A Theory of Integration & Application to Large Games

Nabil I. Al-Najjar

Kellogg School of Management
Northwestern University

Talk at:
Tel Aviv University
December 2007

Motivation: *Large games*

Skip slide

- Continuum of players $\bar{T} = [0, 1]$
- Countably additive measure λ (uniform for simplicity)
- Each player has just two actions $\{u, d\}$
- Strategy profiles are measurable functions

$$\mu : \bar{T} \to \Delta\{u, d\}$$

- Expected payoff of an action

$$U(a, \mu) = u\left(a, \int_{\bar{T}} \mu \, d\lambda\right)$$

Natural class of games with wide range of applications

Schmeidler '73: these games always have pure strategy eq.

Example: continuum player matching pennies

Skip slide

Payoffs are symmetric and given by:

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

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

Scope of purification

Skip slide

1. Intuitively, a much stronger purification result should hold:

   "every realization of every mixed equilibrium is a pure strategy equilibrium almost surely"

2. This point was made by Kalai (2004) who showed that this is approximately true in large finite games

3. In many games, pure strategies do not exist in continuum games, even though the intuition for the purification result still holds

Example .. *continued*

Skip slide

Consider now the following variant where player t wants to mismatch only players before him:

$$\bar{u}_t(d, \bar{\mu}) = \begin{cases} 1 & \text{if } \int_0^t \bar{\mu}\, d\bar{\lambda} < 0.5, \\ -1 & \text{if } \int_0^t \bar{\mu}\, d\bar{\lambda} > 0.5, \\ 0 & \text{otherwise,} \end{cases}$$

and $\bar{u}_t(d, \bar{\mu}) = -\bar{u}_t(u, \bar{\mu})$.

This game has no pure strategy equilibria, even though the logic of purification ought to continue to hold

Conceptual issues

Skip slide

1. What is the state space on which expected payoffs are defined?
2. Dubey and Shapley: Requiring strategy profiles to be measurable violates the spirit of non-cooperative play

Two Ideas and two papers

Talk is based on two papers:

1. Al-Najjar: "Finitely additive representation of L^p spaces." *Journal of Mathematical Analysis and Applications*. (2007)
2. Al-Najjar: "Large games and the law of large numbers." *Games and Economic Behavior.* (forthcoming).

Integration Theory: *the bottom line*

Take any of the "usual continuum spaces" $\bar{T}$ with countably additive $\bar{\lambda}$ used in economics and game theory

Main points

There is a discrete space T, with a countable number of players, the σ-algebra 2^T, and *finitely* additive probability λ such that

- Integration on T is tractable and well-behaved
- The usual theory of integration on $\bar{T}$ can be isometrically isomorphically mapped *into* integration on T.
- No measurability requirements are needed

Large games: *the bottom line*

Skip slide

Main points

- Any continuum game can be mapped into a discrete large game
- mapping preserves strategy profiles, payoffs, and equilibria
- well defined state space, expected payoffs, etc..
- exact purification holds
- limit theorems

Notation

Skip slide

The continuum:

- $\bar{T} = \{0, 1\}^{\mathcal{N}}$
- $\bar{\mathcal{B}}_k$: the algebra generated by the first k coordinates.
- $\cup_{k=1}^{\infty} \bar{\mathcal{B}}_k$ is an algebra
- $\bar{\mathcal{B}}$ is the induced σ-algebra

Discretization

- T_N be a finite subset of $\bar{T}$, $N = 1, 2 \ldots$
- λ_N is the uniform distribution on T_N
- $T = \cup T_N \subset \bar{T}$ is a countable dense subset
- Σ is the power set on T
- $\mathcal{B}_k = \{B \cap T : B \in \bar{\mathcal{B}}_k\}$

Intervals translate by restriction

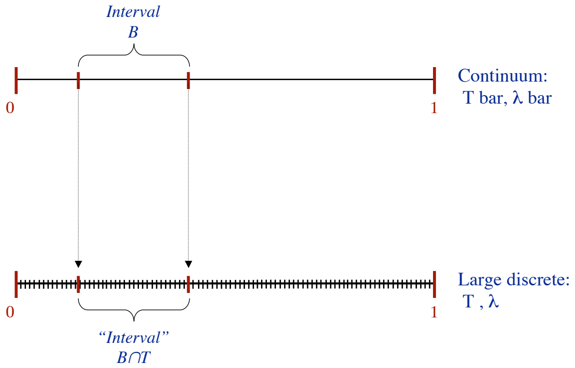

Translate measures of intervals

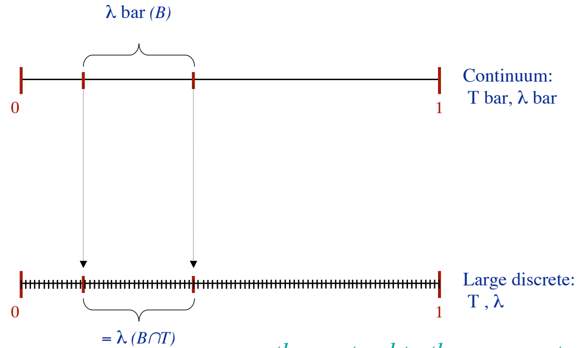

… then extend to the powerset

Step functions and their integrals translate by restriction

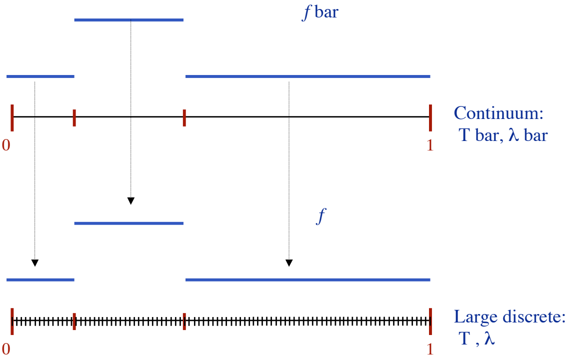

… and so would continuous functions

Simple functions

Skip slide

Definition

$\bar{f} : \bar{T} \to \mathcal{R}$ is *simple* if it is measurable with respect to $\bar{\mathcal{B}}_k$ for some k.

Similar definitions for functions $f : T \to \mathcal{R}$

- Set of all simple functions $\bar{\mathcal{S}}$ (resp. $\mathcal{S}$)
- $\bar{\mathcal{S}}$ and $\mathcal{S}$ are linear spaces
- Define ϕ so that $\phi(\bar{f})$ is the restriction of $\bar{f}$ to T.

Lemma

ϕ is a linear isomorphism from $\bar{\mathcal{S}}$ onto $\mathcal{S}$.

Main embedding theorem: *Preliminaries*

1. $p \in (0, \infty)$
2. Define the $L^p(\lambda)$-norm on functions $f : T \to R$
   - For $p = 1$ we have $\| f \| \equiv \int_T |f| \, d\lambda$.
3. Define $L^p(\bar{\lambda})$-norm on functions $\bar{f} : \bar{T} \to R$ in the same way
4. Let $L^p(\lambda)$ and $L^p(\bar{\lambda})$ be the corresponding spaces of equivalence classes
5. $[f]$ and $[\bar{f}]$ denote the equivalence classes $L^p(\lambda)$ and $L^p(\bar{\lambda})$ resp.

NOTE: Equivalence in $L^p(\lambda)$ acts a little differently from equivalence in $L^p(\bar{\lambda})$!!

Details..

Main embedding theorem

Skip slide

There is a way to construct $\{T_N\}_{N=1}^{\infty}$ such that, for any:

- atomless and countably additive probability measure $\bar{\lambda}$

There is

- a purely finitely additive probability measure λ on (T, Σ)
- $L^p(\lambda)$ is a Banach space
- unique bounded linear operator $\Phi : L^p(\bar{\lambda}) \to L^p(\lambda)$ with

$$\Phi(f) = [\phi(f)], \quad \forall f \in \mathcal{S}.$$

The operator Φ preserves integrals: for every $[\bar{f}] \in L^p(\bar{\lambda})$

$$\int_B \bar{f}\, d\bar{\lambda} = \int_B \Phi(\bar{f})\, d\lambda, \quad \forall B \in \cup_{k=1}^{\infty} \mathcal{B}_k.$$

The green oval is L^p on the continuum; the larger pink oval on the right is the L^p on the discrete space T.

The embedding theorem says that left green oval can be *strictly* embedded into the pink oval.

Question: what do function like μ' correspond to? I later introduce two notions of equivalence to explain this.

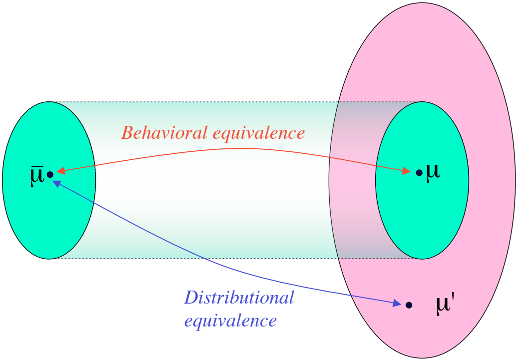

Main embedding theorem *(continued..)*

Skip slide

- Φ is an isometric isomorphism onto a closed linear subspace $H \subset L^p(\lambda)$
- that is, Φ is one-to-one and

$$\| \Phi(\bar{f}) \| = \| \bar{f} \|$$

for every $[\bar{f}] \in L^p(\bar{\lambda})$.

- For $p = 2$, Φ also preserves inner products

Qualifications and Interpretation 1

Skip slide

Note that λ must be purely finitely additive. Details..

From the perspective of the theory of integration, two isometrically isomorphic function spaces are identical.

The map Φ is “canonical” in the sense that it carries

- step functions to step functions
- continuous functions to uniformly continuous functions
- differentiable functions to differentiable functions

Qualifications and Interpretation 2

Skip slide

There is no claim that there is a meaningful way to have a point-to-point map from $\bar{T}$ to T.

- These are very different spaces and there is generally no meaningful point-to-point mapping of subsets

So what is being claimed here?

- The only map claimed is that of equivalence classes of functions
- For example, if A is an arbitrary Borel set in [0,1], there is typically no corresponding event in the discrete space T
- However, to its indicator function $\bar{\chi}_A$ there is a corresponding function χ_A in the discrete model (actually, an L^p equivalence class)

Qualifications and Interpretation 3

Skip slide

- Is $L^p(\lambda)$ strictly larger than $L^p(\bar{\lambda})$?
  - The answer is emphatically: YES.
  - Example: take a typical realizations of an iid process
  - How much larger? *much larger !*
- For the case $p = 2$, $L^2(\bar{\lambda})$ is countable dimensional, while $L^2(\lambda)$ is uncountable dimensional!!
- *..but this seems weird!* Actually, no. Take an i.i.d. process with mean 0.5, then it has uncountably many non-equivalent realizations. The continuum suppresses this richness via measurability.

Idea of the proof

Skip slide

Key issue: We know that $L^p(\bar{\lambda})$ is complete.

But is $L^p(\lambda)$ also complete?

The answer is: Yes, if we construct the space carefully

Construction builds on work in the last few years on the completeness of $L^p(\lambda)$ spaces over fap spaces:

- Gangopadhyay-Rao
- Basile-Bhaskara-Rao
- Blass, Frankiewicz, Plebanek, and Ryll-Nardzewski (2001)

Key points in the construction of λ

Skip slide

1. Construct the finite models T_N so they are dense in $\bar{T}$
2. Define the measure of a subset $A \subset T$ via ultrafilters $\mathcal{U}$ on the integers
3. That is, define
$$\lambda(A) = \lim_{N \to \infty} \lambda_N(A)$$
if this limit exists, otherwise define it as
$$\lambda(A) = \mathcal{U}\text{–}\lim_{N \to \infty} \lambda_N(A)$$
4. Completeness of $L^p(\lambda)$ obtains if
$$\frac{\#(T_N - T_{N-1})}{\# T_N} \to 1, \text{ as } N \to \infty.$$

Key points in the embedding

Skip slide

I want to find a way to map the function $\bar{f}$ to a function f in $L^p(\lambda)(\lambda)$

If $\bar{f}$ is simple or continuous, then this immediate..

If not, take a sequence of simple functions $\bar{f}_n$ that converges to $\bar{f}$; note that this sequence is Cauchy

But since Φ is an isometry, we obtain a Cauchy sequence in $L^p(\lambda)(\lambda)$, which must have limit by completeness

Large games: *strategy profiles*

Finite set of actions A; mixed actions Δ^A

Definitions

Pure strategy profile is *any* functions from T to A.

Mixed strategy profile is *any* functions from T to Δ^A.

Payoffs

Continuum Game

$$\bar{U}(t,\bar{s}) \equiv \bar{u}\left(t,\bar{s}(t),\int_{\bar{T}}\bar{s}\,d\bar{\lambda}\right)$$

where $\bar{u} : \bar{T} \times A \times \Delta^A \rightarrow \mathcal{R}$ is (jointly) continuous.

Discrete Large Game

$$U(t,s) \equiv u\left(t,s(t),\int_{T}s\,d\lambda\right)$$

where $u : T \times A \times \Delta^A \rightarrow \mathcal{R}$ is uniformly (jointly) continuous.

Expected Payoffs and Distributional Strategies

Expected payoffs in the discrete game are defined in the usual way:

$$U(t,\mu) \equiv \int_S u(t, s(t), \textstyle\int_T s\, d\lambda)\; dP^{\mu}.$$

In continuum games Schmeidler, Mas-Colell and Milgrom and Roberts *assume* expected payoff to be:

$$\bar{U}(t,\bar{\mu}) \equiv \bar{u}\left(t, \bar{\mu}(t), \textstyle\int_{\bar{T}} \bar{\mu}\; d\bar{\lambda}\right)$$

Because $\int_{\bar{T}} \bar{s}\, d\bar{\lambda}$ cannot be meaningfully evaluated. This *distributional representation* is a short-cut based on an intuitive LLN.

Law of Large Numbers

- let S be the space of all pure strategy profiles
- $\mathcal{S}$ is the (usual) σ-algebra generated by all events of the form $\{s : s(t) = a\}$
- Any mixed profile μ uniquely defines a *countably additive* probability distribution P^μ on $(S, \mathcal{S})$.

Theorem

For every strategy profile μ

$$P^\mu \left\{ s \in S : \int_B s \, d\lambda = \int_B \mu \, d\lambda, \ \ \forall B \in \cup_{k=1}^{\infty} \mathcal{B}_k \right\} = 1.$$

Interpretation of Distributional Strategies

$$
\begin{array}{ccc}
\bar{\mu} & \xrightarrow{\qquad\qquad ???} & \int_{\bar{T}} \bar{\mu}\, d\bar{\lambda} \\
\downarrow & & \Big\| \; P^{\mu} - a.s. \\
\mu & \xrightarrow{\qquad\qquad\quad} & \int_{T} s\, d\lambda
\end{array}
$$

Equilibrium

Definition of Equilibrium

μ is an *equilibrium* if for every $\epsilon > 0$,

$$\lambda\{t \in T : \mu(t) \in \mathrm{Br}_\epsilon(t, \mu)\} = 1.$$

Theorem

For every equilibrium in the continuum game there is a behaviorally equivalent equilibrium in the large discrete game.

Behavioral equivalence

When are two profiles $\bar{\mu}$, μ 'equivalent?'

$\bar{\mu}$ and μ are *behaviorally equivalent* if for every $\epsilon > 0$ there is a simple profile μ_ϵ such that

$$\int_T |\mu - \mu_\epsilon| d\lambda < \epsilon \ \ \& \ \ \int_{\bar{T}} |\bar{\mu} - \mu_\epsilon| d\bar{\lambda} < \epsilon$$

Representation Theorem

For every profile $\bar{\mu}$ there is a behaviorally equivalent profile μ.
If $\bar{\mu}$ is an equilibrium, then so would μ.

Purification

Theorem

Every realization of every equilibrium μ is a pure strategy equilibrium P^{μ}-almost surely.

For a non-degenerate μ, realizations do NOT have a behaviorally equivalent representation in the continuum, almost surely.

So do these pure strategy realizations correspond to anything?

Distributional equivalence

A weaker notion of equivalence:

μ and μ' are *distributionally equivalent* if

$$\int_B \mu \, d\lambda = \int_B \mu' \, d\lambda, \quad \forall B \in \cup_{k=1}^{\infty} \mathcal{B}_k$$

In the continuum, behavioral and distributional equivalence are the same.

Distributional Representation and Spurious Randomness

Theorem

For every pure strategy s there is a distributionally equivalent *mixed* profile $\bar{\mu}$:

$$\int_B s \, d\lambda = \int_B \bar{\mu} \, d\bar{\lambda}, \quad \forall B \in \cup_{k=1}^{\infty} \mathcal{B}_k$$

If s is an equilibrium, then so is $\bar{\mu}$.

Key Point: *Behavior looks random, but this is an artifact of countable additivity and measurability!!* In the discrete game, behavior is deterministic.

Large Finite games

The LLN paper also explores the relationship with large finite games. Essentially, it is shown that:

- The restriction of any equilibrium in the discrete large game is an ϵ-equilibria of some large finite game.
- Every sequence of ϵ-equilibria in large finite games 'corresponds' to an equilibrium of the discrete large game.

Looking forward

Skip slide

Here is a little teaser: For each t,

- $s(t, \omega) \in \{-1, 1\}$ is a zero-mean random variable
- independence across the t's
- For each t, $T_{N(t)}$ is the first finite model in which t appears

Fix $b > a \geq 0$; what does the (standard) LLN tell us about the random variable:

$$\omega \mapsto \int_a^b s(t, \omega)\, d\lambda \quad ?$$

Answer: *obvious!*

Teaser ..continued

▸ Skip slide

How about the random variable:

$$\omega \mapsto \int_a^b \left[\sqrt{\#T_{N(t)}}\right] s(t,\omega)\, d\lambda \quad ?$$

- The CLT guarantees that this is normally distributed
- Independence guarantees independence of increments
- Standard arguments guarantee that

$$b \mapsto \int_0^b \left[\sqrt{\#T_{N(t)}}\right] s(t,\omega)\, d\lambda$$

is uniformly continuous, for a.e. ω

Time

Skip slide

The above is a rough construction of a driftless Brownian Motion on a discrete time grid T!!

Basic properties carry through because, in the standard theory, they are proved along discrete time approximation of continuous time, and so hold here too.

But what does this buy us? In standard model of the BM, we cannot talk about what happens instant by instant. For example, at each instant the the $s(t, \omega)$ might as well be any distribution with zero mean. This additional structure is lost when looking at continuous time. On a discrete grid, we can do the usual continuous time manipulation AND study the process instant by instant.

Appendix

Equivalence

Define $f(t) = \frac{1}{N}$ for $t \in T_N - T_{N-1}$. Then $f(t) > 0$ for all t, but $\lambda\{t : f(t) < \epsilon\} = 0$ for every $\epsilon > 0$. Essentially, this is a function that is infinitesimally close to zero.

It is easy to see that $\int_T f \, d\lambda = 0$ even though it is strictly positive everywhere.

Note that $\{t : f(t) > \frac{1}{N}\}$ is always finite and hence has λ-measure zero for every N. But $\cup \{t : f(t) > \frac{1}{N}\} = T$ and hence has measure 1.

This cannot happen under countable additivity.

back..

Integration wrt Finitely Additive Measures

See the paper!!

Pure Finite Additivity

For any finite A, we have $\lambda_N(A) \to 0$ hence $\lambda(A) = 0$. This is a consequence of the condition (in the paper):
For every $A_1, \ldots, A_L \subset T$ there is a subsequence $\{N_k\}$ such that

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

Thus, $\lambda(t) = 0$ for every t, T is countable, yet $\lambda(T) = 1$.

back..