- denote the prefix of h satisfying $\text{stage}(h) = \text{stage}(h^-) + 1$.
 - Set $d(h) = \widehat{D}_{\alpha(h^-)}(h)$.
 - If $\text{Val}_i^F(\Gamma^h) - \delta \leq d(h)$, set $\alpha(h) = \alpha(h^-)$. Otherwise, set $\alpha(h) = h$; in this case, we say that d reinitiates at h .

The intuition behind the definition of the function d is the following. We start at $h_0 = \emptyset$ with the function \widehat{D}_{h_0} and stick with it until we encounter a history h_1 such that $\widehat{D}_{h_0}(h_1) < \text{Val}_i^F(\Gamma^{h_1}) - \delta$. From this point on, the function \widehat{D}_{h_0} is no longer useful, and reinitiation occurs, which takes effect from the next stage, $\text{stage}(h_1) + 1$. As of that stage, d follows the function \widehat{D}_{h_1} , until encountering a history h_2 such that $\widehat{D}_{h_1}(h_2) < \text{Val}_i^F(\Gamma^{h_2}) - \delta$. As of stage $\text{stage}(h_2) + 1$, d follows \widehat{D}_{h_2} , and so on.

We show that d satisfies *Conditions M.1–M.3*.

Verifying Conditions M.1 and M.2. For each history $h \in H$, as $d(ha) = \widehat{D}_{\alpha(h)}(ha)$ for all action profiles $a \in A$, we have

$$d^h = \widehat{D}_{\alpha(h)}^h. \quad [15]$$

For the empty history \emptyset , by Eqs. 11, 12, 14, and 15,

$$\begin{aligned} \text{Val}_i^F(\Gamma) - \delta &< \text{Val}_i^F(\Gamma) - \frac{\delta}{2} \leq \widehat{D}_\emptyset(\emptyset) = d(\emptyset) \\ &\leq \text{val}_i^F(\widehat{D}_\emptyset^\emptyset) = \text{val}_i^F(d^\emptyset), \end{aligned}$$

which proves *Conditions M.1* and *M.2* for \emptyset .

Consider now a history $h \neq \emptyset$ at which d does not reinitiate. Since $\alpha(h) = \alpha(h^-)$, by Eqs. 12 and 15,

$$\begin{aligned} \text{Val}_i^F(\Gamma^h) - \delta \leq d(h) &= \widehat{D}_{\alpha(h^-)}(h) = \widehat{D}_{\alpha(h)}(h) \\ &\leq \text{val}_i^F(\widehat{D}_{\alpha(h)}^h) = \text{val}_i^F(d^h), \end{aligned}$$

which proves *Conditions M.1* and *M.2* for h .

Finally, consider a history $h \in H$ at which d reinitiates. By Eqs. 11, 12, 14, and 15, and since $\alpha(h) = h$,

$$\begin{aligned} d(h) &< \text{Val}_i^F(\Gamma^h) - \delta < \text{Val}_i^F(\Gamma^h) - \frac{\delta}{2} \\ &\leq \widehat{D}_h(h) \leq \text{val}_i^F(\widehat{D}_h^h) = \text{val}_i^F(d^h), \end{aligned} \quad [16]$$

which proves *Conditions M.1* and *M.2* for h .

Verifying Condition M.3. Assume that a strategy profile with finite support $\sigma \in \Sigma^F$ satisfies Eq. 10.

Notice that Eq. 10 and *Condition M.1* imply that $d(h) \leq \mathbb{E}_{\sigma(h)}[d^h]$ for each history $h \in H$, and thus, the process $(d(h))_h$ is a bounded submartingale under σ (with respect to the usual filtration on $A^\mathbb{N}$). Hence, in particular, $(d(h))_h$ converges with probability 1 under σ . Notice also that if d reinitiates at a history $h \in H$, then by Eqs. 10 and 16 we have $d(h) + \frac{\delta}{2} \leq \text{val}_i^F(d^h) \leq \mathbb{E}_{\sigma(h)}[d^h]$.

This implies that under σ the expected number of reinitiations is bounded in any subgame. In particular,

$$\mathbb{P}_\sigma(\text{the number of reinitiations is finite} | g) = 1,$$

for every history g .

Fix a history g and $\rho > 0$. We will show that $\text{val}_i^F(d^g) \leq \mathbb{E}_\sigma[f_i | g]$. Let $T \geq \text{stage}(g)$ be a sufficiently large stage such that

$$\mathbb{P}_\sigma(\text{no reinitiation after stage } T | g) \geq 1 - \rho.$$

This implies

$$\mathbb{P}_\sigma(\alpha(p^{\leq T}) = \alpha(p^{\leq t}) \text{ for each } t \geq T | g) \geq 1 - \rho.$$

Hence,

$$\begin{aligned} \text{val}_i^F(d^g) &\leq \mathbb{E}_{\sigma(g)}[d^g] \\ &\leq \mathbb{E}_\sigma \left[\limsup_{t \rightarrow \infty} d(p^{\leq t}) | g \right] \\ &\leq \mathbb{E}_\sigma \left[\limsup_{t \rightarrow \infty} \widehat{D}_{\alpha(p^{\leq T})}(p^{\leq t}) | g \right] + \rho \\ &\leq \mathbb{E}_\sigma[f_i | g] + \rho, \end{aligned}$$

where the last inequality follows from Eq. 13. Since ρ is arbitrary, *Condition M.3* holds as well. \square

Remark: In the proof of *Theorem 3.2* the function d is updated whenever its value drops too much, that is, whenever its continuation cannot serve to show that the value of the current subgame is not far below $\text{Val}_i^F(\Gamma)$. This idea is reminiscent of, e.g., Rosenberg and Vieille (22), where to construct an optimal strategy one follows a discounted optimal strategy until a subgame is reached where the discounted value is far below the undiscounted value and then switches to a discounted optimal strategy with a higher discount factor, for which the discounted value is closer to the undiscounted value.

3.3. Proof of Theorem 2.1. In this section we complete the proof of *Theorem 2.1*. By applying a proper affine transformation to the payoff function of each player i and adjusting ε_i accordingly, we can assume without loss of generality that f_i takes values in $[0, 1]$.

The gist of the argument is as follows. For each player $i \in I$, let $d_i : H \rightarrow [0, 1]$ be a function satisfying the conditions of *Theorem 3.2* with $\delta = \frac{\varepsilon_i}{2}$, and fix arbitrarily some default action $\widehat{a}_i \in A_i$. We define, for each history $h \in H$, a one-shot game $G(h)$. This is a game with finitely many players: if $\text{stage}(h) = t$, then in $G(h)$, only the players $0, \dots, t$ are active; the other players are assumed to be playing their respective default actions. The payoff to an active player i is given by the function d_i^h . Nash equilibria of this game satisfy Eq. 10 for every player $i \in \{0, \dots, t\}$. By *Conditions M.2* and *M.3*, the corresponding strategy profile in the IR game Γ gives each player i , almost surely, a payoff of at least $\text{Val}_i^F(\Gamma) - \frac{\varepsilon_i}{2}$. This in turn implies that there exists a play p with a payoff of at least $\text{Val}_i^F(\Gamma) - \frac{\varepsilon_i}{2}$ for each player $i \in I$. Such a play yields an $\vec{\varepsilon}$ -equilibrium in Γ : if a player i deviates from this play, then she gets punished; i.e., her opponents switch to a strategy profile which ensures that player i 's payoff is at most $\text{Val}_i^F(\Gamma) + \frac{\varepsilon_i}{2}$.

The one-shot game $G(h)$. Fix a history h , and denote $t = \text{stage}(h)$. Let $G(h)$ be the one-shot game where

- The set of players is $\{0, 1, 2, \dots, t\}$.
- The action set of each player $i \in \{0, 1, 2, \dots, t\}$ is A_i .
- The payoff function $r_i^h : \prod_{j=0}^t A_j \rightarrow [0, 1]$ of each player $i \in \{0, 1, 2, \dots, t\}$ is given by

$$r_i^h(a_0, \dots, a_t) = d_i^h(a_0, \dots, a_t, \widehat{a}_{t+1}, \widehat{a}_{t+2}, \dots).$$

By Nash (2), the one-shot game $G(h)$ has an equilibrium $x(h) = (x_i(h))_{i=0}^t \in \prod_{i=0}^t X_i$. Since each player $i = 0,$

$1, 2, \dots, t$ is active in $G(h)$ and only finitely many of her opponents randomize, we have

$$\mathbb{E}_{x(h)}[r_i^h] \geq \text{val}_i^F(d^h) \quad \text{for each } i \in \{0, \dots, t\}. \quad [17]$$

The play p . For $t \in \mathbb{N}$, define the set $H_t^F \subseteq A^t$ as follows. Let H_0^F consist of the empty history \emptyset , and for $t \in \mathbb{N}$, let H_{t+1}^F be the set of histories $(a^0, \dots, a^t) \in A^{t+1}$ such that $a_i^k = \widehat{a}_i$ for $k \leq t$ and $k < i$. This is the set of histories that may result if the inactive players stick to their default actions. The set H_t^F is finite for each $t \in \mathbb{N}$.

For each player $i \in I$, let $\sigma_i \in \Sigma_i$ be the following strategy (compare with equation 9 in ref. 15): for each history $h \in A^t$,

$$\sigma_i(h) = \begin{cases} x_i(h), & \text{if } h \in H_t^F \text{ and } i \leq t, \\ \widehat{a}_i, & \text{otherwise.} \end{cases}$$

The function thus defined is a strategy; i.e., the function $h \mapsto \sigma_i(h)(a_i)$ is $\mathcal{F}(A^t)$ -measurable for every $a_i \in A_i$ and every $t \in \mathbb{N}$. Indeed, this function coincides with the function $h \mapsto x_i(h)(a_i)$ on the finite set of histories H_t^F , and it is constant on $A^t \setminus H_t^F$.

The resulting strategy profile $\sigma = (\sigma_i)_{i \in I}$ is an element of Σ^F . Consider any player $i \in I$. For every history $h \in H$ with $\text{stage}(h) \geq i$, we have, by Eq. 17,

$$\mathbb{E}_{\sigma(h)}[d_i^h] = \mathbb{E}_{x(h)}[r_i^h] \geq \text{val}_i^F(d^h).$$

By *Theorems 3.2* and *3.1* we deduce that for every history $h \in H$ with $\text{stage}(h) \geq i$,

$$\mathbb{E}_\sigma[f_i | h] \geq \text{val}_i^F(d^h) \geq \text{Val}_i^F(\Gamma^h) - \frac{\varepsilon_i}{2} = \text{Val}_i^F(\Gamma) - \frac{\varepsilon_i}{2}.$$

Since this holds for all histories beyond stage i , Lévy's 0–1 law implies that

$$\mathbb{P}_\sigma(f_i \geq \text{Val}_i^F(\Gamma) - \frac{\varepsilon_i}{2}) = 1.$$

Since the set I of players is countable, it follows that there exists a play $p \in A^\mathbb{N}$ such that

$$f_i(p) \geq \text{Val}_i^F(\Gamma) - \frac{\varepsilon_i}{2}, \quad \forall i \in I.$$

The $\vec{\varepsilon}$ -equilibrium in Γ . Based on the play p , the following grim-trigger strategy profile is an $\vec{\varepsilon}$ -equilibrium:

- The players follow the play p .
- If some player i deviates from p at stage t , then from stage $t + 1$ and on, all other players punish player i at her finitistic minmax value. That is, they bring down player i 's payoff to $\text{Val}_i^F(\Gamma) + \frac{\varepsilon_i}{2}$.

4. Discussion

Tail Measurability. In our proof, the tail measurability of the payoff functions was used in two arguments: first, in *Theorem 3.1* to conclude that the finitistic minmax value of a player in a subgame is independent of the history, and second, in the proof of *Theorem 2.1* to ensure that each player i can play arbitrarily during the first i stages without affecting her overall payoff. When the payoffs are not tail-measurable, an $\vec{\varepsilon}$ -equilibrium need not exist; Voorneveld's one-shot game, when viewed as an IR game in which only stage 0 action profile matters, provides an illustration.

Games with Finitely Many Players. Our model encompasses games with finitely many players, as we can assume that only finitely many players have more than one action. Existence of ε -equilibrium was not known even in this setup. Note that even in this case, the tail sigma-algebra $\mathcal{F}^{\text{Tail}}$ is unrelated to the Borel sigma-algebra $\mathcal{F}(A^\mathbb{N})$.

Ongoing Interaction. Conceptually, the current framework differs from the traditional framework of repeated games [Sorin (4)]. In the latter, one specific game—for example, the prisoner's dilemma [see, e.g., Fudenberg and Tirole (23), p. 9]—is played repeatedly, and the players are paid off at each stage. The long-run payoff is then defined as some kind of average of the stage payoffs. Here, on the other hand, stage payoffs are not explicitly defined; instead, the payoffs depend on the entire sequence of the action profiles, as in Gale and Stewart (7). One could still think of stages as representing stage games. These stage games may differ from each other and may evolve over time in a complex, history-dependent fashion; perhaps a prisoner's dilemma at one stage, a stag hunt (see, e.g., ref. 23, p. 20) at the next, a battle of the sexes (see, e.g., ref. 23, p. 18) next, and so on. The stage payoffs, which are not explicit, may be thought of as being implicit in the long-run payoffs. The long-run payoff may exhibit players' goals that cannot be described by an average of stage payoffs; for example, a player may care about the long-run

average payoff as well as the long-run average stage to stage fluctuation.

The theory of repeated games yields insights into phenomena such as altruism, cooperation, trust, loyalty, threats, and revenge [Sorin (4) and Aumann (24)], showing how they arise from the ongoing nature of the interaction. It takes a leap of the imagination to extend these insights to the real world, where people interact in different ways over time and their goals are complex. The current work formalizes this leap.

Data Availability. There are no data underlying this work.

ACKNOWLEDGMENTS. We thank Abraham Neyman, the Associate Editor, and three anonymous reviewers for their valuable comments. G.A.-G. acknowledges the support of the Israel Science Foundation (ISF), grants 217/17 and 722/18, and the National Science Foundation of China–ISF grant 2510/17. E.S. acknowledges the support of the ISF, grant 217/17. This work has been partly supported by European Cooperation in Science and Technology Action CA16228 European Network for Game Theory.

1. C. A. Holt, A. E. Roth, The Nash equilibrium: A perspective. *Proc. Natl. Acad. Sci. U.S.A.* **101**, 3999–4002 (2004).
2. J. F. Nash, Equilibrium points in n -person games. *Proc. Natl. Acad. Sci. U.S.A.* **36**, 48–49 (1950).
3. L. S. Shapley, Stochastic games. *Proc. Natl. Acad. Sci. U.S.A.* **39**, 1095–1100 (1953).
4. S. Sorin, "Repeated games with complete information" in *Handbook of Game Theory with Economic Applications*, R. Aumann, S. Hart, Eds. (Elsevier, ed. 1, 1992), vol. 1, pp. 71–107.
5. N. Vieille, Two-player stochastic games I: A reduction. *Isr. J. Math.* **119**, 55–91 (2000).
6. N. Vieille, Two-player stochastic games II: The case of recursive games. *Isr. J. Math.* **119**, 93–126 (2000).
7. D. Gale, F. M. Stewart, Infinite games with perfect information. *Contributions to Theory Games* **2**, 2–16 (1953).
8. B. Peleg, Equilibrium points for games with infinitely many players. *J. Lond. Math. Soc.* **44**, 292–294 (1969).
9. M. Voorneveld, The possibility of impossible stairways: Tail events and countable player sets. *Games Econ. Behav.* **68**, 403–410 (2010).
10. G. J. Mailath, L. Samuelson, *Repeated Games and Reputations: Long-Run Relationships* (Oxford University Press, 2006).
11. M. Kandori, S. Obayashi, Labor union members play an OLG repeated game. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 10802–10809 (2014).
12. M. O. Jackson, *Social and Economic Networks* (Princeton University Press, 2010).
13. T. O'Donoghue, M. Rabin, Doing it now or later. *Am. Econ. Rev.* **89**, 103–124 (1999).
14. D. A. Martin, The determinacy of Blackwell games. *J. Symb. Log.* **63**, 1565–1581 (1998).
15. E. Solan, N. Vieille, Correlated equilibrium in stochastic games. *Games Econ. Behav.* **38**, 362–399 (2002).
16. J. Rosenthal, Nonmeasurable invariant sets. *Am. Math. Mon.* **82**, 488–491 (1975).
17. D. Blackwell, P. Diaconis, A non-measurable tail set. *Lect. Notes Monogr. Ser.* **30**, 1–5 (1996).
18. F. Horn, H. Gimbert, Optimal strategies in perfect-information stochastic games with tail winning conditions. arXiv [Preprint] (2008). <https://arxiv.org/abs/0811.3978> (Accessed 19 November 2013).
19. K. Chatterjee, T. A. Henzinger, A survey of stochastic ω -regular games. *J. Comput. Syst. Sci.* **78**, 394–413 (2012).
20. A. Maitra, W. Sudderth, Finitely additive stochastic games with borel measurable payoffs. *Int. J. Game Theory* **27**, 257–267 (1998).
21. D. A. Martin, Borel determinacy. *Ann. Math.* **102**, 363–371 (1975).
22. D. Rosenberg, N. Vieille, The maxmin of recursive games with incomplete information on one side. *Math. Oper. Res.* **25**, 23–35 (2000).
23. D. Fudenberg, J. Tirole, *Game Theory* (MIT Press, 1991).
24. R. J. Aumann, War and peace. *Proc. Natl. Acad. Sci. U.S.A.* **103**, 17075–17078 (2006).