

ΣΤΟΙΧΕΙΑ ΘΕΩΡΙΑΣ ΠΑΙΓΝΙΩΝ ΚΑΙ ΛΗΨΗΣ ΑΠΟΦΑΣΕΩΝ

ΔΙΑΛΕΞΗ 7: ΣΥΝΕΧΗ ΠΑΙΓΝΙΑ

Παναγιώτης Μερτικόπουλος

Εθνικό και Καποδιστριακό Πανεπιστήμιο Αθηνών

Τμήμα Μαθηματικών



Χειμερινό Εξάμηνο, 2023–2024



Outline

- 1 A modern game
- 2 Definitions and further examples
- 3 Nash equilibrium and characterizations
- 4 Concave games



Example 1: Spot the fake

Which person is real?





Example 1: Spot the fake

Which person is real?



◆ Spoiler: <https://thispersondoesnotexist.com>



Deep learning 101

The "hello world" of deep learning: how to recognize a hand-written digit?



Figure: A sample from the MNIST database


Neural networks: use labeled data to infer hidden structures ("learn")




The perceptron

A digit recognition *perceptron*:

◆ McCulloch & Pitts (1943)



1. x_1 : is image intensity above 10%?
2. x_2 : does image contain  ?
3. x_3 : does image contain a loop?


Output: $y = f(\sum_i w_i x_i)$: if $y = 1$, image depicts a 0; else image does not depict a 0




The perceptron

A digit recognition *perceptron*:

→ McCulloch & Pitts (1943)



1. x_1 : is image intensity above 10%?
2. x_2 : does image contain ?
3. x_3 : does image contain a loop?

Output: $y = f(\sum_i w_i x_i)$: if $y = 1$, image depicts a 0; else image does not depict a 0


✓ *Simple, but not simplistic:* much better than guessing at random!




The perceptron

A digit recognition *perceptron*:

→ McCulloch & Pitts (1943)



1. x_1 : is image intensity above 10%?
2. x_2 : does image contain ?
3. x_3 : does image contain a loop?

Output: $y = f(\sum_i w_i x_i)$: if $y = 1$, image depicts a 0; else image does not depict a 0

✓ **Simple, but not simplistic:** much better than guessing at random!

✗ What is f ?

✗ How do we extract x_i ?




How would this work in practice?






How would this work in practice?






How would this work in practice?



$$\text{Out}(X) = \sigma(W \cdot X + b)$$




How would this work in practice?






Deep neural networks






Deep neural networks






Deep neural networks



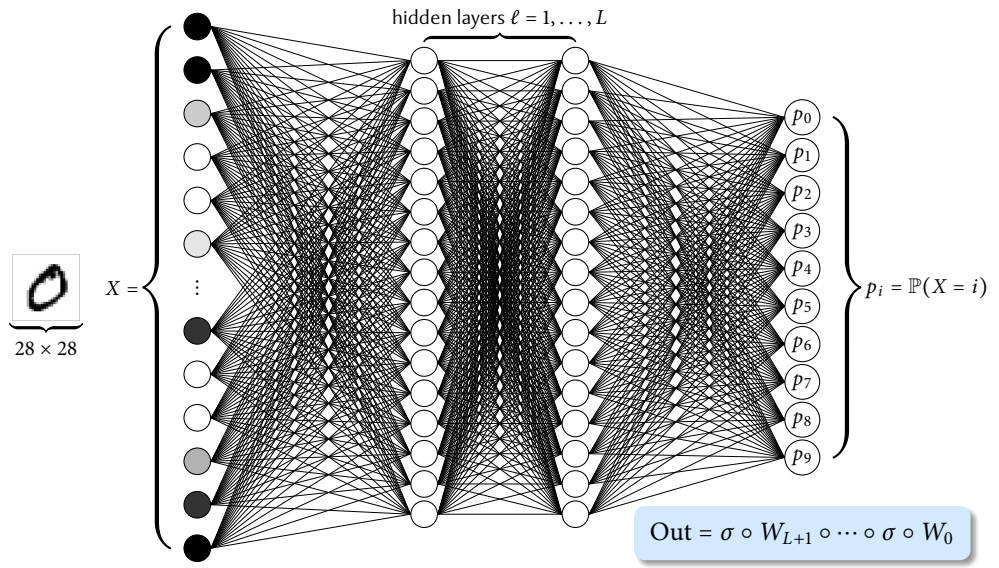


Deep neural networks





Deep neural networks






The schematics of GANs

Z_i noise




The schematics of GANs



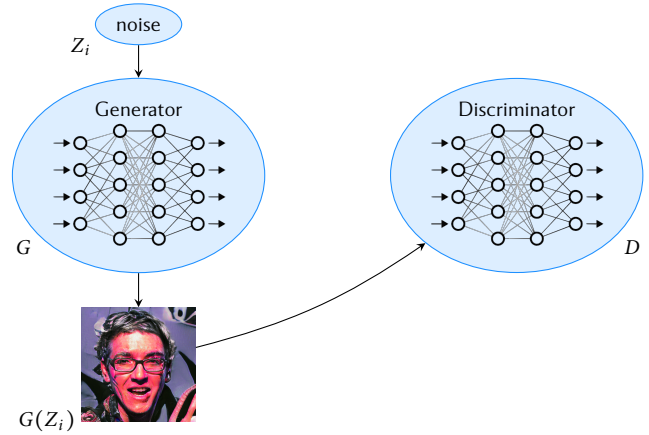


The schematics of GANs



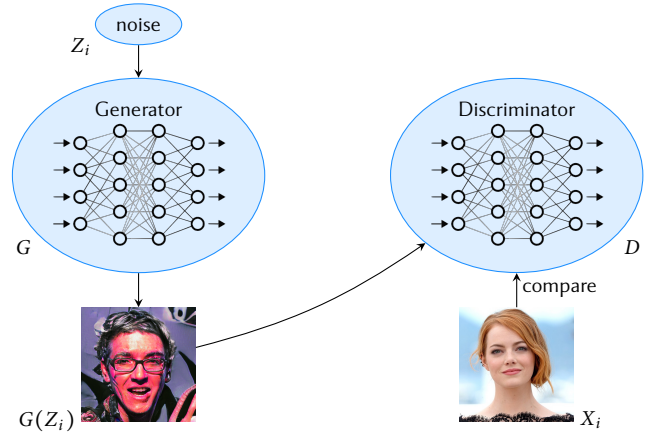


The schematics of GANs



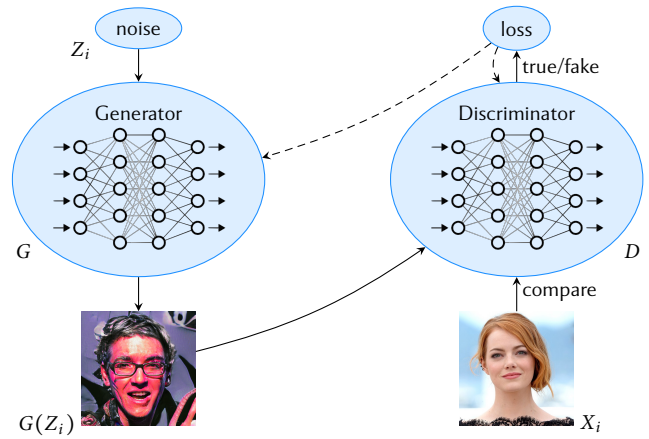


The schematics of GANs





The schematics of GANs



Model likelihood:
$$\ell(G, D) = \prod_{i=1}^N D(X_i) \times \prod_{i=1}^N (1 - D(G(Z_i)))$$



GAN training

How to find good generators (G) and discriminators (D)?

Discriminator: maximize (log-)likelihood estimation

$$\max_{D \in \mathcal{D}} \log \ell(G, D)$$

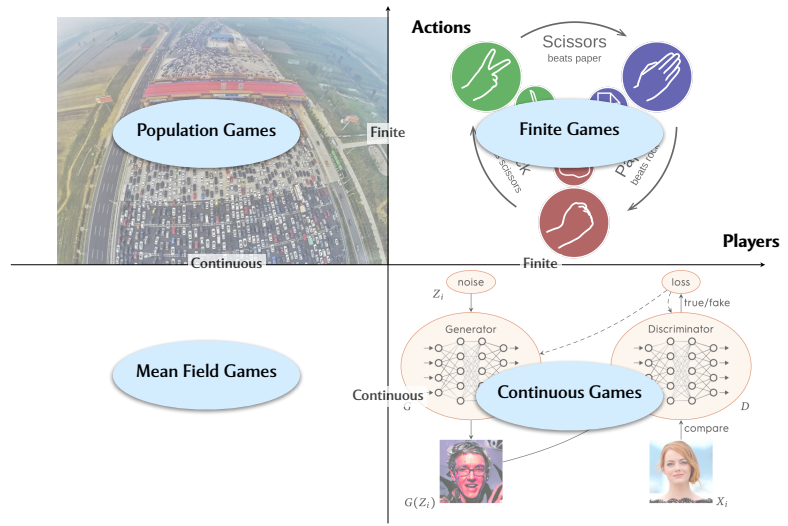
Generator: minimize the resulting divergence

$$\min_{G \in \mathcal{G}} \max_{D \in \mathcal{D}} \log \ell(G, D)$$

A (very complex) zero-sum game!



Taxonomy





Outline

- ① A modern game
- ② Definitions and further examples
- ③ Nash equilibrium and characterizations
- ④ Concave games



Continuous games

Definition (Continuous games)

A **continuous game** is a collection of the following primitives:


- ▶ A finite set of **players** $\mathcal{N} = \{1, \dots, N\}$
- ▶ A closed convex set of **actions** $\mathcal{X}_i \subseteq \mathbb{R}^{d_i}$ for each player $i \in \mathcal{N}$
- ▶ A **payoff function** $u_i: \mathcal{X} := \prod_j \mathcal{X}_j \rightarrow \mathbb{R}$ for each player $i \in \mathcal{N}$
(usually differentiable)

A continuous game with primitives as above will be denoted by $\mathcal{G} \equiv \mathcal{G}(\mathcal{N}, \mathcal{X}, u)$.

finite: $\Gamma \equiv \Gamma(\mathcal{N}, A, u)$

Notes:

- ▶ **Generality:** \mathcal{X}_i could be more general, but almost always closed and convex in practice
- ▶ **Differentiability:** convenient to assume u_i differentiable in an open neighborhood of \mathcal{X} in \mathbb{R}^d






Example: Mixed extensions

Playing with mixed strategies:

- ▶ Players: $\mathcal{N} = \{1, \dots, N\}$
- ▶ Pure strategies: $a_i \in \mathcal{A}_i$
- ▶ Mixed strategies: $x_i \in \mathcal{X}_i \equiv \Delta(\mathcal{A}_i)$
- ▶ Choose mixed strategy $x_i \in \mathcal{X}_i$
- ▶ Choose action $a_i \sim x_i$
- ▶ Get payoff $u_i(a_i; a_{-i})$



Example

The **mixed extension** $\Delta(\Gamma)$ of a finite game $\Gamma \equiv \Gamma(\mathcal{N}, \mathcal{A}, u)$ can be seen as a continuous game with

- ▶ Players: $\mathcal{N} = \{1, \dots, N\}$
- ▶ Action sets: $\mathcal{X}_i = \Delta(\mathcal{A}_i)$
- ▶ Payoff functions: $u_i(x) = \mathbb{E}_{a \sim x}[u_i(a)] = \sum_{a_1 \in \mathcal{A}_1} \dots \sum_{a_N \in \mathcal{A}_N} x_{1,a_1} \dots x_{N,a_N} u_i(a_1, \dots, a_N)$



Example: Cournot competition

A **Cournot oligopoly** consists of the following elements:

- ▶ N firms compete in a market for a given product # no product differentiation
- ▶ Each firm $i = 1, \dots, N$ can produce up to C_i of the good in question # C_i = production capacity
- ▶ Production has a cost of $c_i > 0$ per unit
- ▶ The good is priced as a function $P(x_{\text{tot}})$ of the total production $x_{\text{tot}} = \sum_j x_j$ # usually $P(x) = a + bx$
- ▶ The utility of each firm is given by $u_i(x) = x_i P(x_{\text{tot}}) - c_i x_i$

Example (Cournot competition)

A **Cournot oligopoly** can be seen as a continuous game with


- ▶ Players: $\mathcal{N} = \{1, \dots, N\}$
- ▶ Action sets: $\mathcal{X}_i = [0, C_i]$
- ▶ Payoff functions: $u_i(x) = x_i P(x_{\text{tot}}) - c_i x_i$ # utility = revenue - cost



Example: Resource allocation / Kelly auctions

Proportionally fair resource allocation

• Tullock (1967, 1980); Kelly et al. (1998):






Example: Resource allocation / Kelly auctions

Proportionally fair resource allocation

• Tullock (1967, 1980); Kelly et al. (1998):

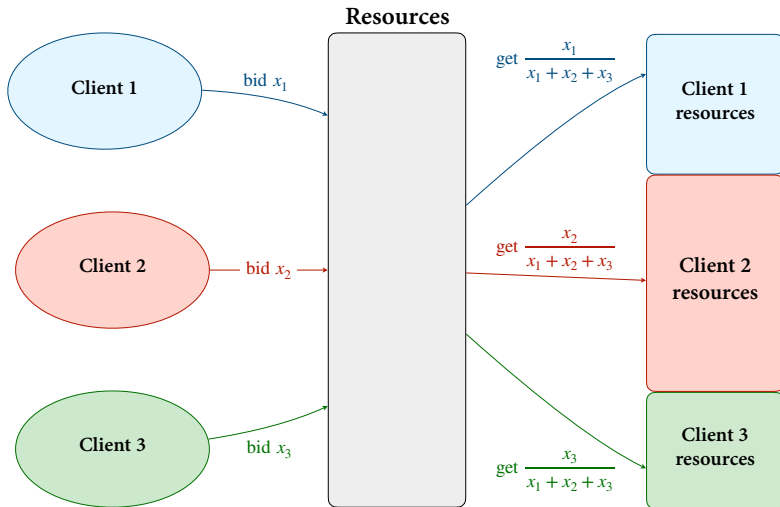




Example: Resource allocation / Kelly auctions

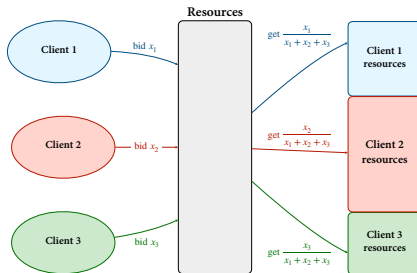
Proportionally fair resource allocation

• Tullock (1967, 1980); Kelly et al. (1998):





Example: Resource allocation / Kelly auctions



Example (Resource allocation)

A **Kelly auction** is a continuous game with

- ▶ Players: $\mathcal{N} = \{1, \dots, N\}$
- ▶ Action sets: $\mathcal{X}_i = [0, B_i]$
- ▶ Payoff functions: $u_i(x) = \frac{g_i x_i}{x_1 + \dots + x_N} - x_i$

B_i = maximum budget of player i

g_i = marginal profit from the good



Outline

- ① A modern game
- ② Definitions and further examples
- ③ Nash equilibrium and characterizations
- ④ Concave games



Nash equilibrium

Nash equilibrium

Let $\mathcal{G} \equiv \mathcal{G}(\mathcal{N}, \mathcal{X}, u)$ be a continuous game. An action profile $x^* = (x_1^*, \dots, x_N^*)$ is a **Nash equilibrium** of \mathcal{G} if

$$u_i(x_i^*; x_{-i}^*) \geq u_i(x_i; x_{-i}^*) \quad \text{for all } x_i \in \mathcal{X}_i \text{ and all } i \in \mathcal{N}. \quad (\text{NE})$$

Intuition:

- ▶ **Stability:** no player has an incentive to deviate
- ▶ **Unilateral resilience:** stable against *individual* player deviations, not multi-player ones



Characterization: best responses

Definition (Best responses)

The action $x_i^* \in \mathcal{X}_i$ is a *best response* to $x_{-i} \in \mathcal{X}_{-i}$ if

$$u_i(x_i^*; x_{-i}) \geq u_i(x_i; x_{-i}) \quad \text{for all } x_i \in \mathcal{X}_i.$$

or, equivalently, if

$$x_i^* \in \arg \max_{x_i \in \mathcal{X}_i} u_i(x_i; x_{-i})$$

The set-valued function $\text{BR}_i: \mathcal{X}_{-i} \rightrightarrows \mathcal{X}_i$ given by

$$\text{BR}_i(x_i) := \arg \max_{x_i \in \mathcal{X}_i} u_i(x_i; x_{-i})$$

is called the *best response correspondence* of player i . Collectively, we also let

$$\text{BR}(x) = \prod_{i \in I} \text{BR}_i(x_i)$$



Characterization: best responses

Definition (Best responses)

The action $x_i^* \in \mathcal{X}_i$ is a *best response* to $x_{-i} \in \mathcal{X}_{-i}$ if

$$u_i(x_i^*; x_{-i}) \geq u_i(x_i; x_{-i}) \quad \text{for all } x_i \in \mathcal{X}_i.$$

or, equivalently, if

$$x_i^* \in \arg \max_{x_i \in \mathcal{X}_i} u_i(x_i; x_{-i})$$

The set-valued function $\text{BR}_i: \mathcal{X}_{-i} \rightrightarrows \mathcal{X}_i$ given by

$$\text{BR}_i(x_i) := \arg \max_{x_i \in \mathcal{X}_i} u_i(x_i; x_{-i})$$

is called the *best response correspondence* of player i . Collectively, we also let

$$\text{BR}(x) = \prod_{i \in I} \text{BR}_i(x_i)$$

Nash equilibrium as fixed points

$$x^* \text{ is a Nash equilibrium of } \mathcal{G} \iff x_i^* \in \text{BR}_i(x_{-i}^*) \text{ for all } i \in \mathcal{N} \iff x^* \in \text{BR}(x^*)$$

Παράδειγμα: Υπολογισμός ΣΣΙ Nash σε ένα δωαήλιο Cournot (Cournot duopoly)

Μας δίνονται $N=2$ εταιρείες με το ίδιο κόστος παραγωγής c ανά μονάδα προϊόντος & γραμμική εξάρτηση ανάμεσα στη συνολική παραγωγή & την τιμή του προϊόντος, δηλ $P(x_{tot}) = a - bx_{tot}$ ($x_{tot} = x_1 + x_2$)


Να υπολογισθούν τα ΣΣΙ Nash αυτού του μοντέλου Cournot.

1^η Βήμα: Υπολογισμός ως συναρτήσεις ωφέλειας:

$$\begin{cases} u_1(x) = x_1 [P(x_{tot}) - c] = [a - b(x_1 + x_2) - c]x_1 \\ \quad \quad \quad = rx_1 - bx_1x_2 - bx_1^2 & (r = a - c) \\ u_2(x) = rx_2 - bx_1x_2 - bx_2^2 \end{cases}$$

2^η Βήμα: Υπολογισμός των βέλτιστων αποκρίσεων:

$$BR_1(x_2) = \operatorname{argmax}_x \{ rx_1 - bx_1x_2 - bx_1^2 \} = \frac{r - bx_2}{2b}$$

$$BR_2(x_1) = \dots = \frac{r - bx_1}{2b}$$


Συμπεραίνουμε ότι το ΣΣΙ Nash του παιχνιδιού είναι $x_1^* = x_2^* = r/3b$



Characterization: payoff gradients

Definition (Individual payoff gradients)

The *individual payoff gradient* of player $i \in \mathcal{N}$ is the vector field

$$v_i(x) = \nabla_{x_i} u_i(x_i; x_{-i})$$

and, collectively, the game's *individual gradient field* is

$$v(x) = (v_1(x), \dots, v_N(x))$$

Notes:

- ▶ In finite games: $\partial_{a_i} u_i(x) = u_i(a_i; x_{-i}) \implies$ **individual gradients = mixed payoff vectors**
- ▶ In general: convenient to assume u_i differentiable in an open neighborhood of \mathcal{X}



Characterization: payoff gradients

Definition (Individual payoff gradients)

The *individual payoff gradient* of player $i \in \mathcal{N}$ is the vector field

$$v_i(x) = \nabla_{x_i} u_i(x_i; x_{-i})$$

and, collectively, the game's *individual gradient field* is

$$v(x) = (v_1(x), \dots, v_N(x))$$

Notes:

- ▶ In finite games: $\partial_{a_i} u_i(x) = u_i(a_i; x_{-i}) \implies$ **individual gradients = mixed payoff vectors**
- ▶ In general: convenient to assume u_i differentiable in an open neighborhood of \mathcal{X}

$$\rightarrow u_i(x) = \mathbb{E}_{\alpha_{-i}} [u_i(a)] = \sum_{\alpha_i \in A_i} \sum_{\alpha_{-i} \in A_{-i}} x_{i\alpha_i} x_{-i\alpha_{-i}} u_i(\alpha_i, \alpha_{-i})$$

$$\frac{\partial u_i}{\partial x_{i\alpha_i}} = \sum_{\alpha_{-i} \in A_{-i}} x_{-i\alpha_{-i}} u_i(\alpha_i, \alpha_{-i}) = u_i(\alpha_i, x_{-i}) = v_{i\alpha_i}(x) = \alpha_i\text{-th coordinate of mixed payoff vector } v_i(x) \text{ of player } i$$



Characterization: payoff gradients

Definition (Individual payoff gradients)

The *individual payoff gradient* of player $i \in \mathcal{N}$ is the vector field

$$v_i(x) = \nabla_{x_i} u_i(x_i; x_{-i})$$

and, collectively, the game's *individual gradient field* is

$$v(x) = (v_1(x), \dots, v_N(x))$$

Nash equilibrium as first-order stationary points

$$u_i(x_i^*; x_{-i}^*) \geq u_i(x_i; x_i^*) \implies \begin{cases} \langle v_i(x^*), x_i - x_i^* \rangle \leq 0 & \text{for all } x_i \in \mathcal{X}_i \text{ and all } i \in \mathcal{N} \\ \langle v(x^*), x - x^* \rangle \leq 0 & \text{for all } x \in \mathcal{X} \end{cases} \quad (\text{FOS})$$



Characterization: payoff gradients

Definition (Individual payoff gradients)

The *individual payoff gradient* of player $i \in \mathcal{N}$ is the vector field

$$v_i(x) = \nabla_{x_i} u_i(x_i; x_{-i})$$

and, collectively, the game's *individual gradient field* is

$$v(x) = (v_1(x), \dots, v_N(x))$$

Nash equilibrium as first-order stationary points

$$u_i(x_i^*; x_{-i}^*) \geq u_i(x_i; x_i^*) \implies \begin{cases} \langle v_i(x^*), x_i - x_i^* \rangle \leq 0 & \text{for all } x_i \in \mathcal{X}_i \text{ and all } i \in \mathcal{N} \\ \langle v(x^*), x - x^* \rangle \leq 0 & \text{for all } x \in \mathcal{X} \end{cases} \quad (\text{FOS})$$

When do we have (??) \implies (NE)?



Outline

- ① A modern game
- ② Definitions and further examples
- ③ Nash equilibrium and characterizations
- ④ Concave games


Concave games

Definition (Concave games)

A *concave game* is a continuous game with *individually concave* payoff functions, i.e.,

$$u_i(x_i; x_{-i}) \text{ is concave in } x_i$$

for all $x_{-i} \in \mathcal{X}_{-i}$ and all $i \in \mathcal{N}$.



$$u(x^*) + \langle u'(x^*), x - x^* \rangle \geq u(x)$$

$$\Rightarrow u(x) - u(x^*) \leq \langle u'(x^*), x - x^* \rangle$$



Concave games

Definition (Concave games)

A *concave game* is a continuous game with *individually concave* payoff functions, i.e.,

$$u_i(x_i; x_{-i}) \text{ is concave in } x_i$$

for all $x_{-i} \in \mathcal{X}_{-i}$ and all $i \in \mathcal{N}$.

Notes:

- ▶ **Gradient dominance:**

$$u_i(x_i; x_{-i}) \leq u_i(x_i^*; x_{-i}) + \langle v_i(x_i^*; x_{-i}), x_i - x_i^* \rangle$$

- ▶ **Stationarity implies optimality:**

$$\langle v_i(x_i^*; x_{-i}), x_i - x_i^* \rangle \leq 0 \implies u_i(x_i; x_{-i}) \leq u_i(x_i^*; x_{-i})$$

- ▶ Closed convex $\arg \max_{x_i \in \mathcal{X}_i} u_i(x_i; x_{-i})$



Concave games

Definition (Concave games)

A **concave game** is a continuous game with **individually concave** payoff functions, i.e.,

$$u_i(x_i; x_{-i}) \text{ is concave in } x_i$$

for all $x_{-i} \in \mathcal{X}_{-i}$ and all $i \in \mathcal{N}$.

Proposition (Variational characterization of Nash equilibria)


Let $\mathcal{G} \equiv \mathcal{G}(i, \mathcal{X}, u)$ be a concave game. Then

$$x^* \text{ is a Nash equilibrium} \iff \langle v(x^*), x - x^* \rangle \leq 0 \text{ for all } x \in \mathcal{X}$$

*Stampacchia
Variational
Inequality (SVI)*



Geometric characterization



At Nash equilibrium, individual payoff gradients are outward-pointing



Existence of Nash equilibria

Theorem (Debreu, 1952)

Every concave game with compact action spaces admits a Nash equilibrium.



Existence of Nash equilibria

Theorem (Debreu, 1952)

Every concave game with compact action spaces admits a Nash equilibrium.

Proof idea:

as in finite games

- ▶ Fixed point characterization of Nash equilibria

$$x^* \text{ is a Nash equilibrium} \iff x^* \in \text{BR}(x^*)$$



Existence of Nash equilibria

Theorem (Debreu, 1952)

Every concave game with compact action spaces admits a Nash equilibrium.

Proof idea:

as in finite games

- ▶ Fixed point characterization of Nash equilibria

$$x^* \text{ is a Nash equilibrium} \iff x^* \in \text{BR}(x^*)$$

- ▶ If the game is concave, $\text{BR}: \mathcal{X} \rightrightarrows \mathcal{X}$ is ~~nonempty~~, closed and convex

If \mathcal{X} is compact $\Rightarrow \text{BR}(x)$ is nonempty



Existence of Nash equilibria

Theorem (Debreu, 1952)

Every concave game with compact action spaces admits a Nash equilibrium.

Proof idea:

as in finite games

- ▶ Fixed point characterization of Nash equilibria

$$x^* \text{ is a Nash equilibrium} \iff x^* \in \text{BR}(x^*)$$

- ▶ If the game is concave, $\text{BR}: \mathcal{X} \rightrightarrows \mathcal{X}$ is nonempty, closed and convex
- ▶ Invoke Kakutani's fixed-point theorem for set-valued functions

Theorem (Kakutani, 1941)

Let \mathcal{C} be a nonempty compact convex subset of \mathbb{R}^d , and let $F: \mathcal{C} \rightrightarrows \mathcal{C}$ be a set-valued function such that:

(P1) $F(x)$ is nonempty, closed and convex for all $x \in \mathcal{C}$

(P2) F is **upper hemicontinuous** at all $x \in \mathcal{C}$, i.e., $\tilde{x} \in F(x)$ whenever $x_t \rightarrow x$ and $\tilde{x}_t \rightarrow \tilde{x}$ for sequences $x_t \in \mathcal{C}$ and $\tilde{x}_t \in F(x_t)$.

Then there exists some $x^* \in \mathcal{C}$ such that $x^* \in F(x^*)$.



Proof of Debreu's theorem

Verify the conditions of Kakutani's theorem for $C \leftarrow \mathcal{X}$ and $F \leftarrow \text{BR}$:

(P1) $\text{BR}(x)$ is a face of \mathcal{X} , so it is nonempty, closed and convex

↪ Why?

(P2) Argue by contradiction

- ▶ Suppose there exist sequences $x_t, \tilde{x}_t \in \mathcal{X}$, $t = 1, 2, \dots$, such that $x_t \rightarrow x$, $\tilde{x}_t \rightarrow \tilde{x}$ and $\tilde{x}_t \in \text{BR}(x_t)$, but $\tilde{x} \notin \text{BR}(x)$.
- ▶ Then there exists a player $i \in \mathcal{N}$ and a deviation $x'_i \in \mathcal{X}_i$ such that

$$u_i(x'_i; x_{-i}) > u_i(\tilde{x}_i; x_{-i})$$

- ▶ But since $\tilde{x}_{i,t} \in \text{BR}(x_{-i,t})$ by assumption, we also have:

$$u_i(x'_i; x_{-i,t}) \leq u_i(\tilde{x}_{i,t}; x_{-i,t})$$

- ▶ Since $x_t \rightarrow x$, $\tilde{x}_t \rightarrow \tilde{x}$ and u_i is continuous, taking limits gives

$$u_i(x'_i; x_{-i}) \leq u_i(\tilde{x}_i; x_{-i})$$

which contradicts our original assumption. □





References I

- [1] Debreu, G. A social equilibrium existence theorem. *Proceedings of the National Academy of Sciences of the USA*, 38(10):886-893, October 1952.
- [2] Kakutani, S. A generalization of Brouwer's fixed point theorem. *Duke Mathematical Journal*, 8(3):457-459, September 1941.
- [3] Kelly, F. P., Maulloo, A. K., and Tan, D. K. H. Rate control for communication networks: shadow prices, proportional fairness and stability. *Journal of the Operational Research Society*, 49(3):237-252, March 1998.
- [4] McCulloch, W. and Pitts, W. A logical calculus of ideas immanent in nervous activity. *Bulletin of Mathematical Biophysics*, 5(4):115-133, 1943.
- [5] Tullock, G. The welfare costs of tariffs, monopolies, and theft. *Western Economic Journal*, 1967.
- [6] Tullock, G. Efficient rent seeking. In Tollison, J. M. B. R. D. and Tullock, G. (eds.), *Toward a theory of the rent-seeking society*. Texas A&M University Press, 1980.

Πόλος μονομερών παραβόλων στα μεγάλου μεγέθους

$$\langle v(x^*), x - x^* \rangle \leq 0 \quad \forall x \in X$$

ΕΙΝΑΙ ΤΟ x^* NASH?



$\text{maximizer} \Rightarrow v(x^*) = 0 \Rightarrow \langle v(x^*), x - x^* \rangle = 0 \quad \forall x$
 \hookrightarrow ΣΙΓΟΥΡΑ ΟΧΙ NASH