Continuous games I

Existence of equilibria and separable games

Tomáš Kroupa

Faculty of Electrical Engineering
Artificial Intelligence Center

CZECH TECHNICAL UNIVERSITY IN PRAGUE

Aperitif: Can you solve this zero-sum game over $[0,1]$?

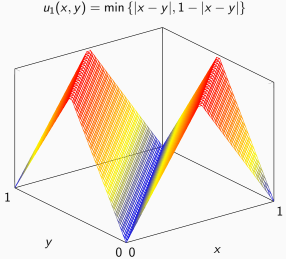

$$S_1 = S_2 = [0,1]$$
$$u_1(x,y) + u_2(x,y) = 0$$

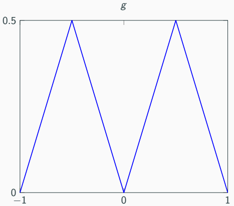

Using the reflection symmetry

$$u_1(x,y) = g(x-y)$$
$$u_1(x,y) = u_1(1-x, 1-y)$$

Every optimal strategy must also be reflection-invariant $\Rightarrow$ uniform distribution

Uniform distributions are optimal

Since the support of uniform distribution is $[0,1]$, it suffices to check that every $x \in [0,1]$ is a best response:

$$\int_0^1 u_1(x,y)\, dy = \int_0^1 g(t)\, dt = \frac{1}{4}$$

General strategic games

Definition

- Player set $N = \{1, \ldots, n\}$
- Strategy set S_i, $\forall i \in N$
- Utility function $u_i \colon S \to \mathbb{R}$, $\forall i \in N$, where $S = S_1 \times \cdots \times S_n$

Strategic game is a tuple $(N, (S_i)_{i \in N}, (u_i)_{i \in N})$.

First questions

1. Applications?
2. What is the concept of mixed strategy?
3. Existence of Nash equilibria?
4. Computing Nash equilibria?

First answers, briefly

1. Strategic situations in which agents decide about a quantity, timing, position, or real-valued parameters of some system:
   - Economic market, Cournot duopoly competition
   - Games of timing–duels
   - Position games
   - Adversarial machine learning games
2. Mixed strategy of player i = probability measure on S_i
3. Existence of Nash equilibria can be guaranteed in a broader class of games than the class of games with finite strategy sets S_i
4. Computing Nash equilibria is a difficult infinite-dimensional optimization problem even for two-person zero-sum games

Our setting – continuous games

Definition

A strategic game $(N, (S_i)_{i \in N}, (u_i)_{i \in N})$ is continuous if

- Each strategy set S_i is a **compact** subset of $\mathbb{R}^{d_i}$ for some $d_i \in \mathbb{N}$
- Each utility function $u_i \colon S \to \mathbb{R}$ is **continuous**

Continuous games include the following important families of games:

Finite games $\subset$ Polynomial games $\subset$ Separable games

Hunting for mixed strategies with small support

- In general, equilibrium strategies in a continuous game can be arbitrarily complicated probability measures!
- Occam's razor or practical applications ask for equilibrium strategies having a **small** (ideally **finite**) support
- Upper bounds on the supports of equilibria exist in finite games:

Lipton et al. (2003)

Consider a two-person game given by payoff matrices $\mathbf{A}_1$ and $\mathbf{A}_2$.
For any Nash equilibrium (μ_1, μ_2) there is a Nash equilibrium

- giving the same payoff to both players as (μ_1, μ_2) and
- in which each player i mixes at most (rank $\mathbf{A}_i + 1$) pure strategies.

Main theme of this course

For a given continuous game

1. Decide existence of a mixed strategy equilibrium in which each player randomizes among **finitely-many** pure strategies only
2. If such an equilibrium exists, then
   - Compute the bounds on the support of equilibrium strategies
   - Compute at least one such equilibrium!

Agenda

1. Short primer on measure theory
2. Continuous games and existence results
3. Separable games

Short primer on measure theory

Probability measures

Finite probability spaces

Ω finite set

$\mathcal{P}(\Omega)$ powerset of Ω

μ probability measure $\mathcal{P}(\Omega) \to [0,1]$

Probability measure is given by $p\colon \Omega \to [0,1]$ such that $\sum_{\omega \in \Omega} p(\omega) = 1$

Infinite probability spaces

Ω compact subset of $\mathbb{R}^n$ for some $n \in \mathbb{N}$

$\mathcal{B}(\Omega)$ σ-algebra of Borel measurable subsets of Ω

μ probability measure $\mathcal{B}(\Omega) \to [0,1]$

Classes of probability measures

Support of μ is the smallest compact set $K \subseteq \Omega$ such that $\mu(K) = 1$.

Example (Dirac)

Let $x \in \Omega$. Probability measure δ_x has the support $\operatorname{spt} \mu = \{x\}$:

$$\delta_x(A) = \begin{cases} 1 & x \in A \\ 0 & x \notin A \end{cases}$$

Example (Atomic)

Let $x_1, \dots, x_k \in \Omega$ and $p_1, \dots, p_k > 0$ satisfy $\sum_{i=1}^k p_i = 1$.
Probability measure $\sum_{i=1}^k p_i \cdot \delta_{x_i}$ is supported by $\{x_1, \dots, x_k\}$.

Example (Limit of atomic measures)

$\mu = \mathrm{w}^* \lim_{n\to\infty} \mu_n$ with each μ_n atomic

Lebesgue integral $\int_\Omega f \ d\mu$

The standard way to integrate a continuous (even measurable) function $f : \Omega \to \mathbb{R}$ with respect to a probability measure μ over Ω.

Example

- If $\mu = \delta_x$, then
$$\int_\Omega f \ d\delta_x = f(x)$$
- If $\mu = \sum_{i=1}^k p_i \cdot \delta_{x_i}$, then
$$\int_\Omega f \ d\mu = \sum_{i=1}^k p_i \cdot f(x_i)$$
- If $\mu = \mathrm{w}^* \lim_{n\to\infty} \mu_n$ with each μ_n atomic, then
$$\int_\Omega f \ d\mu = \lim_{n\to\infty} \int_\Omega f \ d\mu_n$$

Sets of probability measures – comparison

$\Delta(\Omega)$	Set of probability measures over Ω
$\operatorname{ext} C$	Extreme points of a set C

Property of $\Delta(\Omega)$	Ω	
	finite	compact
Convex	✓	✓
Compact	✓	✓
$\operatorname{ext}\Delta(\Omega)$ is finite	✓	no
$\operatorname{ext}\Delta(\Omega) =$ Dirac measures	✓	✓
$\operatorname{ext}\Delta(\Omega)$ is compact	✓	✓
$\operatorname{conv}\operatorname{ext}\Delta(\Omega) = \Delta(\Omega)$	✓	no
$\operatorname{conv}\operatorname{ext}\Delta(\Omega)$ is compact	✓	no
Continuous map $\Delta(\Omega) \to \mathbb{R}$ attains extrema	✓	✓

Existence of equilibria

Mixed strategies and Nash equilibria in a continuous game

- Mixed strategy of player i is a probability measure $\mu_i \in \Delta_i := \Delta(S_i)$
- Expected utility of player i for $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n) \in \Delta := \bigtimes_{i \in N} \Delta_i$ is

$$U_i(\boldsymbol{\mu}) := \int_S u_i \, d(\mu_1 \times \cdots \times \mu_n)$$

Definition

$\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n) \in \Delta$ is a Nash equilibrium if $\forall i \in N \;\; \forall \sigma_i \in \Delta_i$

$$U_i(\sigma_i, \boldsymbol{\mu}_{-i}) \leq U_i(\boldsymbol{\mu})$$

Nash equilibrium is atomic if each μ_i is atomic.

Existence of equilibria

Glicksberg's theorem (1952)

Any continuous game has a Nash equilibrium in mixed strategies.

- For any $\varepsilon > 0$ find $\delta > 0$ satisfying $\forall x_i, x_i' \in S_i$ and $\forall \boldsymbol{x}_{-i} \in S_{-i}$

$$\|x_i - x_i'\| < \delta \qquad \Rightarrow \qquad |u_i(x_i, \boldsymbol{x}_{-i}) - u_i(x_i', \boldsymbol{x}_{-i})| < \varepsilon$$

- We can choose a **finite** subset $T_i \subseteq S_i$ such that for all $x_i \in S_i$, there is $x_i' \in T_i$ with $\|x_i - x_i'\| < \delta$
- The **finite game** $(N, (T_i)_{i \in N}, (u_i | T_i)_{i \in N})$ has an equilibrium
- This is an atomic ε-Nash equilibrium of the continuous game
- For any $\varepsilon_k \to 0$, the sequence of ε_k-Nash equilibria converges to a Nash equilibrium of the continuous game

Characterization of Nash equilibrium $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_n)$

No deviation by pure strategies

$$U_i(x_i, \boldsymbol{\mu}_{-i}) \leq U_i(\boldsymbol{\mu}) \qquad \forall i \in N \ \ \forall x_i \in S_i$$

Support strategies are among best response strategies

$$\operatorname{spt} \mu_i \ \subseteq \ \operatorname*{arg\,max}_{x_i \in S_i} \ U_i(x_i, \boldsymbol{\mu}_{-i}) \qquad \forall i \in N$$

For two-person zero-sum games only

$$\max_{\mu_1 \in \Delta_1} \min_{\mu_2 \in \Delta_2} U_1(\mu_1, \mu_2) = \min_{\mu_2 \in \Delta_2} \max_{\mu_1 \in \Delta_1} U_1(\mu_1, \mu_2)$$

Discontinuous game with no solution

Example (M. Dresher - Games of Strategy)

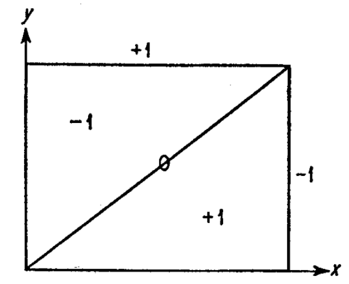

For any $0 < \varepsilon < \frac{1}{2}$:

$$\sup_{\mu_1 \in \Delta_1} \inf_{\mu_2 \in \Delta_2} U_1(\mu_1, \mu_2) \leq \varepsilon - 1 \qquad \text{and} \qquad \inf_{\mu_2 \in \Delta_2} \sup_{\mu_1 \in \Delta_1} U_1(\mu_1, \mu_2) \geq 1 - \varepsilon$$

Discontinuous game with a solution

Example (G. Owen - Game theory)

Two generals commanding an equal number of units fight to take 3 battlefields. The side having more units at a given field will win it. **Assumption**: the two armies are infinitely divisible and the payoff is the number of captured battlefields.

$$
\begin{aligned}
&S_1 = S_2 = \left\{ \boldsymbol{x} \in \mathbb{R}^3 \mid x_1, x_2, x_3 \geq 0,\ x_1 + x_2 + x_3 = 1 \right\} \\
&u_1(\boldsymbol{x}, \boldsymbol{y}) = \operatorname{sgn}(x_1 - y_1) + \operatorname{sgn}(x_2 - y_2) + \operatorname{sgn}(x_3 - y_3) \\
&u_2(\boldsymbol{x}, \boldsymbol{y}) = -u_1(\boldsymbol{x}, \boldsymbol{y})
\end{aligned}
$$

An optimal solution is the density function

$$
f(\boldsymbol{x}) = \begin{cases} \frac{9}{2\sqrt{1-9\|\boldsymbol{x}-\boldsymbol{c}\|^2}} & \|\boldsymbol{x} - \boldsymbol{c}\| \leq \frac{1}{3} \\ 0 & \text{otherwise} \end{cases} \qquad \text{where } \boldsymbol{c} = \left(\tfrac{1}{3}, \tfrac{1}{3}, \tfrac{1}{3}\right)
$$

Continuous game with a bizarre solution

Example (Gross 1952; Karlin 1959)

Two-person zero-sum game with strategy sets $S_1 = S_2 = [0, 1]$ and a utility function of the first player defined by

$$u_1(x, y) = \left(y - \tfrac{1}{2}\right)\left[\frac{1 + (x - \frac{1}{2})(y - \frac{1}{2})^2}{1 + (x - \frac{1}{2})^2(y - \frac{1}{2})^4} - \frac{1}{1 + (\frac{x}{3} - \frac{1}{2})^2(y - \frac{1}{2})^4}\right]$$

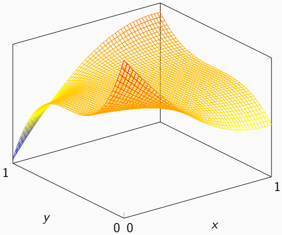

The unique equilibrium strategy for each player is the Cantor distribution.

Why solving continuous games is hard

- Global maximization of a polynomial is hard (1-player game)
- Solution of a particular continuous game is a combination of heuristics and insight into the special structure of the game
- Selected families of continuous games have special equilibria:
    - Convex/concave games
    - Games with bell-shaped utility functions
    - Games invariant under symmetries
    - Games of timing

Separable games

Separable games

A continuous game is separable if there exist $m_1, \ldots, m_n \in \mathbb{N}$ and, for each player $i \in N$,

- continuous functions $f_i^1, \ldots, f_i^{m_i} \colon S_i \to \mathbb{R}$
- real coefficients $a_i^{\boldsymbol{\alpha}}$, for all $\boldsymbol{\alpha} = (\alpha_1, \ldots, \alpha_n) \in [m_1] \times \cdots \times [m_n]$,

such that each utility function u_i is of the form

$$u_i(\mathbf{x}) = \sum_{\boldsymbol{\alpha} \in [m_1] \times \cdots \times [m_n]} a_i^{\boldsymbol{\alpha}} \cdot f_1^{\alpha_1}(x_1) \cdots f_n^{\alpha_n}(x_n), \qquad \mathbf{x} \in S$$

Example

- Every finite game
- Every polynomial game
- Two-player zero-sum game with $u_1(x_1, x_2) = x_1 \sin x_2 + x_1 e^{x_2} + 3x_1^2$

Equivalence of mixed strategies

Mixed strategies $\mu_i, \sigma_i \in \Delta_i$ of a player $i \in N$ in a separable game are

- Payoff equivalent (PE) if, for all $j \in N$ and all $\boldsymbol{x}_{-i} \in S_{-i}$,

$$U_j(\mu_i, \boldsymbol{x}_{-i}) = U_j(\sigma_i, \boldsymbol{x}_{-i})$$

- Moment equivalent (ME) if, for all $\alpha \in [m_i]$,

$$\int_{S_i} f_i^\alpha \, d\mu_i = \int_{S_i} f_i^\alpha \, d\sigma_i$$

For any separable game

1. (ME) implies (PE)
2. If $(\mu_1, \ldots, \mu_n)$ is an equilibrium and each μ_i is (PE) to some σ_i, then $(\sigma_1, \ldots, \sigma_n)$ is also an equilibrium

Atomic equilibria in separable games

Claim

In a separable game any mixed strategy of player i is moment equivalent to an atomic mixed strategy over $\leq m_i + 1$ pure strategies.

The claim immediately implies

Theorem (Karlin, 1959)

Every separable game has an atomic equilibrium in which each player i mixes at most $m_i + 1$ pure strategies.

The main idea of the proof can be shown for polynomial games.

Special case

Example (Polynomial game)

$$u_1(x,y) = 2xy + 3y^3 - 2x^3 - x - 3x^2y^2$$
$$u_2(x,y) = 2x^2y^2 - 4y^3 - x^2 + 4y + x^2y \qquad x, y \in [-1,1]$$

$m_1 = m_2 = 4$, the functions are $1, x, x^2, x^3$ and $1, y, y^2, y^3$, respectively

The idea is to replace the infinite-dimensional

set of all mixed strategies Δ_i

with the finite-dimensional

moment space.

Moment space

Definition

Moment map $M\colon \Delta_i \to \mathbb{R}^4$ is given by

$$M(\mu) = \left(\int_{-1}^{1} x^0 \, d\mu, \ldots, \int_{-1}^{1} x^3 \, d\mu\right) \qquad \forall \mu \in \Delta_i$$

Moment space is the set $M(\Delta_i) \subseteq \mathbb{R}^4$.

Moment space can be identified with $\text{conv}\left\{(x, x^2, x^3) \mid x \in [-1, 1]\right\}$

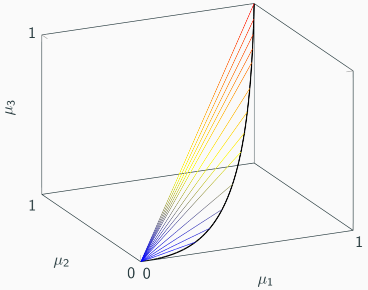

Remarks

- The bound $m_i + 1$ is not tight, which motivated the introduction of the notion of rank for any continuous game (Stein et al., 2008)
- For $\varepsilon > 0$, there exists an algorithm computing an ε-equilibrium of a two-person separable game with Lipschitz utility functions
- For two-person zero-sum polynomial games, an equilibrium can be found by solving a single semidefinite program (Parrilo, 2006)

References

M. Dresher.
Games of strategy: Theory and applications.
Prentice-Hall Applied Mathematics Series. Prentice-Hall Inc., Englewood Cliffs, N.J., 1961.

N. D. Stein.
Characterization and computation of equilibria in infinite games.
Master's thesis, MIT, http://hdl.handle.net/1721.1/40326, 2007.

N. D. Stein, A. Ozdaglar, and P. A. Parrilo.
Separable and low-rank continuous games.
International Journal of Game Theory, 37(4):475–504, 2008.