Large games and the law of large numbers

Nabil I. Al-Najjar

Department of Managerial Economics and Decision Sciences, Kellogg School of Management, Northwestern University, Evanston, IL 60208, USA

Abstract

This paper introduces *discrete large games* where the set of players is a countable dense 'grid' with a finitely additive distribution. In these games *any* function from player names to mixed actions is a legitimate strategy profile. No extraneous continuity or measurability conditions are assumed. Randomness can be modeled explicitly and an exact law of large numbers holds. Equilibria enjoy a strong purification property: *every* realization of *every* mixed strategy equilibrium is a pure strategy equilibrium almost surely.

Every continuum-player game has a discrete large game representation that preserves the original pay-offs, strategy profiles and equilibria. It is argued that strategy profiles in continuum-player games have an ambiguous meaning because measurability requirements force the smoothing out of individual variations. These variations have clear strategic meaning in finite-player games and can be expressed in discrete large games, but not when the set of players is the continuum.

© 2008 Elsevier Inc. All rights reserved.

JEL classification: C7; C6; H41

E-mail address: al-najjar@northwestern.edu.

URL: http://www.kellogg.northwestern.edu/faculty/alnajjar/htm/index.htm.

"All models are wrong;
some are useful."[1]

1. Introduction

1.1. Aims and motivation

Two methodologies dominate the study of anonymous strategic interactions in large populations. One approach focuses on the asymptotic properties of games with finite but increasing number of players. The other looks directly at idealized continuum-player games, in which arguments are often simpler and the insights more transparent.

The use of abstractions like continuum-player games is founded on the belief that they capture the same fundamental strategic phenomena as large finite-player games. As is well known, and discussed in greater details below, this intuition often lacks formal basis, and is sometimes plainly wrong.

This paper introduces *discrete large games* as a tool for analyzing large-scale strategic interactions. The new framework dispenses with extraneous measurability-type conditions that have no behavioral or strategic content while retaining the tractability of continuum-player games. The paper will also argue that discrete large games can be remarkably similar to finite-player games both conceptually and analytically.

1.2. Illustrative example and main questions

The main points can be illustrated using the following example:

Example 1 *($\bar{\Gamma}_m$; Mis-matching game).* The set of players is [0, 1] with the uniform distribution $\bar{\lambda}$. Each player chooses among two actions $\{U, D\}$. A strategy profile is a measurable function $\bar{\mu} : [0, 1] \to [0, 1]$, assigning to each player t a number $\bar{\mu}(t)$ interpreted as the probability of choosing U. Payoffs are symmetric and given by:

$$\bar{u}(U, \bar{\mu}) = \begin{cases} 1 & \text{if } \int \bar{\mu} \, d\bar{\lambda} < 0.5, \\ -1 & \text{if } \int \bar{\mu} \, d\bar{\lambda} > 0.5, \\ 0 & \text{otherwise,} \end{cases} \tag{1}$$

and $\bar{u}(D, \bar{\mu}) = -\bar{u}(U, \bar{\mu})$.

Interpreting $\int \bar{\mu} \, d\bar{\lambda}$ as the aggregate frequency with which U is chosen, a player gets 1 if he mismatches the majority choice, gets -1 if he matches, and 0 otherwise. Clearly, any equilibrium must satisfy $\int \bar{\mu} \, d\bar{\lambda} = 0.5$.

The game $\bar{\Gamma}_m$ is an example of a large game in the sense of Schmeidler (1973). The use of games like $\bar{\Gamma}$ is predicated on the belief that they capture the essentials of strategic interaction in large finite-player games. Viewed in this light, continuum-player games raise a number of issues:

1* *What are strategy profiles?* In finite-player games, any function from player names to actions is a legitimate strategy profile. On the other hand, in continuum-player games like $\bar{\Gamma}_m$

---

[1] George Box, industrial statistician.

profiles are required to be measurable functions. This imposes extraneous joint restrictions on players' choices that are anathema to the idea of non-cooperative play. Subsequent discussion illustrates that this restriction has important consequences.

2* *What are mixed strategies and how to compute expected payoffs?* In finite player games mixed strategies are random variables on a state space, with expected payoffs defined as their integrals. In a continuum-player game a strategy profile assigns to each player a number $\bar{\mu}(t)$ in $[0, 1]$ interpreted as the probability of choosing a particular action. But there is no state space, so expressions of expected payoffs and aggregate play are "acts of faith" that rest on a feeling for what the right answer ought to be.

3* *Can strategy profiles be purified?* In any sensible large finite-player version of $\bar{\Gamma}_m$ almost every realization of every mixed strategy equilibrium is an approximate equilibrium. Kalai (2004) calls this property self-purification and shows that it holds quite generally as the number of players increases. One would expect this property to hold exactly in large games.[2] The problem is that this requires a version of the law of large numbers, which is problematic in the continuum.

1.3. Outline and results

A *discrete large game* is a game with a countable set of players $T \subset [0, 1]$ and an atomless finitely additive probability measure λ defined on all subsets of T. Intuitively, think of the space of players as an infinitely fine discrete grid, rather than the full continuum.[3]

Informally, this paper shows that every continuum-player game $\bar{\Gamma}$ corresponds to a discrete large game Γ such that:

- The distribution of players and payoff structure are preserved.
- The set of strategy profiles in a continuum-player game $\bar{\Gamma}$ can be embedded, in a unique and "behaviorally equivalent" fashion, as a subset of the set of strategy profiles in Γ. Behavioral equivalence is a strong criterion that preserves the underlying strategic behavior.[4]
- The embedding of strategy profiles preserves equilibrium: A profile $\bar{\mu}$ is an equilibrium for $\bar{\Gamma}$ if and only if the corresponding profile μ is an equilibrium for Γ.
- Equilibria of Γ are limits of approximate equilibria in finite-player games.
- A single set of players T can be chosen so it works simultaneously for any payoff specification, any atomless distribution $\bar{\lambda}$, and all strategy profiles.

The upshot is that any pattern of strategic behavior arising in a continuum-player game has a mirror image in the corresponding discrete large game. And any equilibrium in the latter is obtained as limit of ϵ-equilibria in large finite-player games. Given this, the framework of discrete large games provides a way to resolve issues 1*–3* above:

---

[2] Schmeidler's (1973) well-known result makes the weaker claim that any large game possesses a pure strategy equilibrium. For example, the profile $\bar{\mu}'$ in which $\bar{\mu}'(t) = U$ if $t \leqslant 0.5$ and $\bar{\mu}'(t) = D$ otherwise is a pure strategy equilibrium for this game.

[3] The σ-algebra used is the power set. The term "discrete" may be justified by thinking of the power set as the Borel σ-algebra generated by the discrete topology.

[4] In a continuum-player game, two behaviorally equivalent profiles must be equal almost everywhere. The idea is to define a similar concept to compare profiles in two games with different sets of players (one countable, the other continuum). See Definition 10.

1. *Strategy profiles* include all function from players to actions, exactly as in finite-player games. Extraneous joint restrictions on players' actions, like measurability, are unnecessary since every function is measurable.
2. *Mixed strategies and expected payoffs in discrete large games* can be introduced in a manner identical to finite-player games.
3. *All equilibria in discrete large games can be purified.* Specifically, every realization of any mixed strategy equilibrium is a pure strategy equilibrium with identical distribution on every subinterval of players, almost surely.

1.4. Example revisited

For the mis-matching game $\bar{\Gamma}_m$ the corresponding discrete large game Γ_m is obtained by taking any countable dense subset $T \subset [0, 1]$ as a set of players. Let λ be any uniform, finitely additive probability distribution λ on the set of all subsets of T. Here, 'uniform' means that it assigns to any 'interval' of players $T \cap [a, b]$ a measure equal to its length $|a - b|$. Payoffs in the discrete large game Γ_m at a profile μ are defined exactly as in Eq. (1), except that we now use the integral $\int \mu \, d\lambda$ with respect to the finitely additive probability λ.

Notational Convention: Throughout the paper we shall maintain the notational convention of using 'bars' to indicate objects in continuum models.

The analysis of the discrete large games is remarkably similar to that of finite-player games. Any function $\mu : T \to \{U, D\}$ is a legitimate pure strategy profile because any subset of players, hence any function, is measurable (just as in finite-player games). As in these games, the natural state space here is that of all pure strategy profiles, denoted S. Assuming players randomize independently, any mixed strategy profile defines a unique (countably additive) probability distribution on this state space. Expected payoffs can be computed in the standard way.

As an example, consider the symmetric equilibrium profile $\bar{\mu}$ in $\bar{\Gamma}_m$ where $\bar{\mu}(t) = 0.5, t \in \bar{T}$. Its behaviorally equivalent strategy profile in Γ_m is just the restriction of $\bar{\mu}$ to T, namely the profile $\mu(t) = 0.5, t \in T$. Under μ, players are flipping i.i.d. coins to choose between U and D, giving rise to a unique probability distribution P on the set of pure strategy profiles S. The law of large numbers then implies that P-a.e. state s has the property that, in any interval $[a, b]$, half of the players choose U and the other half chooses D. One can show that almost any such s is a pure strategy equilibrium (Theorem 3).

1.5. Interpretation of continuum-player games

The properties illustrated in the example above hold quite generally for discrete large games, but do not hold in continuum player-games. One goal of this paper is to understand why.

The short answer is that the mathematical structure of the continuum suppresses patterns of behavior that are meaningful in large finite-player games. To explain this, recall our claim that the set of all strategy profiles in a continuum-player game $\bar{\Gamma}$ can be embedded as a subset of strategy profiles in its discrete large game representation Γ. The essential point is that *this embedding is strict* (Theorem 8), so Γ has strictly richer set of strategy profiles than $\bar{\Gamma}$. Fig. 1 illustrates this.

To see the implications, consider again the symmetric equilibrium $\bar{\mu}$ in the mis-matching game $\bar{\Gamma}_m$ and the corresponding behaviorally equivalent profile μ in Γ_m. The two profiles de-
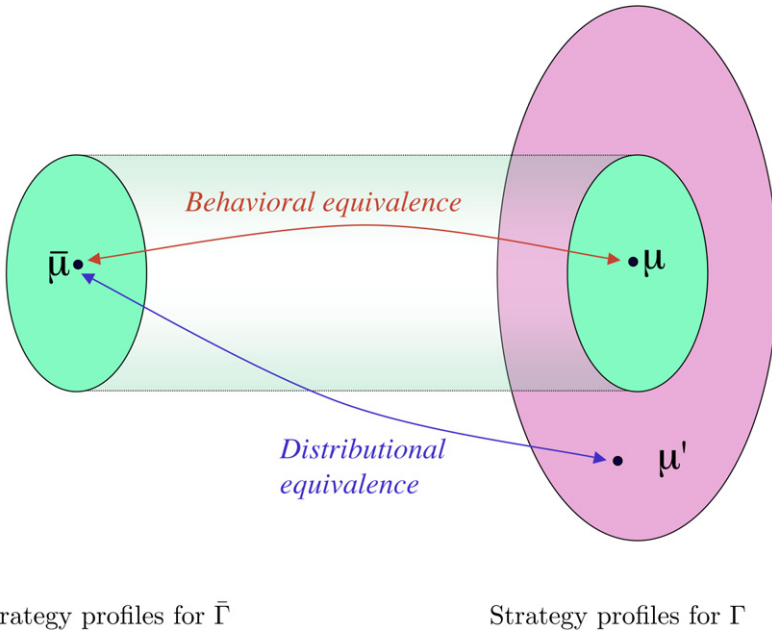
Fig. 1. Comparison of spaces of strategy profiles.

scribe identical "player-by-player" choice of strategy.[5] But since the embedding is strict, there is a strategy profile μ' with no behaviorally equivalent counterpart in $\bar{\Gamma}_m$. Theorem 8 shows that any such μ' is *distributionally equivalent* to a unique profile in $\bar{\Gamma}_m$ which we assume, for the purpose of this informal discussion, to be $\bar{\mu}$.

Distributional equivalence is a much weaker criterion, implying only:

$$\int_{[a,b]} \mu' \, d\lambda = \int_{[a,b]} \bar{\mu} \, d\bar{\lambda}, \quad \text{for every interval of players } [a,b]. \tag{2}$$

This criterion preserves *aggregate* players' behavior, but not what individual players actually do.

How would one generate profiles like μ' and what would they look like? An obvious method to generate examples is by taking "typical" draws from the set of pure strategies, i.e., a pure strategy profile μ' satisfying

$$\int_{[a,b]} \mu' \, d\lambda = 0.5 \, |b - a|.$$

As far as *aggregate* behavior is concerned, it is not possible to distinguish between the behaviors implied by μ', $\bar{\mu}$ or μ.[6] On the other hand, μ' and $\bar{\mu}$ imply radically different player-by-player behavior: Under $\bar{\mu}$ all players take the same randomized action, while in μ' no one randomizes. In fact, μ' is a *pure strategy equilibrium*.

It is well known that no pure strategy profile in $\bar{\Gamma}_m$ can be distributionally equivalent to $\bar{\mu}$ in the sense of Eq. (2). The thesis of this paper is that difficulties arising in modeling strategy

---

[5] This is a delicate issue since Γ and $\bar{\Gamma}$ consist of different set of players. Definition 10 ensures that this can be done.

[6] Note that μ must be distributionally equivalent to both $\bar{\mu}$ and μ'.

profiles, mixing, or purification in continuum-player games are manifestations of the phenomenon that these games have deficient strategy spaces. Informally, discrete large games supply the missing strategy profiles.

1.6. Summary and plan

If we take models literally, then *all* of the models presented in this paper are wrong: Sets of players are neither a continuum, a finitely additive probability space, nor even a finite set of arbitrary cardinality. The point of the various models is not to be correct in a naive descriptive sense, but to provide better understanding of strategic interactions in concrete situations with "many" players.

In light of this, the contribution of this paper may be viewed in two ways:

- Discrete large games provide a tractable and self-contained model of strategic interactions in large populations; or
- Discrete large games provide a foundation to justify and interpret commonly used shortcuts, heuristics and loose intuitions in continuum models.

Among the long-standing issues that can be simply and transparently understood with the help of discrete large games are purification, independent randomizations, distributional strategies and the definition of strategy profiles.

The paper proceeds as follows: The model of discrete large games is formally introduced in the next section and basic theorems on expected payoffs and purification are established. Section 3 proves a basic representation theorem for sequences of finite-player games by discrete large games and establishes the connection between their equilibria. Section 4 introduces continuum-player games and shows that its space of strategy profiles can be embedded into that of its discretization. This is followed by a discussion of distributional strategies and equilibria in these games. Section 5 concludes with a discussion of the literature.

2. Discrete large games

This section introduces the new model of large games. Finite-player games are introduced first to serve as a benchmark and motivation.

2.1. Players and attributes

We shall think of a player as a *vector of attributes*, representing either factors correlated with the player's payoff functions (e.g., age, sex, income level, neighborhood, etc.) or arbitrary labels to distinguish between different players.

There is a countably infinite number of possible attributes, and assume, for simplicity, that each attribute is binary. Then $\bar{T} \equiv \{0, 1\}^{\mathcal{N}}$ is the space of all vectors of attributes. By interpreting each element of $\bar{T}$ as a binary expansion of a real number, this model covers the usual case in which the space of players is taken as the interval $[0, 1]$. Taking $\bar{T}$ as the space of players, as we do here, is both more general and convenient.[7]

---

[7] The interval $[0, 1]$ as index set of players is endowed with a particular metric, while no specific metric structure is introduced on $\bar{T}$.

Let $\mathcal{B}_k$ denote the algebra of sets in $\bar{T}$ generated by the first k attributes. Then $\mathcal{B} \equiv \bigcup_{k=1}^{\infty} \mathcal{B}_k$ is an algebra that consists of all sets of players that can be defined in terms of a finite set of attributes. A function $f : \bar{T} \to \mathcal{R}^L$ is *simple* if it is finite-valued and measurable with respect to $\mathcal{B}$.

2.2. Finite-player games

Let T_N be a finite subset of $\bar{T}$, interpreted as a set of players with cardinality denoted $\#T_N$. A player chooses an action in a finite set A common to all players and fixed for all the games considered in this paper. Let Δ denote the set of probability distributions on A (mixed strategies) and identify each $a \in A$ with the probability distribution that puts unit mass on a.

The frequency of a set $C \subset \bar{T}$ relative to T_N is

$$\lambda_N(C) \equiv \frac{\#(C \cap T_N)}{\#T_N}.$$

Note that we can (and will) view λ_N as a distribution on $\bar{T}$ with support T_N.

Although this may appear a little pedantic at first, we shall use integral notation (rather than sums) in finite-player games to simplify comparison with infinite-player games where integrals are unavoidable. Thus, for any function $f : T_N \to \mathcal{R}^L$, we shall write[8]:

$$\int_{\bar{T}} f \, d\lambda_N \equiv \frac{1}{\#T_N} \sum_{t \in T_N} f(t).$$

Definition 1. A *finite-player game (with anonymous interaction)*, denoted $\Gamma_N = (T_N, u_N)$, is a game where[9]

1. The set of players is T_N;
2. A pure strategy profile is any function $s : T_N \to A$;
3. Player t's payoff at a pure strategy profile s is:

$$U_N(t, s) \equiv u_N\left(t, s(t), \int s \, d\lambda_N\right), \tag{3}$$

   where $u_N : T_N \times A \times \Delta \to \mathcal{R}$ is a continuous function.[10] To simplify notation, we shall write $u_N(t, s(t), s)$ instead of $u_N(t, s(t), \int s \, d\lambda_N)$.

Anonymity of interaction is reflected in the fact that player's payoff depends only on the action he takes and the average play $\int s \, d\lambda_N$.

2.3. Discrete large games: definition

The definition of discrete large games closely parallels that of finite-player games:

---

[8] To simplify notation, we will often drop reference to the space on which we are integrating since all measures used are defined on all of $\bar{T}$. The reference is restored if ambiguity is likely.

[9] The set of actions A is not included since it is fixed throughout the paper.

[10] Obviously, only continuity in Δ is relevant here.

Definition 2. A *discrete large game*, denoted $\Gamma \equiv (T, \lambda, u)$, is a game where:

1. The space of players is $(T, 2^T, \lambda)$ where $T \subset \bar{T}$ is countable and λ is a finitely additive probability measure on the set of all subsets of T, 2^T,[11] such that:
   (a) $\lambda(t) = 0$ for every $t \in T$ and
   (b) $\lambda(B) > 0$ for every non-empty $B \in \mathcal{B}$.
2. A pure strategy profile is any function $s : T \to A$.
3. Player t's payoff at a pure strategy profile s is:

$$U(t, s) \equiv u\left(t, s(t), \int s \, d\lambda\right) \tag{4}$$

   where $u : T \times A \times \Delta \to \mathcal{R}$ is uniformly continuous.[12] To simplify notation, we shall write $u(t, s(t), s)$ instead of $u(t, s(t), \int s \, d\lambda)$.

Condition 1(a) captures the idea that each player becomes negligible as the number of players increases, so in the limiting model Γ every player has zero weight. This implies that λ cannot be countably additive, because $\lambda(t) = 0$ for each t and T is countable, yet $\lambda(T) = 1$.[13] Condition 1(b) simplifies the exposition.

Condition 2 says that no set of players or strategy profiles is a priori ruled out. In particular, there is no implied coordination of actions since extraneous measurability considerations are superfluous. Finally, uniform continuity is needed in condition 3 because T is not compact.

Notation: We shall view λ as a distribution on $\bar{T}$ that puts unit mass on the countable set T. Also, we will not notationally distinguish between $B \in \mathcal{B}$ and $B \cap T$ or between a simple $f : \bar{T} \to \mathcal{R}$ and its restriction to T since there is a one-one, onto mapping between such functions.[14] In particular, we may view $\mathcal{B}$ as an algebra on T.

2.3.1. The finitely additive integral

As before, $\int_T s \, d\lambda$ represents the aggregate play, except that the integral now is with respect to the finitely additive probability λ.

The reader taken aback by the use of finitely additive probabilities should bear in mind that the basic notion (and much of the theory) of integration is the *same* for finitely and countably additive measures. They are based on the same concepts and definitions and the basic properties of the resulting integral hold irrespectively of whether the underlying measure is countably additive or not.[15] The key difference appears in connection with limit theorems, which may break down for *arbitrary* finitely additive measures.[16] Proposition A.2 shows that $(T, 2^T, \lambda)$ is particularly well-behaved in this respect.

---

[11] Existence of such games is shown as a consequence of Theorem 7.

[12] Here, uniform continuity is defined with respect to any metric for the product topology on $\bar{T}$. Uniform continuity can be introduced in a metric-free fashion, as shown in Lemma A.4.

[13] In fact, λ is *purely* finitely additive, as defined in Bhaskara Rao and Bhaskara Rao (1983, p. 240): If η is any *countably* additive measure on $(T, 2^T)$ such that $\eta(A) \leqslant \lambda(A)$ for every $A \subset T$, then $\eta(A) = 0$ for every $A \subset T$. The fact that λ is purely finitely additive follows from Theorem 10.3.2 in Bhaskara Rao and Bhaskara Rao.

[14] Given any simple function $\bar{f}$ on $\bar{T}$, its restriction to T is a (uniquely defined) simple function f. Conversely, since T is dense in $\bar{T}$, any simple function f on T has a unique extension to all of $\bar{T}$.

[15] Textbooks like Dunford and Schwartz (1958) develop the theory of integration for finitely additive measures first, then introduce countable additivity as a special case.

[16] E.g. the Radon–Nikodym Theorem, completeness of L^p spaces etc. fail for general finitely additive measure spaces.

For completeness, I provide a brief account of integration for $(T, 2^T, \lambda)$. Define the integral of a function $f : T \to \mathcal{R}^L$ with finite range $\{x_1, \ldots, x_K\}$ by

$$\int_T f \, d\lambda \equiv \sum_{k=1}^{K} x_k \lambda \big(f^{-1}(x_k) \big). \tag{5}$$

That is, the integral is just the average of its values weighted by their frequencies. Note that this is well-defined for *any* finite-valued function, because every subset of T is measurable. For a sequence $\{f_n\}$ of functions, we write $f_n \to f$ if for every $\epsilon > 0$

$$\lim_{n \to \infty} \lambda \big\{ t \in T : \big| f_n(t) - f(t) \big| > \epsilon \big\} = 0.^{17}$$

This is the analogue of 'convergence in measure' in the standard theory.

Definition 3. The integral of an arbitrary function $f : T \to \mathcal{R}^L$ is

$$\int_T f \, d\lambda = \lim_{n \to \infty} \int_T f_n \, d\lambda,$$

where $f_n : T \to \mathcal{R}^L$, $n = 1, 2, \ldots$, is any sequence of simple functions such that $f_n \to f$ and $\lim_{n,m \to \infty} \int |f_n - f_m| \, d\lambda = 0$.

Bhaskara Rao and Bhaskara Rao (1983) show that this definition does not depend on the particular approximating sequence $\{f_n\}$ (Proposition 4.4.10), that all of the familiar properties of integration continue to hold (Theorem 4.4.13), and that any bounded function is integrable (Theorem 4.4.18).[18] Lemma A.2 in Appendix A provides a method for calculating the integral $\int f \, d\lambda$ of a bounded function f in terms of its average in finite games, $\int f \, d\lambda_N$. Finally, Proposition A.2 shows that, under a mild condition, the function space $L^1(T, 2^T, \lambda)$ is complete, hence a Banach space.

2.4. The law of large numbers

Mixed strategy profiles are defined in exactly the same way as in finite-player games, namely as *any* function $\mu : T \to \Delta$, with $\mu(t)$ denoting player t's mixed action. Take as state space the set of all pure strategy profiles, $S \equiv A^T$ with the σ-algebra $\mathcal{S}$ generated by all events of the form $\{s \in S : s(t) = a\}$, $t \in T$, $a \in A$. Given that players randomize independently, the profile μ determines, using standard arguments, a unique *countably* additive probability distribution P^μ on $\mathcal{S}$. Note that P^μ is determined entirely by the profile μ and the independence assumption; P^μ has otherwise nothing to do with λ.

In discrete large games a strong version of the law of large numbers holds:

---

[17] One may take $|\cdot|$ to be any of the usual equivalent norms on $\mathcal{R}^L$. Bhaskara Rao and Bhaskara Rao (1983, Definition 4.3.1) refer to this as *hazy convergence*. Note that their definition is simplified here because λ is positive and 2^T is the power set, so there is no need to take absolute values or outer measures.

[18] To apply Theorem 4.4.18, note that any bounded function is dominated by a constant (hence, integrable) function and all functions are measurable because the σ-algebra used is the power set.

Theorem 1 *(Law of Large Numbers). For every strategy profile μ*

$$P^\mu \left\{ s \in S : \int_B s \, d\lambda = \int_B \mu \, d\lambda, \ \forall B \in \mathcal{B} \right\} = 1. \tag{6}$$

That is, almost every sample realization s has the property that its integral $\int_B s \, d\lambda$ is equal to the population mean $\int_B \mu \, d\lambda$ on all $B \in \mathcal{B}$. The proof relies on the usual law of large numbers for sequences of independent random variables. Note that, as in finite-player games, no assumptions on μ are imposed. Profiles are quite arbitrary; players need not randomize identically or even with probabilities that depend continuously on their labels. In addition, the law of large numbers holds on all sets $B \in \mathcal{B}$ simultaneously.

2.5. Expected payoffs

In finite-player games, a mixed strategy profile is any function $\mu_N : T_N \to \Delta$. Such profile uniquely defines a probability distribution P^{μ_N} on the set of pure profiles $S_N \equiv A^{T_N}$ that can be used to compute expected payoffs:

$$U_N(t, \mu_N) \equiv \int_{S_N} u_N\big(t, s(t), s\big) \, dP^{\mu_N}. \tag{7}$$

The expected payoff of player t at a profile μ can now be defined in a manner identical to the one used in finite games (see Eq. (7)):

$$U(t, \mu) \equiv \int_S u\big(t, s(t), s\big) \, dP^\mu. \tag{8}$$

That is, we compute the payoff $u(t, s(t), s)$ at a given pure profile s, then average these payoffs out using P^μ.

Theorem 2 *(Expected Payoffs). The expected payoff $U(t, \mu)$ in Eq. (8) is well-defined for any t and μ and satisfies*:

$$U(t, \mu) = \sum_{a \in A} \mu(t)(a) u\left(t, s(t), \int \mu \, d\lambda\right).^{19} \tag{9}$$

In the sequel, we shall write $u(t, a, \mu)$ instead of $u(t, s(t), \int \mu \, d\lambda)$ for notational simplicity.

The theorem shows that the law of large numbers implies that the outer integral in Eq. (8) can be eliminated, leaving a non-stochastic expression. For instance, if player t chooses a pure action a, Eq. (8) simply becomes:

$$U(t, \mu) \equiv u(t, a, \mu).$$

---

[19] $\mu(t)(a)$ is the probability that player t chooses a.

2.6. Equilibrium

Given a strategy profile μ, we define ϵ-best responses as usual by:

$$\text{Br}_\epsilon(t, \mu) \equiv \big\{\alpha \in \Delta\colon \ U(t, \alpha, \mu) \geqslant U(t, \alpha', \mu) - \epsilon, \ \forall \alpha' \in \Delta\big\}. \tag{10}$$

We use $\text{Br}(t, \mu)$ to denote the set of exact best responses.

Definition 4. Fix a discrete large game Γ. A strategy profile μ is:

1. An ϵ-*equilibrium*, $\epsilon \geqslant 0$, if:

$$\lambda\big\{t \in T\colon \ \mu(t) \in \text{Br}_\epsilon(t, \mu)\big\} \geqslant 1 - \epsilon. \tag{11}$$

2. An *equilibrium* if:

$$\lambda\big\{t \in T\colon \ \mu(t) \in \text{Br}_\epsilon(t, \mu)\big\} > 1 - \epsilon, \quad \text{for every } \epsilon > 0. \tag{12}$$

3. An *exact equilibrium* if

$$\lambda\big\{t \in T\colon \ \mu(t) \in \text{Br}(t, \mu)\big\} = 1. \tag{13}$$

The notions of ϵ- and exact equilibria are the obvious large game counterparts of the corresponding notions in finite-player games. What is new is the notion of equilibrium, which requires a profile to be ϵ-equilibrium for every $\epsilon > 0$. The reader may find this unusual because when the number of players is either finite or the continuum, any equilibrium is necessarily exact.

The following example illustrates the different concepts:

Example 2. The set of players is T with atomless distribution λ. Each player chooses among two actions $\{U, D\}$. The payoff of a player depends only on his own action (so there is no strategic interaction), with $u(t, U) = 1$ and $u(t, D) = 0$. Consider a profile μ, where the probability of playing U is:

$$\mu(t) = 1 - \frac{1}{N}, \quad \text{for } t \in T_N - T_{N-1}, \ N \geqslant 2. \tag{14}$$

The profile μ is an equilibrium that is not exact. To see this, let

$$E_k \equiv \big\{t \in T\colon \ \mu(t) \in \text{Br}_{1/k}(t, \mu)\big\}$$

denote the set of players who cannot gain more than $\frac{1}{k}$ by deviating from μ. Since E_k includes all but at most finitely many players, we have $\lambda(E_k) = 1$ for every k. This means that μ is indeed an equilibrium, but it is not exact, since no single player exactly best responds. What is true, of course, is that almost every player is arbitrarily close to best responding. In fact, players' deviations from optimizing have no aggregate effects in the sense that the frequency of playing U in any typical realization is 1.[20]

To understand the source of this phenomenon, consider the set $\bigcap_k E_k$. This is the set of players who best respond to μ exactly. Saying that the profile μ is an exact equilibrium is equivalent to

---

[20] This is a direct consequence of the law of large numbers in Eq. (6). Note that this is exactly what would happen in a standard application of the law of large numbers to independent coin flips with the probability of heads decreasing to zero. All that is needed is that the conditions in, for instance, Shiryayev (1984, Theorem 2, p. 364) are satisfied.

saying that $\lambda(\bigcap_k E_k) = 1$. If λ were countable additivity, then $\lambda(\bigcap_k E_k) = \lim \lambda(E_k) = 1$, which is clearly not true here since $\bigcap_k E_k$ is empty. This explains why the equilibrium concepts in (12) and (13) coincide with a continuum of players but not in a discrete large game.

So what is the right equilibrium notion? I adopt the criterion that the properties of an abstract model of large games should be evaluated in terms of how well they reflect those of large finite-player games. Theorems in support of this claim are supplied in Section 3.

Finally, the reader should also note that theorems asserting properties for all equilibria hold a fortiori for exact equilibria. This includes, for instance, the Purification Theorem 3 and Theorem 5, among others.

2.7. Purification

We first need the following notions of equivalence of strategy profiles:

Definition 5. Two profiles, μ, μ' in Γ are:

1. *distributionally equivalent* if for every $B \in \mathcal{B}$

$$\int_B \mu \, d\lambda = \int_B \mu' \, d\lambda; \tag{15}$$

2. *behaviorally equivalent* if for every $B \in \mathcal{B}$

$$\int_B |\mu - \mu'| \, d\lambda = 0.^{21}$$

That is, two distributionally equivalent profiles represent identical *average* play for any set of players that can be defined in terms of a finite set of attributes. If, in addition, they satisfy the stronger property of being behaviorally equivalent then, roughly, they also describe the same behavior *player-by-player*, and not just on average.

We can now state a strong and natural purification result:

Theorem 3 *(Purification Theorem). For every equilibrium μ of Γ*

$$P^\mu \left\{ s \in S: \begin{array}{l} s \text{ is a pure strategy equilibrium} \\ \text{distributionally equivalent to } \mu \end{array} \right\} = 1. \tag{16}$$

If μ is an exact equilibrium, then the assertion in Eq. (16) can be strengthened to: "*s is an exact pure strategy equilibrium distributionally equivalent to μ.*"

Thus, every realization of a mixed strategy equilibrium μ is a pure strategy equilibrium, P^μ-almost surely. There is simply no way to distinguish between the mixed profile μ and any of its typical realizations s if all one has is information about aggregate play, in the sense that:

$$\int_B (\mu - s) \, d\lambda = 0, \quad \forall B \in \mathcal{B}.$$

---

[21] This says that μ and μ' differ by a λ-null function, which is the standard notion of equivalence in L^1-norm. For finitely additive probabilities, this is strictly weaker than saying that μ and μ' agree with λ-probability 1. See Bhaskara Rao and Bhaskara Rao (1983, Definition 4.2.4) and subsequent discussion.

On the other hand, s and μ will typically be behaviorally distinct. For example, if under μ every player randomizes 50/50 between two actions a and b, then $\mu(t) - s(t)$ is never zero since $s(t) \in \{a, b\}$.[22]

In fact,

$$\int_T |\mu - s| \, d\lambda = \frac{1}{2}\lambda\big(\{t: s(t) = a\}\big) + \frac{1}{2}\lambda\big(\{t: s(t) = b\}\big) = \frac{1}{2}.$$

Theorem 8, part 3, shows that the situation depicted in this example is quite general.

In the context of finite-player games, Kalai (2004) showed, among other things, that with high probability, realizations of mixed strategy equilibria are ϵ-equilibria when the number of players is large. He refers to this property as self-purification. As we shall discuss later, a version of this property cannot hold for continuum-player games. On the other hand, every equilibrium in a discrete large game can, in Kalai's terminology, be exactly self-purified almost surely.

3. Limit theorems

At this point the reader should be skeptical whether the abstract concept of discrete large games is anything more than a technical artifact. In this and the next section I will argue that they represent the "right" idealization of large finite-player games.

3.1. Representation theorem

First we formally define what it means for a discrete large game to represent a "well-behaved" sequence of large finite-player games[23]:

Definition 6. $\{\Gamma_N\}_{N=1}^{\infty}$ is a *proper sequence of finite-player games* if

1. $T_N \subset T_{N+1}$ for every N;
2. The function

   $$u: \bigcup_{N=1}^{\infty} T_N \times A \times \Delta \to \mathcal{R}$$

   given by $u(t, \cdot, \cdot) = u_N(t, \cdot, \cdot)$ for any N with $t \in T_N$ is well-defined and uniformly continuous;

3. $$\frac{\#T_N - \#T_{N-1}}{\#T_N} \to 1, \quad \text{as } N \to \infty; \tag{17}$$
4. For every non-empty $B \in \mathcal{B}$, $\lim_{N \to \infty} \lambda_N(B)$ exists and is strictly positive;
5. For every $t \in \bar{T}$, $\lim_{N \to \infty} \lambda_N(t) = 0$.

Conditions 1, 2, 4 and 5 are obvious and they are needed, in one form or another, for $\{\Gamma_N\}_{N=1}^{\infty}$ to be representable by a limiting large game model. Condition 3 says that the number of players along the sequence of finite-player games grows rapidly. It is an innocuous requirement in that

---

[22] As standard, we identify each pure action with the mixed action in Δ that assigns to it unit mass.

[23] Lemma A.6 in Appendix A shows that this definition is not vacuous.

one can always extract a subsequence in which it is satisfied. But it is a requirement with surprising and far-reaching technical implications (Proposition A.2) and guarantees, among other things, strategic embedding (Theorem 8).

Definition 7. A discrete large game $\Gamma = (T, \lambda, u)$ *represents* a sequence of finite-player games $\{\Gamma_N\}_{N=1}^{\infty}$ if $T = \bigcup_{N=1}^{\infty} T_N$ and

1. $u_N(t, \cdot, \cdot) = u(t, \cdot, \cdot)$ for every N and $t \in T_N$;
2. For every finite collection of subsets $A_1, \ldots, A_L \subset T$ there is a subsequence $\{N_k\}$ such that

$$\lambda(A_l) = \lim_{k \to \infty} \lambda_{N_k}(A_l), \quad \text{for } l = 1, \ldots, L. \tag{18}$$

The next theorem is offered here for completeness and without proof since it is subsumed by Theorem 7:

Theorem 4. *Every proper sequence of finite-player games can be represented by a discrete large game.*

3.2. Limits of equilibria

Our first limit theorem says that Γ does not overlook any strategic phenomena arising as approximate equilibria in large finite-player games:

Theorem 5. *Suppose that Γ represents a proper sequence of finite-player games $\{\Gamma_N\}_{N=1}^{\infty}$. If, for each N, μ_N is an ϵ_N-equilibrium for Γ_N with $\epsilon_N \downarrow 0$, then there is an equilibrium μ for Γ and a subsequence $\{N_k\}$ such that $\mu_{N_k}(t) \to \mu(t)$ for every $t \in T$.*

To illustrate this theorem, consider a finite-player version of the game in Example 2. Let μ_N denote the strategy profile obtained by applying Eq. (14) to players in T_N. Starting from the sequence $\{\mu_N\}$, the construction of Theorem 5 gives rise to the profile μ in Example 2. The profile μ summarizes behavior along a sequence of approximate equilibria in finite-player games. As such, the weaker notion of equilibrium used is sensible, indeed necessary. Only by appealing to additional principles, such as countable additivity in a continuum-player game, can one obtain the stronger conclusion that the limit of the sequence $\{\mu_N\}$ is exact.

We finally note that this theorem also establishes the existence of equilibrium for those discrete large games that arise as limits of proper sequences of finite-player games.

3.3. Equilibria in the limit

The next theorem shows that equilibria in Γ arise as limits of ϵ-equilibria in large finite-player games:

Theorem 6. *Suppose that μ is an equilibrium for Γ and let μ_N denote its restriction to T_N. Then for every $\epsilon > 0$, there is N large enough such that μ_N is an ϵ-equilibrium for Γ_N.*

That is, the discrete large game Γ does not contain any equilibrium that arises as an artifact of the infinite model.

4. What is the "right" large game model?

There is no such thing as the "right" model in any narrow descriptive sense. A more useful criterion is whether a model can help us understand strategic interactions in large finite populations. Here I compare discrete large games with the continuum model that dominates the literature.

4.1. Continuum-player games

Here we take the space of players to be $\bar{T}$. Let $\bar{\mathcal{B}}$ denote the σ-algebra generated by $\mathcal{B}$ and $\bar{\lambda}$ is a (countably additive) probability measure on $(\bar{T}, \bar{\mathcal{B}})$.

Definition 8. A *continuum-player game* $\bar{\Gamma} \equiv (\bar{T}, \bar{\lambda}, \bar{u})$ is one where:

1. The space of players is $(\bar{T}, \bar{\mathcal{B}}, \bar{\lambda})$ with:
    1.1. $\bar{\lambda}(t) = 0$ for every $t \in \bar{T}$ and
    1.2. $\bar{\lambda}(B) > 0$ for every $B \in \mathcal{B}$.
2. A pure strategy profile is any measurable function $s : \bar{T} \to A$.
3. Player t's payoff at a pure strategy profile s is:

$$\bar{U}(t, s) \equiv \bar{u}\left(t, s(t), \int s \, d\bar{\lambda}\right) \tag{19}$$

where $\bar{u} : \bar{T} \times A \times \Delta \to \mathcal{R}$ is a continuous function. To simplify notation, we shall write $\bar{u}(t, s(t), s)$ instead of $\bar{u}(t, s(t), \int s \, d\bar{\lambda})$.

Most aspects of this definition are similar to those found in the definition of discrete large games, and thus have similar interpretation. The key difference is in the second requirement: in a continuum-player game some (in fact, many) strategy profiles are ruled out as "illegitimate." There is no such restriction when the number of players is either finite or discrete.

As before, we are interested in interpreting $\bar{\Gamma}$ as representing strategic situation in large finite-player games:

Definition 9. A continuum-player game $\bar{\Gamma}$ *represents* a proper sequence of finite-player games $\{\Gamma_N\}_{N=1}^{\infty}$ if

1. $u_N(t, \cdot, \cdot) = \bar{u}(t, \cdot, \cdot)$ for every N and $t \in T_N$;
2. For every subset $B \in \mathcal{B}$

$$\bar{\lambda}(B) = \lim_{N \to \infty} \lambda_N(B). \tag{20}$$

It is easy to see that condition (20) uniquely defines $\bar{\lambda}$: If $\bar{\lambda}'$ is any countably additive measure on $\bar{T}$ satisfying (20), then $\bar{\lambda}$ and $\bar{\lambda}'$ agree on all cylinder sets and must therefore coincide. Note that this is not true for discrete large game where measures are finitely additive, hence the stronger requirement (18).

Theorem 7 *(Representation Theorem). Every proper sequence of finite-player games has a discrete large game representation $\Gamma = (T, \lambda, u)$ and a continuum-player game representation $\bar{\Gamma} = (\bar{T}, \bar{\lambda}, \bar{u})$.*[24] *Any two such representations must satisfy:*

1. $\lim_{N \to \infty} \lambda_N(B) = \lambda(B) = \bar{\lambda}(B)$, *for every $B \in \mathcal{B}$;*[25]
2. $\bar{u}$ *and u agree on T.*

In this case we say that Γ discretizes $\bar{\Gamma}$.

The remainder of this section compares the two representations Γ and $\bar{\Gamma}$.

4.2. Embedding

We show that the continuum representation $\bar{\Gamma}$ of a given proper sequence of finite-player games can be "strategically embedded" into the discrete representation Γ. Loosely, this means that any pattern of strategic behavior arising in $\bar{\Gamma}$ has a corresponding pattern in Γ.

First we need to introduce notions of equivalence of strategy profiles in Γ and $\bar{\Gamma}$. Distributional equivalence is straightforward, but behavioral equivalence is not. The reason is that strategy profiles in $\bar{\Gamma}$ and Γ are defined on different spaces of players, so the formal definition proceeds indirectly.

Definition 10. A strategy profile μ in Γ and $\bar{\mu}$ in $\bar{\Gamma}$ are:

1. *distributionally equivalent* if for every $B \in \mathcal{B}$

$$\int_B \mu \, d\lambda = \int_B \bar{\mu} \, d\bar{\lambda}; \tag{21}$$

2. *behaviorally equivalent* if for every $\epsilon > 0$ there is a simple profile μ_ϵ[26] such that

$$\int_T |\mu - \mu_\epsilon| \, d\lambda < \epsilon \quad \text{and} \quad \int_{\bar{T}} |\bar{\mu} - \mu_\epsilon| \, d\bar{\lambda} < \epsilon. \tag{22}$$

It is routine to verify that the notions of behavioral and distributional equivalence coincide when comparing profiles in $\bar{\Gamma}$.[27] Not so in discrete large games: as the next theorem shows, there

---

[24] We should also note that if $\{\Gamma_N\}$ can be at all represented by a continuum-player game $\bar{\Gamma}$, then this representation is unique. If $\bar{\Gamma}' = (\bar{T}, \bar{\lambda}', \bar{u}')$ is another representation, then by condition 2 the two measures $\bar{\lambda}$ and $\bar{\lambda}'$ agree on an algebra that generates $\bar{\mathcal{B}}$ and so they must be equal. Also $\bar{u}$ and $\bar{u}'$ are two uniformly continuous functions that agree on a dense subset $\cup T_N$, so they must be equal.

[25] This implies that for any simple function f, $\int_{\bar{T}} f \, d\bar{\lambda} = \int_T f \, d\lambda$. Recall from Footnote 14 that there is a bijection between simple functions on T and $\bar{T}$, and thus we do not distinguish between them notationally.

[26] Recall that we do not distinguish between simple profiles in $\bar{\Gamma}$ and Γ since there is a one-one, onto mapping between simple strategies in the two games. See Footnote 14.

[27] To make this precise, two profiles $\bar{\mu}$ and $\bar{\mu}'$ in $\bar{\Gamma}$ are distributionally equivalent if:

$$\int_B \bar{\mu} \, d\bar{\lambda} = \int_B \bar{\mu}' \, d\bar{\lambda}.$$

are many distributionally equivalent strategy profiles that correspond to behaviorally distinct individual play.[28]

First some definitions: Two strategy profiles μ, μ' in Γ are (L^1-) *equivalent* if $\int |\mu - \mu'|\, d\lambda = 0$, and they are *(essentially) distinct* otherwise.[29] Every time we qualify a conclusion by "essentially" we mean "up to equivalence" in the above sense. In particular, μ is *essentially pure* if there is $s : T \to A$ such that:

$$\int |s - \mu|\, d\lambda = 0.$$

Theorem 8. *Suppose that Γ and $\bar{\Gamma}$ are, respectively, a discrete and continuum representation of a proper sequence of finite-player games. Then*:

1. *Every strategy profile $\bar{\mu}$ in $\bar{\Gamma}$ has an essentially unique behaviorally equivalent profile μ in Γ.[30]*
2. *Every strategy profile μ in Γ has an essentially unique distributionally equivalent strategy profile $\bar{\mu}$ in $\bar{\Gamma}$.*
3. *Every strategy profile μ in Γ that is not essentially pure has uncountably many distinct distributionally equivalent profiles in Γ.*

The conclusion of the theorem is graphically depicted in Fig. 1. Part 1 roughly says that there is no strategic pattern of behavior that arises in a continuum game that cannot be captured by Γ. Part 2 says that any profile in Γ can be "averaged out" to a profile in $\bar{\Gamma}$. Part 3 reflects the phenomenon illustrated in the figure: many distinct profiles in Γ may map into a single profile in $\bar{\Gamma}$ and are thus indistinguishable in that game. For example, the profile in which each player randomizes 50/50 over two actions is distributionally indistinguishable from any of its pure strategy realizations. Continuum-player games simply miss many "legitimate" strategy profiles and equilibria that arise as limits of corresponding profiles and equilibria in finite-player games.

Technically, part 1 of the theorem hinges on the completeness of the space of functions defined on the finitely additive probability space $(T, 2^T, \lambda)$. See Proposition A.2; Al-Najjar (2007) contains a more comprehensive treatment. To get a rough intuition, suppose we are given a sequence of simple strategy profiles $\bar{\mu}_k$ that converges (in L^1, say) to some limit $\bar{\mu}$ in $\bar{\Gamma}$. Our goal is to find a profile μ to map $\bar{\mu}$ to. Because simple strategies transfer uniquely between Γ and $\bar{\Gamma}$, take the corresponding sequence of profiles μ_k in Γ. An obvious candidate is the limit of the

---

If so, the measures defined by the indefinite integrals $\int_B \bar{\mu}\, d\bar{\lambda}$ and $\int_B \bar{\mu}'\, d\bar{\lambda}$ coincide on $\mathcal{B}$. They are, furthermore countably additive on $\mathcal{B}$, so they have the same (and unique) countably additive extension to $\bar{\mathcal{B}}$. But then both $\bar{\mu}$ and $\bar{\mu}'$ are derivatives of this measure, and must therefore coincide $\bar{\lambda}$-a.e.

[28] Distributional equivalence in Γ can be defined similarly: two profiles μ and μ' are distributionally equivalent if:

$$\int_B \mu\, d\lambda = \int_B \mu'\, d\lambda.$$

[29] This is the usual notion of equivalence in L^1-norm, familiar in countably additive integration. There is a subtle difference, however. If two functions f, f' on T are equivalent then it easily follows that $\lambda\{t : |f - f'| > \epsilon\} = 0$ for every $\epsilon > 0$. If λ were countably additive, then this would imply that f and f' are equal λ-a.e.; i.e., $\lambda\{t : |f - f'| > 0\} = 0$. If λ is only finitely additive, then almost every equality is strictly stronger. Consider for example the function $f(t) = \frac{1}{N}$ for every $t \in T_{N+1} - T_N$, $N = 1, \ldots$. This function is strictly positive, yet it is in this sense equivalent to the function that is identically zero.

[30] That is, if μ' is any other behaviorally equivalent profile, then μ and μ' must be L^1-equivalent.

sequence of simple profiles $\{\mu_k\}$ if such limit exists. Now we know that $\{\bar{\mu}_k\}$ converges, so it must be Cauchy (in $L^1(\bar{T}, \bar{\lambda})$). From this, it immediately follows that $\{\mu_k\}$ must be Cauchy in $L^1(T, \lambda)$. The question now reduces to whether every Cauchy sequence in $L^1(T, \lambda)$ has a limit, which is precisely the statement that this function space is complete.

Our final theorem establishes the connection between equilibria of Γ and $\bar{\Gamma}$:

Theorem 9.

1. *If μ is an equilibrium of Γ then its unique distributionally equivalent profile $\bar{\mu}$ in $\bar{\Gamma}$ is also an equilibrium;*
2. *If the strategy profile $\bar{\mu}$ is an equilibrium of $\bar{\Gamma}$, then its essentially unique behaviorally equivalent profile μ in Γ is also an equilibrium.*

4.3. The problem of defining mixed strategy profiles

Expected payoffs in a finite-player game given a mixed profile μ_N can be naturally defined as in Eq. (7):

$$U_N(t, \mu_N) \equiv \int_{S_N} u_N\big(t, s(t), s\big) \, dP^{\mu_N}.$$

This is impossible in continuum-player games. To see this, consider a profile $\bar{\mu}$ in which each player plays each action with strictly positive probability. The analogue of the set of pure strategy profiles here is the set of *all* functions $\bar{S} \equiv A^{\bar{T}}$. Endow $\bar{S}$ with the smallest σ-algebra $\bar{\mathcal{S}}$ containing all events of the form $\{s \in \bar{S}: s(t) = a\}$, for $t \in \bar{T}$, $a \in A$. As in finite-player games, a mixed strategy profile $\bar{\mu}$ uniquely defines a probability distribution $P^{\bar{\mu}}$ on $\bar{S}$.[31]

While this may initially appear encouraging, the distribution $P^{\bar{\mu}}$ is all but useless in defining expected payoffs because the average population play "$\int s \, d\bar{\lambda}$" is not well-defined for a typical realization s yielding the meaningless expression:

$$\text{``}\bar{U}(t, \bar{\mu}) \equiv \int_{\bar{S}} \bar{u}(t, s(t), s) \, dP^{\bar{\mu}}.\text{''}$$

4.3.1. Distributional strategies

To get around this issue, Schmeidler (1973), Mas-Colell (1984), and Milgrom and Weber (1985) define player t's payoff at a profile $\bar{\mu}$ as:

$$\bar{U}(t, \bar{\mu}(t), \bar{\mu}) \equiv \bar{u}\big(t, \bar{\mu}(t), \bar{\mu}\big), \tag{23}$$

where it is understood that $\bar{u}$ depends on $\bar{\mu}$ only through its integral $\int \bar{\mu} \, d\bar{\lambda}$.

I will refer to this as the *distributional approach* and to $\bar{\mu}$ as a *distributional strategy*. Implicit in Eq. (23) is the intuition that players randomize independently and that some version of the law of large numbers holds. However, no explicit state space is provided, so this statement should be viewed as informal intuition taken on faith.

---

[31] This is a standard construction that relies on Kolmogorov Extension Theorem; see, for instance, Shiryayev (1984, p. 244).

The heuristic formula introduced in (23) above may be used to define equilibria. Thus, player t's ϵ-*best responses* consist of the set:

$$\text{Br}_\epsilon(t, \bar{\mu}) \equiv \big\{ \alpha \in \Delta \colon \bar{U}(t, \alpha, \bar{\mu}) \geqslant \bar{U}(t, \alpha', \bar{\mu}) - \epsilon, \ \forall \alpha' \in \Delta \big\}. \tag{24}$$

A strategy profile $\bar{\mu}$ is an *equilibrium (in distributional strategies)* if

$$\bar{\lambda}\big\{ t \in \bar{T} \colon \bar{\mu}(t) \in \text{Br}(t, \bar{\mu}) \big\} = 1. \tag{25}$$

4.3.2. A new interpretation of distributional strategies

The analysis so far shows that modeling mixed strategies, randomizations, and expected payoffs in discrete large games is straightforward, but quite problematic in the continuum. Here I show why this is the case.

Assume that Γ discretizes the continuum-player game $\bar{\Gamma}$ and μ is some strategy profile in Γ. Let $\bar{\mu}$ be its essentially unique distributionally equivalent strategy profile in $\bar{\Gamma}$ (guaranteed by Theorem 8). For concreteness, assume that in μ, hence in $\bar{\mu}$, almost every player randomizes. Let s be a typical realization, in the sense of Eq. (6), and note that s is a *pure* strategy profile in its own right.

From Eq. (6) and Theorem 8 we have:

$$\int_B s \, d\lambda = \int_B \mu \, d\lambda = \int_B \bar{\mu} \, d\bar{\lambda}, \quad \forall B \in \mathcal{B}.$$

Then the three profiles s, μ and $\bar{\mu}$ are all distributionally equivalent, reflecting identical *average* play on any set of players that can be defined in terms of a finite set of attributes. On the other hand, s and $\bar{\mu}$ are *not* behaviorally equivalent: there can be no sequence of simple profiles that can simultaneously approximate both a profile like $\bar{\mu}$, where almost every player randomizes, and a profile like s, where every player plays a pure strategy.

In summary, to each strategy profile $\bar{\mu}$ in a continuum-player game corresponds an essentially unique behaviorally equivalent profile μ but, typically, a continuum of behaviorally distinct, but distributionally equivalent profiles. A distributional strategy like $\bar{\mu}$ is a heuristic device to represent profiles, like s above, where behavior fluctuates too wildly to permit a behaviorally equivalent representation in the continuum. The cause of the difficulty is the measurability condition on strategy profiles in $\bar{\Gamma}$ which requires $\bar{\mu}$ to smooth over these individuals variations (while leaving the average intact).

This provides an interpretation of what the distributional strategy $\bar{\mu}$ is in effect doing: Since '$\int s \, d\bar{\lambda}$' cannot be meaningfully evaluated in the continuum, the distributional approach substitutes it by its 'plausible' value $\int \bar{\mu} \, d\bar{\lambda}$. This is based on the heuristic intuition that this is what one ought to obtain in any reasonable discrete version of the continuum. We made this intuition precise by identifying a unique behaviorally equivalent profile μ in the discrete large game, and computed expected aggregate play by taking expectations of $\int s \, d\lambda$ with respect to P^μ. This turns out to be equal to the value assumed by the distributional approach, $\int \bar{\mu} \, d\bar{\lambda}$. That is, the diagram in Fig. 2 commutes.

It should be noted that this is just an *interpretation*: The problem of finding a meaningful definition for the integral of sample realizations "$\int s \, d\bar{\lambda}$" is intractable. There is no resolution to the essential conflict between the freedom required by independent randomizations, and the straightjacket that measurability imposes. Rather, the claim is that one can think of someone using a continuum-player game as in effect using a shorthand for an underlying discrete large game in which integrals of realizations can be legitimately evaluated.
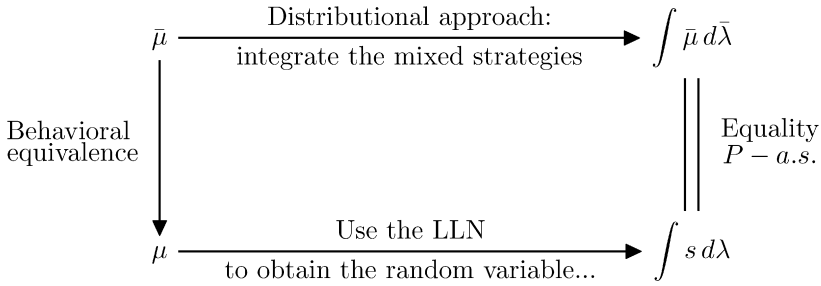
Fig. 2. Interpretation of distributional strategies.

4.3.3. Spurious randomness

Consider a discrete large game Γ and let s be one of its pure strategy profiles. From Theorem 8 there exists a distributionally equivalent strategy profile $\bar{\mu}$ in the corresponding continuum-player game $\bar{\Gamma}$. That is:

$$\int_B \bar{\mu} \, d\bar{\lambda} \equiv \int_B s \, d\lambda, \quad \forall B \in \mathcal{B}.$$

If $\bar{\mu}$ is not essentially pure, then its interpretation within $\bar{\Gamma}$ is ambiguous: Do we have in mind a setting where players randomize, one were players play pure strategies, or some combination of the two? The source of this ambiguity is the measurability requirement in the continuum, which forces a smoothing out of individual variations, variations that have clear strategic meaning in finite-player games. In cases where we are interested only in aggregate play, this smoothing may be harmless. But there are environments (e.g., ones involving private information, mechanism design, and the like) where individual-level behavior is important. For an example involving mechanism design, see Al-Najjar (2004).

In summary, randomness in behavior in $\bar{\Gamma}$ *may be* spurious, a mere reflection of the technical requirement of measurability rather than any relevant behavioral or strategic property. The discrete large games framework permits a much finer distinctions to be made.

5. Literature and concluding remarks

The scope of the theory of integration is the class of functions measurable with respect to countably additive probability spaces. In most economic and game theoretic applications, the underlying space is complete, separable and metric, and thus comes with an inherent similarity structure that plays a meaningful role in modeling primitives like payoffs, local interactions, etc.

Since any measurable function must be approximately continuous,[32] measurability implicitly imposes a great deal of similarity in players' behavior.[33] While similarity in primitives may be natural, similarity in behavior is an entirely different matter. Behavior is endogenous and, in non-cooperative strategic settings, requiring it to be similar across individuals is quite unnatural. As Dubey and Shapley (1994) note, measurability builds into the very definition of strategy profiles a degree of coordination among players that goes against the spirit of non-cooperative play.

---

[32] This is the well-known and standard Luzin's theorem (see Royden, 1988, p. 74).

[33] The same argument, of course, applies to modeling idiosyncratic utility, endowment, and informational states, and not just behavior.

This issue comes to a head when attempts are made to model idiosyncratic randomness, be it in types, randomized strategy choices, or random matching. Stochastically independent choices inevitably lead to action profiles that vary too wildly to be measurable, making it impossible to formulate a viable law of large numbers (Feldman and Gilles, 1985 and Judd, 1985). Al-Najjar (1995) and Uhlig (1996) presented attempts to sidestep the problem by appealing to vector-valued integration. Although the above-mentioned contributions provide convenient shortcuts around the problem, their value is limited when analysis is concerned with the properties of realizations (as in the case of purification, for instance).

The framework of this paper makes it meaningful to talk about players' similarity of attributes without requiring similarity in their action choices. Essentially, discrete large games develop formally the familiar intuition of using fine grids to represent continuous variables like prices, actions, players, etc. Al-Najjar (2004) develops some of these ideas, but relies on conditions that turn out to be unnecessarily restrictive and does not deal with strategic interaction.[34] Al-Najjar et al. (2006) use a similar idea to model undescribable events and complexity in a contracting context.

In game theory, one response was to use *distributional strategies*, as Schmeidler (1973), Mas-Colell (1984) and Milgrom and Weber (1985) have suggested. A distributional strategy, like the symmetric profile $\bar{\mu}$ in $\bar{\Gamma}_m$, is an assignment of lotteries to individual players. Interestingly, a similar idea, that of *relaxed controls*, is introduced in optimal control theory to deal with the problem of chattering controls in continuous time problems. See Kushner (1990).

Distributional strategies are convenient short-cuts to overcome the deficiency in strategy profiles inherent in continuum models. As our earlier discussions illustrated, this is a shortcut that comes at a price. Any strategy profile in a discrete large game corresponds to a unique distributional strategy in the corresponding continuum-player game (e.g., profiles μ' and $\bar{\mu}$ in Fig. 1). But much is lost in the translation, since: one has to be content with an aggregate, rather than player-by-player, description of what players do. The importance of what is lost is a methodological and practical question to be judged on a case by case basis.

The idea of using a countable set of players to model large games is not new. Guesnerie (1981, 1995), and Bewley (1986) proposed using a randomly drawn sample of agents from an underlying continuum game as a set of players. Feldman and Gilles (1985) suggested indexing agents by the integers with a purely finitely additive density charge.

Finitely additive probabilities have a long, but at times tortured, history. Following De Finetti, Savage (1954) and Dubins and Savage (1965) advocated the use of finitely additive measures in the foundation of decision theory and stochastic processes. These and other works make a normative, foundational case that requiring countable additivity introduces extraneous measurability and other technical restrictions that lack convincing behavioral justification.

Interest in finitely additive probabilities waned because they can be quite unruly, yielding well-known conceptual paradoxes and resisting essential mathematical tools. In this paper I use a class of finitely additive probabilities that is tractable and easy to interpret as discretizations of standard countably additive probabilities. Some of the mathematical machinery is introduced in Al-Najjar (2007), although without dealing with strategic interaction, distributional strategies, equilibria and the like.

---

[34] Specifically, in Al-Najjar (2004) attention is restricted to functions whose integral is the limit of their averages in large finite models. The problem is that the set of such functions is not closed under usual operations and thus intractable. The setup here, by contrasts, gives rise to a well-behaved and standard function space.

Finally, in a series of papers, Khan and Sun advocated the use of hyperfinite Loeb spaces to model large games. A representative paper is Khan and Sun (1999) where they model players as hyperreal numbers in a hyperfinite internal probability space. These spaces are founded on non-standard entities whose existence hinges on concepts like superstructures, $\aleph_0$-saturation, enlargement and principles like comprehensiveness, over- and under-flow, among many others. To understand, let alone evaluate, the nature of Khan and Sun's contribution would require the reader to be armed with a bewildering non-standard technical machinery. While this may be valuable in some areas of mathematics, the present paper shows that key issues in the theory of large games can be addressed in a simple and transparent way. Aside from a few exceptions, our arguments use sequences, limits, ordinary law of large numbers and the like.[35] A reader with little more than basic knowledge of analysis and probability theory should have no difficulty interpreting the model of discrete large games, verifying the correctness of the theorem, and making an informed judgment about the scope of the theory. More recently, Sun (2006) introduced a model where the space of players "can still be taken to be the unit interval [0, 1] with some atomless probability measure" (p. 34). He does not mention that his 'unit interval' is just a hyperfinite space whose elements have been arbitrarily relabeled by real numbers.[36] The new hyperfinite-relabeled-as-continuum space has little to do with the space [0, 1] other than cardinality; it preserves none of the usual structures we ordinarily associate with the interval [0, 1] such as ordering, metric and continuity concepts. Given the incidental nature of this relabeling, it would be more straightforward to acknowledge this for what it is, namely a hyperfinite space with the advantages and disadvantages these spaces entail.

Acknowledgments

Appendix A

A.1. Mathematical preliminaries

A.1.1. Measures and ultrafilters

We use ultrafilters on $\mathcal{N}$ to obtain a finitely additive extension of the density measure.[37] See Blass et al. (2001). First, we recall some standard definitions. Let $\mathcal{N}$ denote the natural numbers. A collection of subsets $\mathcal{U} \subset 2^{\mathcal{N}}$ is a *free ultrafilter* on $\mathcal{N}$ if:

1. $A \in \mathcal{U}$ and $A \subset B$ imply $B \in \mathcal{U}$;
2. $A,\ B \in \mathcal{U}$ implies $A \cap B \in \mathcal{U}$.
3. $\emptyset \notin \mathcal{U}$.

---

[35] The only non-elementary proofs (i.e., those requiring more knowledge than first-year graduate treatments at, say, the levels of Royden, 1988 and Billingsley, 1995) are those in Section A.1.1 which involve ultrafilters and embedding of L^p spaces. In general, one has to be on guard, throughout the paper, that the assumption of countable additivity does not surreptitiously sneak into the proofs.

[36] See the role of the relabeling function φ in the proof of his Proposition 5.6, p. 67.

[37] Another method is to use Banach limits; see Bhaskara Rao and Bhaskara Rao (1983, pp. 39–41).

4. No finite subset of $\mathcal{N}$ belongs to $\mathcal{U}$.
5. For every $A \subset \mathcal{N}$ either $A \in \mathcal{U}$ or $\mathcal{N} - A \in \mathcal{U}$.

Conditions 1 and 2 define a filter, while 3 requires that it is 'proper' (i.e., , $\mathcal{U} \neq 2^{\mathcal{N}}$). Condition 4 says that $\mathcal{U}$ is free (non-principal), while condition 5 says that it is maximal (i.e., $\mathcal{U}$ is an *ultra*filter).

Given a free ultrafilter $\mathcal{U}$, define the $\mathcal{U}$-limit of a sequence of real numbers $\{x_n\}$ by:

$$\mathcal{U} - \lim_{n \to \infty} x_n = x \iff \forall \epsilon > 0, \quad \{n: |x_n - x| < \epsilon\} \in \mathcal{U}. \tag{26}$$

Clearly, if $x_n \to x$ in the usual sense, then $x_n \xrightarrow{\mathcal{U}} x$. However, $\mathcal{U} - \lim x_n$ exists for any bounded sequence of real numbers, so $\mathcal{U}$-convergence is indeed a generalization of the usual concept.

Fix any enumeration $\{t_1, t_2, \ldots\}$ such that for every N, $T_N = \{t_1, \ldots, t_{\#T_N}\}$. Any set $A \subset T$ defines a sequence of real numbers:

$$\frac{\#(A \cap \{t_1, \ldots, t_n\})}{n}.$$

Define the *usual density* of A by

$$\lim_{n \to \infty} \frac{\#(A \cap \{t_1, \ldots, t_n\})}{n}$$

whenever the limit exists. It is well known that given any free ultrafilter $\mathcal{U}$, the set function $\lambda_{\mathcal{U}} : 2^T \to [0, 1]$ given by

$$\lambda_{\mathcal{U}}(A) \equiv \mathcal{U} - \lim_{n \to \infty} \frac{\#(A \cap \{t_1, \ldots, t_n\})}{n}, \tag{27}$$

defines a finitely additive extension of the usual density to the power set of T.

We are not interested in the limit in Eq. (27) directly, but rather in the limit of terms of the form $\frac{\#(A \cap T_N)}{\#T_N}$ that are defined on the subset of integers:

$$\mathcal{M} = \{\#T_1, \#T_2, \ldots, \#T_N, \ldots\} \subset \mathcal{N}$$

rather than all of $\mathcal{N}$. More generally, for a sequence $\{x_N\}$ defined only on $\mathcal{M}$,

$$\mathcal{U} - \lim_{N \to \infty} x_N = x \iff \forall \epsilon > 0, \quad \{\#T_N: |x_N - x| < \epsilon\} \in \mathcal{U}. \tag{28}$$

If $\mathcal{M} \in \mathcal{U}$, then the convergence criteria in Eqs. (26) and (28) are consistent. That is, for any pair of sequences $\{x'_n\}$ and $\{x_N\}$ such that $x_N \equiv x'_{\#T_N}$ for each N, $\mathcal{U} - \lim_{n \to \infty} x'_n = \mathcal{U} - \lim_{N \to \infty} x_N$ if $\mathcal{M} \in \mathcal{U}$.

The following proposition explains that λ satisfies Eq. (18) if and only if it is generated by a free ultrafilter containing $\mathcal{M}$:

Proposition A.1. λ *is a finitely additive probability measure on* T *satisfying the condition* (*see Definition* 7): "*For every finite collection of subsets* $A_1, \ldots, A_L \subset T$ *there is a subsequence* $\{N_k\}$ *such that*:

$$\lambda(A_l) = \lim_{k \to \infty} \lambda_{N_k}(A_l), \quad for\ l = 1, \ldots, L, \text{"}$$

if and only if there is a free ultrafilter $\mathcal{U}$ *containing* $\mathcal{M}$ *such that* $\lambda = \lambda_{\mathcal{U}}$.

Proof. For the 'if' part, suppose that $\mathcal{U}$ contains $\mathcal{M}$ and generates λ. Fix the subsets $A_1, \ldots, A_L$. Since $\mathcal{U}$ generates λ,

$$\left\{ n \colon \left| \frac{\#(A_l \cap \{t_1, \ldots, t_n\})}{n} - \lambda(A_l) \right| < \frac{1}{k} \right\} \in \mathcal{U}$$

for every positive integer k. Since $\mathcal{M} \in \mathcal{U}$,

$$\left\{ n \colon \left| \frac{\#(A_l \cap \{t_1, \ldots, t_n\})}{n} - \lambda(A_l) \right| < \frac{1}{k} \right\} \cap \mathcal{M} = \left\{ \#T_N \colon \left| \lambda_N(A_l) - \lambda(A_l) \right| < \frac{1}{k} \right\}$$

also belongs to $\mathcal{U}$. Consequently, $\bigcap_{l=1}^{L} \{ \#T_N \colon |\lambda_N(A_l) - \lambda(A_l)| < \frac{1}{k} \} \in \mathcal{U}$ and is thus non-empty. Let N_k denote any element of that set. Repeating this process for every integer k, we may obtain a strictly increasing subsequence $\{N_k\}$ such that for every l and every $\epsilon > 0$ there is K such that $k > K$ implies $|\lambda_{N_k}(A_l) - \lambda(A_l)| < \frac{1}{k} < \epsilon$, as required.

For the 'only if' part, define $\mathcal{U}^* \subset 2^{\mathcal{N}}$ to consist of all subsets of the form $\{\#T_N \colon |\lambda_N(A) - \lambda(A)| < \epsilon\}$ for some $A \subset T$ and $\epsilon > 0$. Clearly, $\mathcal{M} \in \mathcal{U}^*$. It is also easy to see that a consequence of Eq. (18) is that if $U = \{\#T_N \colon |\lambda_N(A) - \lambda(A)| < \epsilon\}$ and $U' = \{\#T_N \colon |\lambda_N(A') - \lambda(A')| < \epsilon'\}$ are two sets in $\mathcal{U}^*$, then $U \cap U' \neq \emptyset$. Thus, $\mathcal{U}^*$ is a collection of non-empty subsets of $\mathcal{N}$ with the finite intersection property and can therefore be extended to an ultrafilter $\mathcal{U}$ on $\mathcal{N}$. It only remains to show that $\mathcal{U}$ generates λ; that is, $\lambda(A) = \mathcal{U} - \lim_{N \to \infty} \frac{\#(A \cap T_N)}{\#T_N}$ for every $A \subset T$. But, by construction, we have $\{\#T_N \colon |\lambda_N(A) - \lambda(A)| < \epsilon\} \in \mathcal{U}^* \subset \mathcal{U}$ for any $\epsilon > 0$. $\quad \square$

A.1.2. Completeness of the strategy space

One reason why finitely additive probabilities have been shunned is their lack of tractability. In particular, many of the classic theorems in Analysis that one takes for granted in a countably additive context, like the completeness of L^p spaces and the Radon–Nikodym Theorem, fail for general finitely additive probabilities.

With appropriate care, however, one can show that finitely additive spaces can be constructed that do not suffer from these problems:

Proposition A.2. *Suppose that the finitely additive probability λ on T satisfies $\lambda = \lambda_{\mathcal{U}}$ for some free ultrafilter $\mathcal{U}$ that contains $\mathcal{M}$ and that*

$$\frac{\#T_N - \#T_{N-1}}{\#T_N} \to 1, \quad as \ N \to \infty. \tag{17}$$

Then the following properties obtain:

1. *Completeness of $L^1(T, 2^T, \lambda)$: For any sequence of functions $f_m \colon T \to \mathcal{R}$, $m = 1, 2 \ldots$ with $\lim_{m,m' \to \infty} \int |f_m - f_{m'}| \, d\lambda = 0$ there is a function f such that $\lim_{m \to \infty} \int |f_m - f| \, d\lambda = 0$.*
2. *Radon–Nikodym Property: For every finitely additive measure ν on $(T, 2^T)$ that is absolutely continuous with respect to λ,[38] that is:*

$$\lambda(A) = 0 \implies \nu(A) = 0, \quad for \ all \ A \subset T,$$

---

[38] There are various versions of absolute continuity for finitely additive measures. See Bhaskara Rao and Bhaskara Rao (1983, Section 6.1). However, they are all equivalent here because 2^T is a σ-algebra and λ and ν are both non-negative.

there is a function $f : T \to \mathcal{R}$ such that

$$v(A) = \int_A f \, d\lambda, \quad \text{for all } A \subset T.$$

A stronger form of the proposition may be found in Al-Najjar (2007). The key step is a result by Blass et al. (2001) who showed that the requirement that $\mathcal{U}$ satisfies condition (17) together with results by Gangopadhyay and Rao (1999) and Basile and Bhaskara Rao (2000) imply the completeness of $L^1(T, \Sigma, \lambda)$. It is well known (see Basile and Bhaskara Rao, 2000) that completeness of L^1 is equivalent to the completeness of L^p spaces for any p, and to the space $(T, 2^T, \lambda)$ having the Radon–Nikodym Property.

A.1.3. Useful lemmas
Lemma A.1.

1. *For any bounded sequence $\{x_N\}$ there is a subsequence $\{x_{N_k}\}$ such that*

$$\lim_{k \to \infty} x_{N_k} = \mathcal{U} - \lim_{N \to \infty} x_N.$$

2. *For any pair of bounded sequences $\{x_N\}$ and $\{y_N\}$,*

$$\lim_{N \to \infty} |x_N - y_N| = 0 \implies \mathcal{U} - \lim_{N \to \infty} x_N = \mathcal{U} - \lim_{N \to \infty} y_N.$$

The next lemma provides a useful criterion to calculate integrals.

Lemma A.2. *For any bounded function $f : T \to \mathcal{R}$,*

$$\int_T f \, d\lambda = \mathcal{U} - \lim_{N \to \infty} \int_T f \, d\lambda_N$$

$$\equiv \mathcal{U} - \lim_{N \to \infty} \frac{1}{\#T_N} \sum_{n=1}^{\#T_N} f(t_n).$$

Proof. For every $\epsilon > 0$ there is a simple function f' such that $|f(t) - f'(t)| < \frac{\epsilon}{2}$, hence $|\int_T f \, d\lambda - \int_T f' \, d\lambda| < \frac{\epsilon}{2}$. Write $f' = \sum_{k=1}^{K} r_k \chi_{B_k}$ where $B_1, \ldots, B_K$ is a partition of T, and $r_1, \ldots, r_K$ are constants. Then:

$$\mathcal{U} - \lim_{N \to \infty} \frac{1}{N} \sum_{n=1}^{N} f(t_n) = \mathcal{U} - \lim_{N \to \infty} \frac{1}{N} \sum_{n=1}^{N} \left[f(t_n)(\chi_{B_1} + \cdots + \chi_{B_K}) \right]$$

$$= \sum_{k=1}^{K} \mathcal{U} - \lim_{N \to \infty} \frac{1}{N} \sum_{n=1}^{N} f(t_n) \chi_{B_k}$$

$$> \sum_{k=1}^{K} r_k \, \mathcal{U} - \lim_{N \to \infty} \frac{1}{N} \sum_{n=1}^{N} \chi_{B_k} - \frac{\epsilon}{2}$$

$$= \sum_{k=1}^{K} r_k \lambda(B_k) - \frac{\epsilon}{2}$$

$$= \int\limits_T f' \, \mathrm{d}\lambda - \frac{\epsilon}{2}.$$

Using a similar argument, we conclude that $|\mathcal{U} - \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} f(t_n) - \int_T f' \, \mathrm{d}\lambda| < \frac{\epsilon}{2}$ and, consequently, that $|\mathcal{U} - \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} f(t_n) - \int_T f \, \mathrm{d}\lambda| < \epsilon$.   □

A.2. Proof for Section 2: discrete large games

Proof of Theorem 1. Fix $B \in \mathcal{B}$. Using the version of the strong law of large numbers in Shiryayev (1984, Theorem 2, p. 364) there is $\Omega_B \subset S$ with $P^\mu(\Omega_B) = 1$ such that for every $s \in \Omega_B$

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \chi_B(t_n) \big[s(t_n) - \mu(t_n) \big] = 0.$$

This implies[39]:

$$\mathcal{U} - \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \chi_B(t_n) s(t_n) = \mathcal{U} - \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \chi_B(t_n) \mu(t_n).$$

By Lemma A.2, we have $\int_B s(t) \, \mathrm{d}\lambda = \int_B \mu \, \mathrm{d}\lambda$.

Define $\Omega' = \bigcap \{\Omega_B : B \in \mathcal{B}\}$ and note that it is a countable intersection of probability 1 events. Therefore, $P^\mu(\Omega') = 1$, proving the theorem.   □

Proof of Theorem 2. The conclusion follows by noting that:

$$U(t, \mu) \equiv \int\limits_S u\big(t, s(t), s\big) \, \mathrm{d}P^\mu$$

$$= \sum_{a\in A} \mu(t)(a) \int\limits_S u(t, a, s) \, \mathrm{d}P^\mu \tag{29}$$

$$= \sum_{a\in A} \mu(t)(a) \, u(t, a, \mu) \tag{30}$$

---

[39] Proof: write $F_N \equiv \frac{1}{N} \sum_{n=1}^{N} \chi_B(t_n) s(t_n)$ and $M_N \equiv \frac{1}{N} \sum_{n=1}^{N} \chi_B(t_n) \mu(t_n)$ and fix any $\epsilon > 0$. Then

$$\left\{ N : |F_N - \mathcal{U} - \lim_{N\to\infty} F_N| < \frac{\epsilon}{2} \right\} \in \mathcal{U} \quad \text{and} \quad \left\{ N : |F_N - M_N| < \frac{\epsilon}{2} \right\} \in \mathcal{U}$$

so

$$\left\{ N : |F_N - \mathcal{U} - \lim_{N\to\infty} F_N| < \frac{\epsilon}{2} \right\} \cap \left\{ N : |F_N - M_N| < \frac{\epsilon}{2} \right\} \in \mathcal{U}$$

and

$$\left\{ N : |M_N - \mathcal{U} - \lim_{N\to\infty} F_N| < \epsilon \right\} \in \mathcal{U}$$

hence

$$\mathcal{U} - \lim_{N\to\infty} M_N = \mathcal{U} - \lim_{N\to\infty} F_N.$$

where Eq. (29) follows from the assumption that players randomize independently, and Eq. (30) follows from Theorem 1. $\quad\square$

Lemma A.3. *If μ is an equilibrium, then for every positive integer k:*

$$\lambda\left\{t:\ v_k(t)\equiv\sum_{a\notin\mathrm{Br}_{1/k}(t,\mu)}\mu(t)(a)<\frac{1}{k}\right\}>1-\frac{1}{k}.$$

In particular,

$$\int v_k\,\mathrm{d}\lambda\leqslant\frac{1}{k}.$$

Proof. Let $v^*(t,\mu)\equiv\max_a u(t,a,\mu)$ and fix k. Then $\mu(t)\in\mathrm{Br}_{1/k^2}(t,\mu)$ implies that

$$\begin{aligned}
v^*(t,\mu)-\frac{1}{k^2}&\leqslant\sum_a\mu(t)(a)u(t,a,\mu)\\
&\leqslant v^*(t,\mu)\bigl(1-v_k(t)\bigr)+\left(v^*(t,\mu)-\frac{1}{k}\right)v_k(t)\\
&=v^*(t,\mu)-\frac{1}{k}v_k(t),
\end{aligned}$$

from which it follows that $v_k(t)\leqslant\frac{1}{k}$. Since μ is an equilibrium, we have

$$\lambda\bigl\{t:\ \mu(t)\in\mathrm{Br}_{1/k^2}(t,\mu)\bigr\}=\lambda\left\{t:\ v_k(t)\leqslant\frac{1}{k}\right\}=1,$$

so

$$\int v_k\,\mathrm{d}\lambda\leqslant\frac{1}{k}.\qquad\square$$

Proof of Theorem 3. From Theorem 1, almost every realization is a pure strategy s such that $\int_T s\,\mathrm{d}\lambda=\int_T\mu\,\mathrm{d}\lambda$. For any such s and any $\epsilon>0$, $\mathrm{Br}_\epsilon(t,\mu)=\mathrm{Br}_\epsilon(t,s)$ for every t.

Let $z_k(t,\mu)$ denote the random variable that takes the value 1 if $s(t)\in\mathrm{Br}_{1/k}(t,\mu)$ and zero otherwise. Then $Ez_k(t,\mu)=1-v_k(t)$. By the law of large numbers, Theorem 1, and Lemma A.3:

$$\int z_k\,\mathrm{d}\lambda=1-\int v_k\,\mathrm{d}\lambda\geqslant1-\frac{1}{k},\quad P^\mu\text{-a.s.}$$

Note that Theorem 1 is formulated for strategy profiles but, as is clear from the proof, it is valid for any independent, finite-valued family of random variables on T.

This says that

$$P^\mu\left\{s:\ \lambda\bigl\{t:\ s(t)\in\mathrm{Br}_{1/k}(t,\mu)\bigr\}\geqslant1-\frac{1}{k}\right\}=1.$$

Since P^μ is countably additive,

$$P^\mu\left(\bigcap_{k=1}^\infty\left\{s:\ \lambda\bigl\{t:\ s(t)\in\mathrm{Br}_{1/k}(t,\mu)\bigr\}\geqslant1-\frac{1}{k}\right\}\right)=1.$$

Any pure strategy profile s in this intersection has the property that $\lambda\{t\colon s(t) \in \mathrm{Br}_{1/k}(t, \mu)\} \geqslant 1 - \frac{1}{k}$ for every k, and is hence an equilibrium.[40]

The argument for the case where μ is exact: $z(t, \mu)$ now defines the random variable that takes the value 1 if $s(t) \in \mathrm{Br}(t, \mu)$ and zero otherwise. The law of large numbers would then imply:

$$\int z \, d\lambda = 1 - \int v \, d\lambda = 1, \quad P^\mu\text{-a.s.}$$

so

$$P^\mu\{s\colon \lambda\{t\colon s(t) \in \mathrm{Br}(t, \mu)\} = 1\} = 1.$$

The remainder of the argument is identical to the above. □

A.3. Lemmas on uniform continuity

In the body of the paper, uniform continuity is defined with respect to an arbitrary choice of a metric on $\bar{T}$. Since such metric plays no substantive role, the next lemma replaces uniform continuity with a slightly weaker and more convenient (but slightly more complicated to state) metric-free condition. This is what we use in the remainder of the paper.

Lemma A.4. *If a payoff function $u\colon T \times A \times \Delta^A \to \mathcal{R}$ is uniformly continuous with respect to a metric d that generates the product topology on $\bar{T} \times A \times \Delta^A$ then for every $\epsilon > 0$ there is $\epsilon' > 0$ and a partition $\{B_1, \ldots, B_K\} \subset \mathcal{B}$ such that*

$$\left| \bar{u}(t, a, \alpha) - u(t', a', \alpha') \right| < \epsilon \tag{31}$$

whenever $a = a'$, $|\alpha - \alpha'| < \epsilon'$ and $t,\ t' \in B$ for some $B \in \{B_1, \ldots, B_K\}$.

For future reference, we note that this lemma holds for payoff functions $\bar{u}$ in continuum-player games.

Proof. Suppose that $\bar{u}$ is uniformly continuous and fix $\epsilon > 0$. Then there is $\epsilon' > 0$ such that Eq. (31) holds whenever $a = a'$, $|\alpha - \alpha'| < \epsilon'$ and $|t - t'| < \epsilon'$. But then find any partition $B \in \{B_1, \ldots, B_K\}$ with d-diameter not exceeding ϵ'. □

Lemma A.5. *If u is uniformly continuous then for every $\epsilon > 0$ there is $\epsilon' > 0$ and a partition $\{B_1, \ldots, B_K\} \subset \mathcal{B}$ such that*

$$\alpha_1' \in \mathrm{Br}_{\epsilon'}(t', \alpha_2') \quad \Longrightarrow \quad \alpha_1 \in \mathrm{Br}_\epsilon(t, \alpha_2) \tag{32}$$

whenever $|\alpha_1 - \alpha_1'| < \epsilon'$, $|\alpha_2 - \alpha_2'| < \epsilon'$ and $t,\ t' \in B$ for some $B \in \{B_1, \ldots, B_K\}$.

Proof. Choose $\bar{\epsilon} < \frac{\epsilon}{4}$ and apply uniform continuity relative to $\bar{\epsilon}$, finding $\epsilon' \in (0, \frac{\epsilon}{2})$ and a partition $B \in \{B_1, \ldots, B_K\}$ so that

$$\left| u(t, \alpha_1, \alpha_2) - u(t', \alpha_1', \alpha_2') \right| < \bar{\epsilon},$$

whenever $|\alpha_1 - \alpha_1'| < \epsilon'$, $|\alpha_2 - \alpha_2'| < \epsilon'$ and $t,\ t' \in B$ for some $B \in \{B_1, \ldots, B_K\}$.

---

[40] Note that this implies that any such s must satisfy Eq. (12). This does not imply, however, that s is an exact equilibrium.

Take any $\beta \in \text{Br}(t, \alpha_2)$. Then $\alpha_1' \in \text{Br}_{\epsilon'}(t', \alpha_2')$ implies:

$$u(t', \alpha_1', \alpha_2') \geqslant u(t', \beta, \alpha_2') - \epsilon'.$$

Thus, if t, t' belong to some B and $|\alpha_1 - \alpha_1'| < \epsilon'$ and $|\alpha_2 - \alpha_2'| < \epsilon'$, uniform continuity implies that:

$$u(t, \alpha_1, \alpha_2) + \bar{\epsilon} > u(t', \alpha_1', \alpha_2') \geqslant u(t', \beta, \alpha_2') - \epsilon' > u(t, \beta, \alpha_2) - \bar{\epsilon} - \epsilon'.$$

That is:

$$u(t, \alpha_1, \alpha_2) > u(t, \beta, \alpha_2) - 2\bar{\epsilon} - \epsilon'.$$

But then $-2\bar{\epsilon} - \epsilon' \geqslant -2\frac{\epsilon}{4} - \frac{\epsilon}{2} = -\epsilon$, as required. $\quad\square$

For $t \in \bar{T}$ and $k = 1, 2, \dots$ define $B(t, k)$ to be the smallest (by set inclusion) element of $\mathcal{B}_k$ containing t. Given a profile μ in Γ, define the *k-conditional strategy* in Γ by

$$\mu_k(t) = \frac{1}{\lambda(B(t, k))} \int\limits_{B(t,k)} \mu \, d\lambda, \quad t \in T. \tag{33}$$

Similarly, for a profile $\bar{\mu}$ in $\bar{\Gamma}$, define

$$\bar{\mu}_k(t) = \frac{1}{\bar{\lambda}(B(t, k))} \int\limits_{B(t,k)} \bar{\mu} \, d\bar{\lambda}, \quad t \in \bar{T}. \tag{34}$$

Proposition A.3.

1. *If μ is an equilibrium in Γ and μ_k is its k-conditional strategy, then*: *For every $\epsilon > 0$ there is K such that $k > K$ implies that $\mu_k(t) \in \text{Br}_\epsilon(t, \mu_k)$ for every $t \in T$.*[41]
2. *A similar conclusion holds for any equilibrium $\bar{\mu}$ of $\bar{\Gamma}$ and corresponding k-conditional strategies $\bar{\mu}_k$.*[42]

Proof. We begin with part 1. Since u is uniformly continuous, using Lemma A.4 given $\epsilon > 0$ there is K large enough such that for every $k > K$

$$\max_{t, t' \in B(t,k)} \max_{\alpha_1, \alpha_2 \in \Delta} \left| u(t, \alpha_1, \alpha_2) - u(t', \alpha_1, \alpha_2) \right| < \frac{\epsilon}{4}. \tag{35}$$

Given $k > K$ and $\bar{t} \in T$,

$$a \in \text{supp}\, \mu_k(\bar{t}) \quad \Longleftrightarrow \quad \int\limits_{B(\bar{t},k)} \mu(t)(a) \, d\lambda > 0$$

$$\Longrightarrow \quad \exists l > \frac{2}{\epsilon}, \ \exists C \subset B(\bar{t}, k), \ \lambda(C) > 0:$$
$$\forall t \in C, \ a \in \text{Br}_{\frac{1}{l}}(t, \mu) \tag{36}$$

---

[41] This is stronger than: For every $\epsilon > 0$ there is K such that $k > K$ implies that μ_k is an ϵ-equilibrium. What is being shown here is that *every* player ϵ-best responds.

[42] That is, using the same quantifiers, $\bar{\mu}_k(t) \in \text{Br}_\epsilon(t, \bar{\mu}_k)$ for every $t \in \bar{T}$.

$$\implies \quad \forall t' \in B(\bar{t}, k):$$

$$\max_{\beta \in \Delta} u(t', \beta, \mu) - u(t', a, \mu) < \frac{1}{l} + 2\frac{\epsilon}{4} \tag{37}$$

$$\implies \quad \forall t' \in B(\bar{t}, k):$$

$$\max_{\beta \in \Delta} u(t', \beta, \mu_k) - u(t', a, \mu_k) < \epsilon \tag{38}$$

$$\implies \quad \forall t' \in B(\bar{t}, k): \ a \in \mathrm{Br}_\epsilon(t', \mu_k).$$

Comments:

- The implication in (36) follows from the assumption that μ is an equilibrium and Lemma A.3. To see this, recall that for any t, $a \notin \mathrm{Br}_{1/l}(t, \mu)$ implies, in the notation of Lemma A.3, that $\mu(t)(a) \leqslant v_l(t)$. Thus, if for every $l > \frac{2}{\epsilon}$ we had $a \notin \mathrm{Br}_{1/l}(t, \mu)$ for almost every $t \in B(\bar{t}, k)$, then $\int_{B(\bar{t},k)} \mu(t)(a)\, d\lambda \leqslant \int_{B(\bar{t},k)} v_l(t)\, d\lambda \leqslant \frac{1}{l}$. Choosing l large enough yields a contradiction. Note that equilibrium in Γ does not imply that there is t' such that $a \in \mathrm{Br}(t', \mu)$, but only that for every $\epsilon' > 0$, $a \in \mathrm{Br}_\epsilon(t', \mu)$ for almost every t'.
- The implication in (37) uses Eq. (35).
- The implication in (38) uses the fact that $\int \mu_k\, d\lambda = \int \mu\, d\lambda$ for every k, by definition.

The argument for part 2 is similar (and in fact a little simpler, since there is no need to use Lemma A.3). $\quad \square$

The following lemma shows that Definition 9 is not vacuous:

Lemma A.6. *Proper sequences of finite-player games exist.*

Proof. Fix any continuum-player game $\bar{\Gamma}$. Define the infinite product space $(\bar{T}^\infty, \bar{\mathcal{B}}^\infty, \bar{\lambda}^\infty)$ of all infinite independent samples from $\bar{T}$ using $\bar{\lambda}$. For any $B_l \in \mathcal{B}$, let $C(B_l) \subset \bar{T}^\infty$ be the set of all samples in which the asymptotic frequency of B_l is $\bar{\lambda}(B_l)$. A consequence of the law of large numbers is that $\bar{\lambda}^\infty(C(B_l)) = 1$, hence

$$\bar{\lambda}^\infty\left(\bigcap\{C(B_l): \ B_l \in \mathcal{B}\}\right) = 1.$$

Let $\{t_1, \ldots\}$ be any element of $\cap\{C(B_l): \ B_k \in \mathcal{B}\}$.

Define the sequence of finite games so that $T_{N_k} = \{t_1, \ldots, t_{N_k}\}$, for any strictly increasing sequence of integers $\{N_k\}$ that grows fast enough so condition (17) holds. Let $u_N(t, \cdot, \cdot) = \bar{u}(t, \cdot, \cdot)$ for every $t \in T_N$ and N. The requirement that for every $B \in \mathcal{B}$, $\lim_{N \to \infty} \lambda_N(B) = \bar{\lambda}(B)$ is satisfied by construction. $\quad \square$

A.4. Proofs for Sections 3 and 4

Proof of Theorem 5. By Lemma A.1 there is a subsequence $\{N_k\}$ such that

$$\lim_{k \to \infty} \int \mu_{N_k}\, d\lambda_{N_k} = \mathcal{U} - \lim_{N \to \infty} \int \mu_N\, d\lambda_N \equiv \alpha.$$

Let $N_k(t)$ be the smallest k such that $t \in T_{N_k}$. That is, $\Gamma_{N_k(t)}$ is the first game in which player t appears. Define μ by:

$$\mu(t) \equiv \mu_{N_k(t)}(t).$$

As an intermediate step, I show that the constructed profile μ satisfies:

$$\int \mu \, d\lambda = \alpha. \tag{$*$}$$

By construction of μ, we have, for every N:

$$\left| \int \mu \, d\lambda_N - \int \mu_N \, d\lambda_N \right| \leqslant \frac{1}{\#T_N} \sum_{n=1}^{\#T_N} \left| \mu(t_n) - \mu_N(t_n) \right|$$

$$= \frac{1}{\#T_N} \sum_{n=1}^{\#T_{N-1}} \left| \mu(t_n) - \mu_N(t_n) \right| \tag{39}$$

$$\leqslant \frac{\#T_{N-1}}{\#T_N},$$

where Eq. (39) follows from the fact that μ and μ_{N_k} coincide on $T_{N_k} - T_{N_{k-1}}$ by definition.

By the assumption of properness of the sequence of finite games, we have $\frac{\#T_{N-1}}{\#T_N} \to 0$ as $N \to \infty$. In particular (since $\mathcal{U}$ is free), for every $\epsilon > 0$ we have:

$$\left\{ N : \left| \int \mu_N \, d\lambda_N - \int \mu \, d\lambda_N \right| < \epsilon \right\} \in \mathcal{U}.$$

From the definition of "$\mathcal{U} - \lim$" we also have, for every $\epsilon > 0$:

$$\left\{ N : \left| \int \mu_N \, d\lambda_N - \alpha \right| < \epsilon \right\} \in \mathcal{U}$$

Combining these two facts, we have:

$$\left\{ N : \left| \int \mu \, d\lambda_N - \alpha \right| < \epsilon \right\} \in \mathcal{U};$$

that is,

$$\mathcal{U} - \lim_{N \to \infty} \int \mu \, d\lambda_N \equiv \alpha.$$

The desired conclusion $(*)$ now follows by applying Lemma A.2.

The last step in the proof is to show that μ is an equilibrium. Fix $\epsilon > 0$. The continuity of $u(t, \cdot, \cdot)$, the fact that $\mu_{N_k}(t) \in \text{Br}_\epsilon(t, \mu_{N_k})$ for every k and t, and $(*)$ imply that we can find k large enough so that $\mu_{N_k}(t) \in \text{Br}_{2\epsilon}(t, \mu)$ for every $t \notin T_{N_{k-1}}$.[43] Using the definition of μ we conclude that $\mu(t) \in \text{Br}_{2\epsilon}(t, \mu)$ for every t outside the finite (hence λ-measure zero) set of players $T_{N_{k-1}}$, so μ is indeed an equilibrium. $\quad \square$

Proof of Theorem 6. Fix $\epsilon > 0$. Using the uniform continuity of u and applying Lemma A.5, we conclude that there is $\epsilon > \epsilon' > 0$ such that for every t and N for which $| \int \mu \, d\lambda - \int \mu_N \, d\lambda_N | < \epsilon'$, $\mu(t) \in \text{Br}_{\epsilon'}(t, \mu)$ implies $\mu(t) \in \text{Br}_\epsilon(t, \mu_N)$.

It only remains to show that there is N such that $| \int \mu \, d\lambda - \int \mu_N \, d\lambda_N | < \epsilon'$ and $\lambda_N \{t : \mu(t) \in \text{Br}_{\epsilon'}(t, \mu)\} > 1 - \epsilon$. Since μ is an equilibrium, $\lambda \{t : \mu(t) \in \text{Br}_{\epsilon'}(t, \mu)\} = 1$. By Eq. (27), the set $A \equiv \{\#T_N : \lambda_N \{t : \mu(t) \in \text{Br}_{\epsilon'}(t, \mu)\} > 1 - \epsilon\}$ belongs to $\mathcal{U}$. By Lemma A.2, $\mathcal{U}$ also contains the

---

[43] Note that we are not using continuity in t nor are we invoking Lemma A.5.

set $B \equiv \{\#T_N : |\int \mu \, d\lambda - \int \mu_N \, d\lambda_N| < \epsilon'\}$. Any N such that $\#T_N \in A \cap B$ (which is non-empty, since $\mathcal{U}$ is a filter) has the desired properties. $\square$

Proof of Theorem 7. Define Γ so that $T = \bigcup_{N=1}^{\infty} T_N$ and $u(t, \cdot, \cdot) = u_N(t, \cdot, \cdot)$ for every N and $t \in T_N$, which is possible since $\{\Gamma_N\}$ is a proper sequence. Define λ as $\lambda = \lambda_{\mathcal{U}}$ for some free ultrafilter containing $\mathcal{M}$.

The last two conditions in the definition of a proper sequence imply $\lambda(B) = \lim_{N\to\infty} \lambda_N(B) > 0$ for every $B \in \mathcal{B}$ and $\lambda(t) = 0$ for every $t \in \bar{T}$. Thus, Γ is indeed a discrete large game. Proposition A.1 implies that Γ satisfies Eq. (18).

To construct $\bar{\Gamma}$ we first define $\bar{u}$ as the (necessarily) unique uniformly continuous extension of

$$u : \bigcup_{N=1}^{\infty} T_N \times A \times \Delta \to \mathcal{R}.$$

To define $\bar{\lambda}$, note that λ is finitely additive on the algebra $\mathcal{B}$. But then λ (viewed as a measure on $\bar{T}$) must be countably additive on $\mathcal{B}$, and thus admits a unique countably additive extension $\bar{\lambda}$ to $\bar{\mathcal{B}}$ (Billingsley, 1995, Theorems 2.3 and 3.1 respectively)). The requirement:

$$\lim_{N\to\infty} \lambda_N(B) = \lambda(B) = \bar{\lambda}(B), \quad \text{for every } B \in \mathcal{B}$$

is now satisfied by construction. $\square$

Proof of Theorem 8. 1. Let $\{\bar{\mu}_k\}$ denote the sequence of k-conditional strategies corresponding to $\bar{\mu}$. Let μ_k denote the restriction of $\bar{\mu}_k$ to T (although we do not usually notationally distinguish between simple functions on $\bar{T}$ and their restrictions to T, we do so here for expository clarity). From Levy's Theorem (Shiryayev, 1984, Theorem 3, p. 478) we have that $\bar{\mu}_k \to \bar{\mu}$ in $L^1(\bar{\lambda})$. In particular, it is a Cauchy sequence. Then $\{\mu_k\}$ must also be Cauchy since the L^1 norm is preserved for simple functions. By Proposition A.2, $L^1(\lambda)$ is complete so $\{\mu_k\}$ must have a limit μ. Clearly, μ and $\bar{\mu}$ are behaviorally equivalent via the sequence of strategies $\{\bar{\mu}_k\}$.

2. Let $\{\bar{\mu}_k\}$ be the sequence of conditional strategies on $\bar{T}$ induced by μ. This sequence is a bounded martingale with respect to $\{\mathcal{B}_k\}$. By the Martingale Convergence Theorem (see, e.g., Shiryayev, 1984, Theorem 1, p. 476) there is a $\bar{\mathcal{B}}$-measurable function $\bar{\mu}$ such that $\bar{\mu}_k(t) \to \bar{\mu}(t)$ for $\bar{\lambda}$-a.e. t.

3. By way of contradiction, suppose that $s'_1, s'_2 \ldots,$ is an enumeration that exhausts the set of pure profiles that are distributionally equivalent to μ. For each l define $S_l = \{s : \int_T |s'_l - s| \, d\lambda > 0\}$. Then, under P^μ, $\{|s'_l(t) - s(t)|; t \in T\}$ is a collection of bounded and independent random variables so for each l, P^μ-a.s.,

$$\int_T |s'_l - s| \, d\lambda = \int_T E|s'_l - s| \, d\lambda \tag{40}$$

$$= \int_T |s'_l - \mu| \, d\lambda$$

$$> 0. \tag{41}$$

where Eq. (40) follows from the law of large numbers and inequality (41) follows from the assumption that μ is not essentially pure. This means that $P^\mu(S_l) = 1$ for each l. Since P^μ

is countably additive, the set $\bigcap_l S_l$ has P^μ-probability 1 and is therefore non-empty. But any $s'' \in \bigcap_l S_l$ must be distinct from each element of the sequence of s_l's. A contradiction. $\quad\square$

By Theorems 1 and 3 almost every realization of P^μ is a pure strategy equilibrium that is distributionally equivalent to μ.

Proof of Theorem 9. 1. Let μ_k, $k = 1, \ldots$ denote the k-conditional strategies μ in Γ and let $\bar{\mu}_k$ be the corresponding strategies in $\bar{\Gamma}$. By Proposition A.3, for every $\epsilon > 0$ there is K such that $k > K$ implies that $\mu_k(t) \in \text{Br}_\epsilon(t, \mu_k)$ for every $t \in T$. This in turn implies that $\bar{\mu}_k$ is an ϵ-equilibrium for $\bar{\Gamma}$. To see this, for any $\bar{t} \in \bar{T}$, take a sequence $\{t_l\}$ in T converging to $\bar{t}$. By our earlier argument, $\mu_k(t_l) \in \text{Br}_\epsilon(t_l, \mu_k)$ for each l. By the continuity of $\bar{u}$, we also have $\bar{\mu}_k(\bar{t}) \in \text{Br}_\epsilon(\bar{t}, \mu_k)$. Note that this also shows that, for each k, $\bar{\mu}_k$ is ϵ_k-equilibrium where $\epsilon_k \downarrow 0$.

To show that $\bar{\mu}$ is an equilibrium for $\bar{\Gamma}$, first recall the martingale argument made in part 2 of the proof of Theorem 8 and let $Z \subset \bar{T}$ denote the set with $\bar{\lambda}(Z) = 1$ for which $\{\bar{\mu}_k\}$ converges to $\bar{\mu}$. Then, for every $t \in Z$ and $b \in A$,

$$\bar{u}\big(t, \bar{\mu}_k(t), \bar{\mu}_k\big) \geqslant \bar{u}(t, b, \bar{\mu}_k) - \epsilon_k$$

so

$$\bar{u}\big(t, \bar{\mu}(t), \bar{\mu}\big) \geqslant \bar{u}(t, b, \bar{\mu}),$$

by continuity of $\bar{u}$ and the integral. Thus, $\bar{\mu}$ is an equilibrium of $\bar{\Gamma}$.

2. Suppose that $\bar{\mu}$ is an equilibrium for $\bar{\Gamma}$. We first collect some facts. By Proposition A.3, for every $\epsilon_1 > 0$ and all large enough k, $\bar{\mu}_k(t) \in \text{Br}_{\epsilon_1}(t, \bar{\mu}_k)$ for every $t \in \bar{T}$. Since $T \subset \bar{T}$ and μ and $\bar{\mu}$ are behaviorally equivalent, we also have:

$$\mu_k(t) \in \text{Br}_{\epsilon_1}(t, \mu_k), \quad \forall t \in T. \tag{42}$$

Also, since $\|\mu_k - \mu\| \to 0$ implies $\mu_k \to \mu$ (Bhaskara Rao and Bhaskara Rao, 1983, Theorems 4.6.10 or 4.6.14) given $\epsilon_2 > 0$ for all large enough k

$$\lambda\big(\{t\colon |\mu(t) - \mu_k(t)| < \epsilon_2\}\big) > 1 - \epsilon_2. \tag{43}$$

Combining the above facts, given any $\epsilon > 0$ choose ϵ' to satisfy Eq. (32) then choose $\epsilon_1, \epsilon_2 < \epsilon$ and k large enough so that Eq. (42) and (43) are satisfied. This implies that $\lambda(\{t\colon \mu(t) \in \text{Br}_\epsilon(t, \mu)\}) > 1 - \epsilon$. Since ϵ is arbitrary, μ is an equilibrium for Γ. $\quad\square$

References

Al-Najjar, N.I., 1995. Decomposition and characterization of risk with a continuum of random variables. Econometrica 63, 1195–1224. Corrigendum; Econometrica 67 (1999) 919–920.

Al-Najjar, N.I., 2004. Aggregation and the law of large numbers in large economies. Games Econ. Behav. 47, 1–35.

Al-Najjar, N.I., 2007. Finitely additive representation of L^p spaces. J. Math. Anal. Appl. 330, 891–899.

Al-Najjar, N.L., Anderlini, L., Felli, L., 2006. Undescribable events. Rev. Econ. Stud. 73, 849–868.

Basile, A., Bhaskara Rao, K.P.S., 2000. Completeness of L_p-spaces in the finitely additive setting and related stories. J. Math. Anal. Appl. 248 (2), 588–624.

Bewley, T., 1986. Stationary monetary equilibrium with a continuum of independently fluctuating consumers. In: Hildenbrand, W., Mas-Colell, A. (Eds.), Contributions to Mathematical Economies: In honor of Gerard Debreu. North-Holland, New York.

Bhaskara Rao, K.P.S., Bhaskara Rao, M., 1983. Theory of Charges. In: Pure Appl. Math., vol. 109. Academic Press, New York.

Billingsley, P., 1995. Probability and Measure, third ed. Wiley Ser. Probab. Math. Statist. Wiley, New York.

Blass, A., Frankiewicz, R., Plebanek, G., Ryll-Nardzewski, C., 2001. A note on extensions of asymptotic density. Proc. Amer. Math. Soc. 129 (11), 3313–3320.

Dubey, P., Shapley, L.S., 1994. Noncooperative general exchange with a continuum of traders: Two models. J. Math. Econ. 23 (3), 253–293.

Dubins, L.E., Savage, L.J., 1965. How to Gamble if You Must. Inequalities for Stochastic Processes. McGraw–Hill, New York.

Dunford, N., Schwartz, J., 1958. Linear Operators, Part I. Interscience, New York.

Feldman, M., Gilles, C., 1985. An expository note on individual risk without aggregate uncertainty. J. Econ. Theory 35, 26–32.

Gangopadhyay, S., Rao, B.V., 1999. Completeness of L_1 spaces over finitely additive probabilities. Colloq. Math. 80 (1), 83–95.

Guesnerie, R., 1981. On taxation and incentives: Further reflections on the limits of redistribution. Ecole des hautes études en sciences sociales.

Guesnerie, R., 1995. A Contribution to the Pure Theory of Taxation. Cambridge Univ. Press, Cambridge.

Judd, K.L., 1985. The law of large numbers with a continuum of IID random variables. J. Econ. Theory 35, 19–25.

Kalai, E., 2004. Large robust games. Econometrica 72, 1631–1665.

Khan, M., Sun, Y., 1999. Non-cooperative games on hyperfinite Loeb spaces. J. Math. Econ. 31, 455–492.

Kushner, H.J., 1990. Weak Convergence Methods and Singularly Perturbed Stochastic Control and Filtering Problems. In: Syst. Control: Found. Appl., vol. 3. Birkhäuser Boston, Boston, MA.

Mas-Colell, A., 1984. On a theorem of Schmeidler. J. Math. Econ. 13, 201–206.

Milgrom, P., Weber, R., 1985. Distributional strategies for games with incomplete information. Math. Operations Res. 10, 619–632.

Royden, H.L., 1988. Real Analysis, third ed. Macmillan Publishing Company, New York.

Savage, L.J., 1954. The Foundations of Statistics. John Wiley & Sons, New York.

Schmeidler, D., 1973. Equilibrium points of nonatomic games. J. Statist. Phys. 7, 295–300.

Shiryayev, A.N., 1984. Probability. In: Grad. Texts Math., vol. 95. Springer-Verlag, New York. Translated from Russian by R.P. Boas.

Sun, Y., 2006. The exact law of large numbers via Fubini extension and characterization of insurable risks. J. Econ. Theory 126 (1), 31–69.

Uhlig, H., 1996. A law of large numbers for large economies. Econ. Theory 8, 41–50.