6.S890: Topics in Multiagent Learning

Lecture 2 – Prof. Daskalakis

Fall 2023

Recall: Prisoner's Dilemma

	Deny (cooperate)	Confess (betray)
Deny (cooperate)	-1, -1	-3, 0
Confess (betray)	0, -3	-2, -2

("-1" = "1 year in jail")

Our prediction: both prisoners will **confess**

Why?

No matter what the other player may play, confessing is optimal for me.

Playing confess is a **dominant strategy equilibrium**

Recall: Rock-Paper-Scissors

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	-1,1	1,-1	0,0

Our prediction: both players play **uniformly at random**

Why?

If my opponent plays uniformly at random, then playing uniformly at random is optimal for me.

Player u.a.r. is a **Nash equilibrium**

Remarks:

1. Nash is a much *weaker* solution of a game compared to dominant strategy equilibrium
   - need assumption/knowledge about other player's strategy to justify my strategy
2. No dominant strategy equilibrium exists in Rock-Paper-Scissors
3. No Nash equilibrium exists in pure (i.e. non-randomizing) strategies
4. There is a unique Nash equilibrium in this game

Football vs Theater

		5/6	1/6
		Insist on Theater	Accept Football
1/6	Accept Theater	1, 5	0, 0
5/6	Insist on Football	0, 0	5, 1

Our prediction here?

there are two obvious Nash equilibria

there is a 3rd Nash equilibrium $x_{blue} = \left(\frac{1}{6}, \frac{5}{6}\right)$ $x_{orange} = \left(\frac{5}{6}, \frac{1}{6}\right)$

cool fact: in two-player (non-degenerate games) there is always an odd number of Nash eq

Our focus (part I): Normal-Form Games

Normal-form Games: *Single-shot, simultaneous move, complete information* Games

Complete-information means:

- Every player knows their own objective as well as the objective of every other player

	Rock	Paper	Scissors
Rock	0,0	-1,1	1,-1
Paper	1,-1	0,0	-1,1
Scissors	-1,1	1,-1	0,0

	Deny (cooperate)	Confess (betray)
Deny (cooperate)	-1, -1	-3, 0
Confess (betray)	0, -3	-2, -2

	Insist on Theater	Accept Football
Accept Theater	1, 5	0, 0
Insist on Football	0, 0	5, 1

[if I throw away structure and represent this game as a huge table, whose rows/columns are all possible algorithms (a.k.a. contingency plans) that the two players can use]

More Abstract Game Formulation

- **Def:** A *finite n-player game* is described by:
  - a set of pure strategies/actions per player: S_p
  - a utility/payoff function per player: $u_p : \times_q S_q \to \mathbb{R}$
- **Def:** A *randomized/mixed strategy* for player p is any $x_p \in \Delta^{S_p}$
  - assigns probability $x_p(j)$ to each $j \in S_p$
  - i.e. Δ^{S_p} is the simplex whose vertices are identified with the elements of S_p
- **Def:** a player's *expected utility* is
  - $u_p(x_1, \ldots, x_n) = \sum_{s \in \times_q S_q} u_p(s)\, x_1(s_1) \cdot \ldots \cdot x_n(s_n)$

		5/6	1/6
		Theater!	Football fine
1/6	Theater fine	1, 5	0, 0
5/6	Football!	0, 0	5, 1

$S_{blue} = \{Theater\ fine, Football!\}$

$S_{orange} = \{Theater!, Football\ fine\}$

$x_{blue} = \left(\frac{1}{6}, \frac{5}{6}\right) \qquad x_{orange} = \left(\frac{5}{6}, \frac{1}{6}\right)$

$u_{blue} = \frac{1}{6} \cdot \frac{5}{6} \cdot 1 + \frac{5}{6} \cdot \frac{1}{6} \cdot 5 = \frac{5}{6}$

$u_{orange} = \frac{1}{6} \cdot \frac{5}{6} \cdot 5 + \frac{5}{6} \cdot \frac{1}{6} \cdot 1 = \frac{5}{6}$

More Abstract Game Formulation

- **Def:** A *finite n-player game* is described by:
  - a set of pure strategies/actions per player: S_p
  - a utility/payoff function per player: $u_p : \times_q S_q \to \mathbb{R}$
- **Def:** A *randomized/mixed strategy* for player p is any $x_p \in \Delta^{S_p}$
  - assigns probability $x_p(j)$ to each $j \in S_p$
  - i.e. Δ^{S_p} is the simplex whose vertices are identified with the elements of S_p
- **Def:** a player's *expected utility* is
  - $u_p(x_1, \ldots, x_n) = \sum_{s \in \times_q S_q} u_p(s)\, x_1(s_1) \cdot \ldots \cdot x_n(s_n)$
- A piece of very useful notation: if $x_1, \ldots, x_n$ are player strategies, then x_{-i} denotes the strategies of all players except player i's
- **Def:** a collection of mixed strategies $x_1, \ldots, x_n$ is a *Nash equilibrium* iff
  - $\forall i, x_i'$: $\quad u_i(x_i, x_{-i}) \geq u_i(x_i', x_{-i})$
- **Def:** a collection $x_1, \ldots, x_n$ is a *dominant strategy equilibrium* iff
  - $\forall i, x_i', x_{-i}'$: $\quad u_i(x_i, x_{-i}') \geq u_i({x_i}', x_{-i}')$

More Abstract Game Formulation

- **Def:** a collection $x_1, \dots, x_n$ is a *Nash equilibrium* iff
  - $\forall i, x_i'$: $\quad u_i(x_i, x_{-i}) \geq u_i(x_i', x_{-i})$
- **Def:** a collection $x_1, \dots, x_n$ is a *dominant strategy equilibrium* iff
  - $\forall i, x_i', x_{-i}'$: $\quad u_i(x_i, x_{-i}') \geq u_i(x_i', x_{-i}')$

	Theater!	Football fine
Theater fine	1, 5	0, 0
Football!	0, 0	5, 1

$$x_{blue} = \left(\frac{1}{6}, \frac{5}{6}\right) \qquad x_{orange} = \left(\frac{5}{6}, \frac{1}{6}\right)$$

$$u_{blue}(x_{blue}, x_{orange}) = \frac{1}{6} \cdot \frac{5}{6} \cdot 1 + \frac{5}{6} \cdot \frac{1}{6} \cdot 5 = \frac{5}{6}$$

$$u_{blue}('theater\ fine', x_{orange}) = \frac{5}{6} \cdot 1 + \frac{1}{6} \cdot 0 = \frac{5}{6}$$

$$u_{blue}('football!', x_{orange}) = \frac{5}{6} \cdot 0 + \frac{1}{6} \cdot 5 = \frac{5}{6}$$

$$\Rightarrow u_{blue}(x_{blue}', x_{orange}) = \frac{5}{6}, \forall\ x_{blue}'$$

Nash's Theorem

[Nash 1950]: Every finite game (i.e. with a finite number of players and a finite number of pure strategies per player) has a Nash equilibrium.

- We'll prove it!
- We'll make use of Brouwer's fixed point theorem, following a proof that Nash produced in 1951; his original proof used Kakutani's fixed point theorem.

Menu

- Refresher and game-theoretic formalism
- Nash's theorem
- von Neumann's theorem

Menu

- **Refresher and game-theoretic formalism**
- **Nash's theorem**
- von Neumann's theorem


Brouwer

Brouwer's Fixed Point Theorem

[Brouwer 1910]: Let $f : D \to D$ be a continuous function from a convex and compact subset D of the Euclidean space to itself. Then there exists an $x \in D$ s.t. $x = f(x)$.

closed and bounded

Below we show a few examples, when *D* is the 2-dimensional disk.

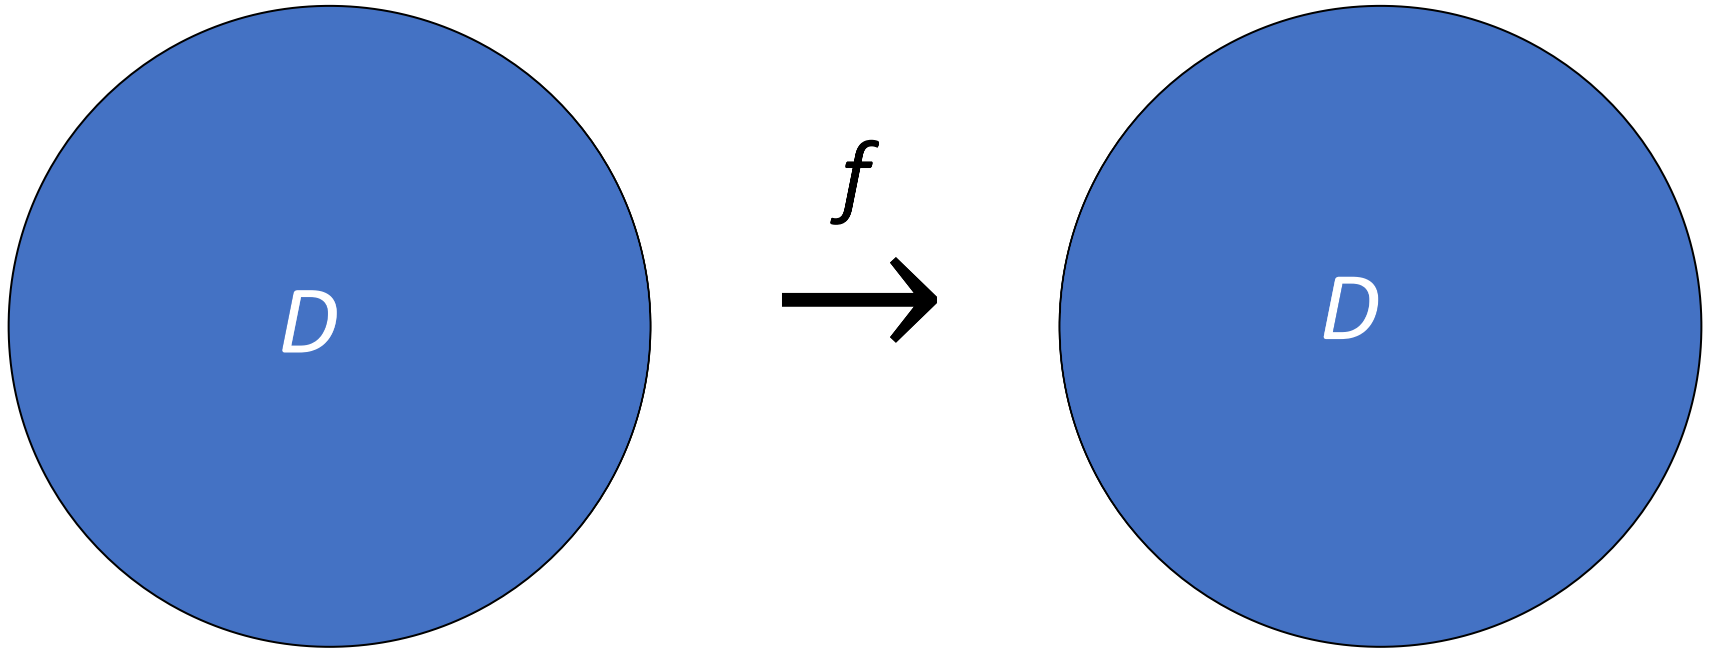

N.B. All conditions in the statement of the theorem are necessary.

Brouwer’s Fixed Point Theorem

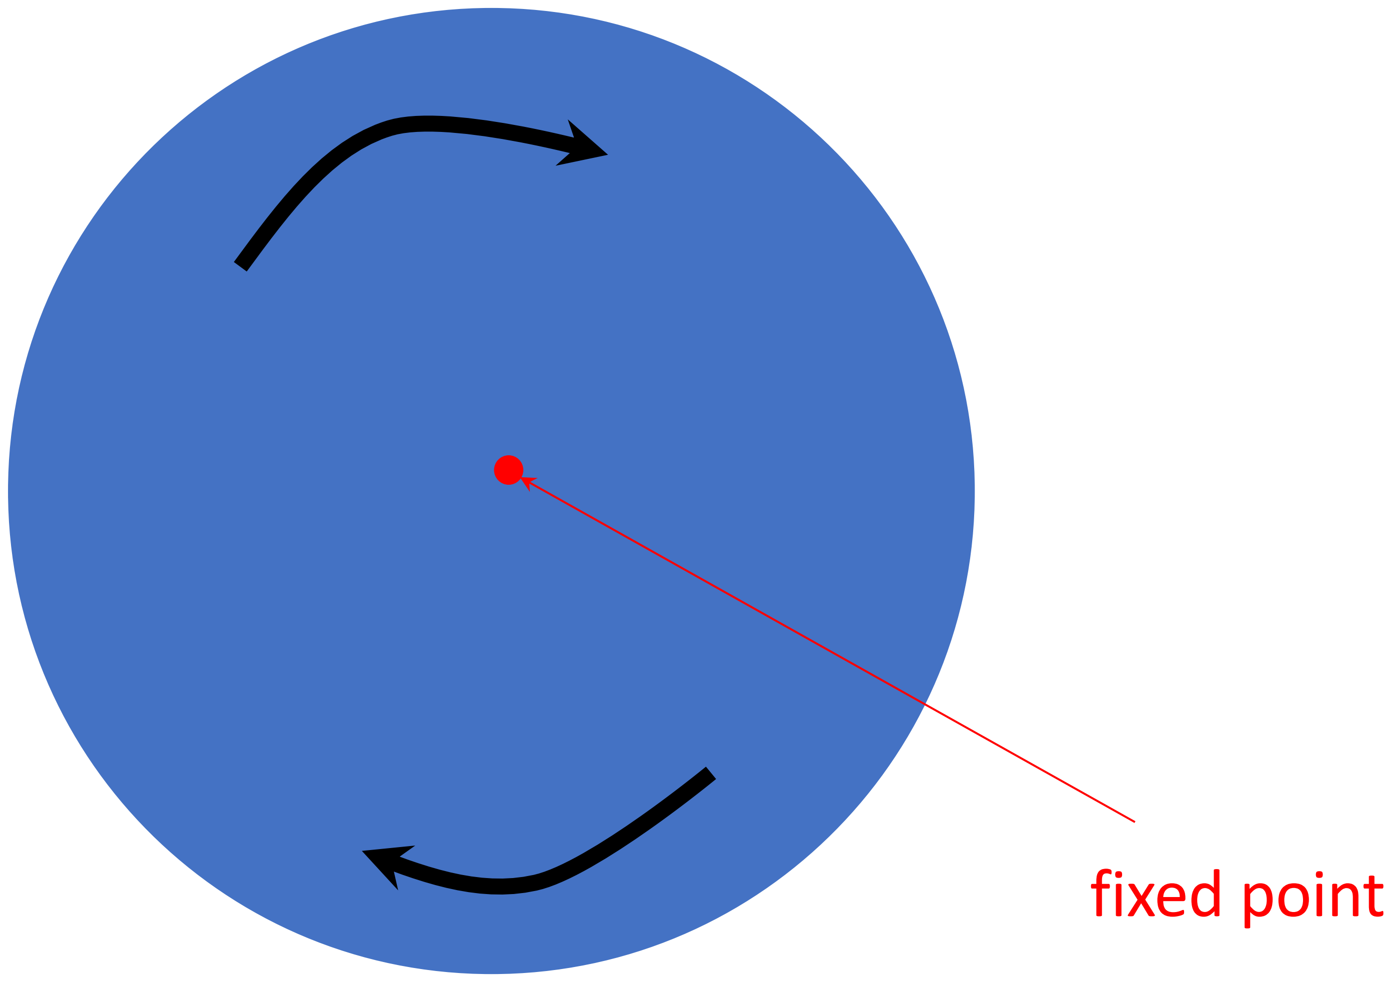

Brouwer’s Fixed Point Theorem

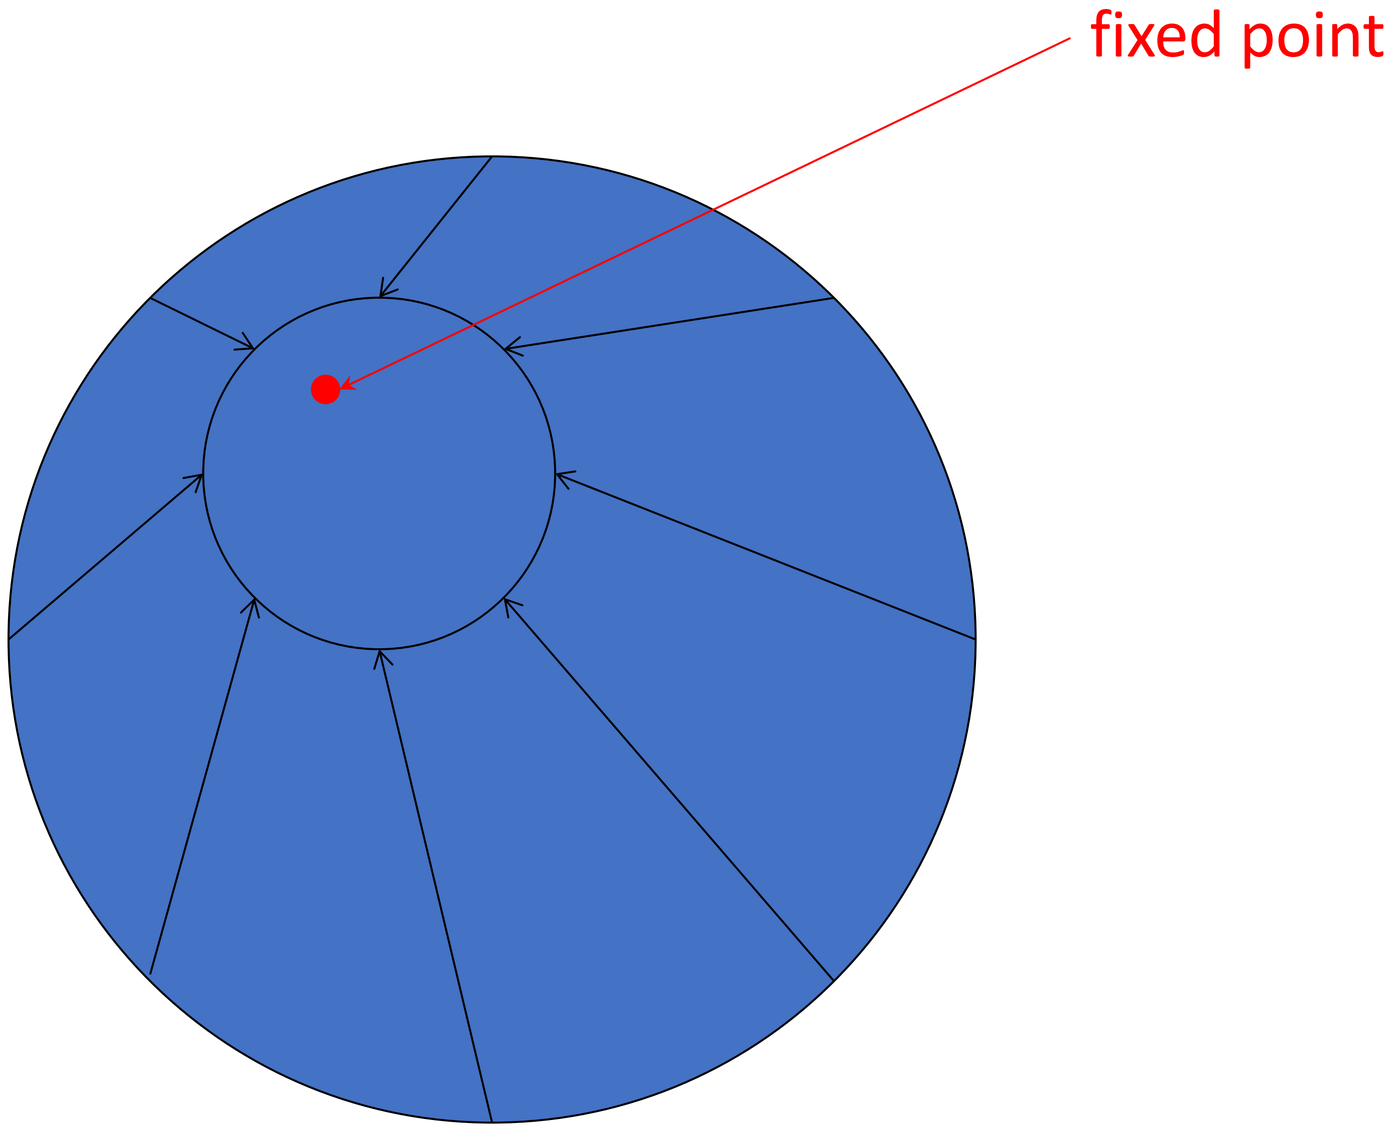

Brouwer's Fixed Point Theorem

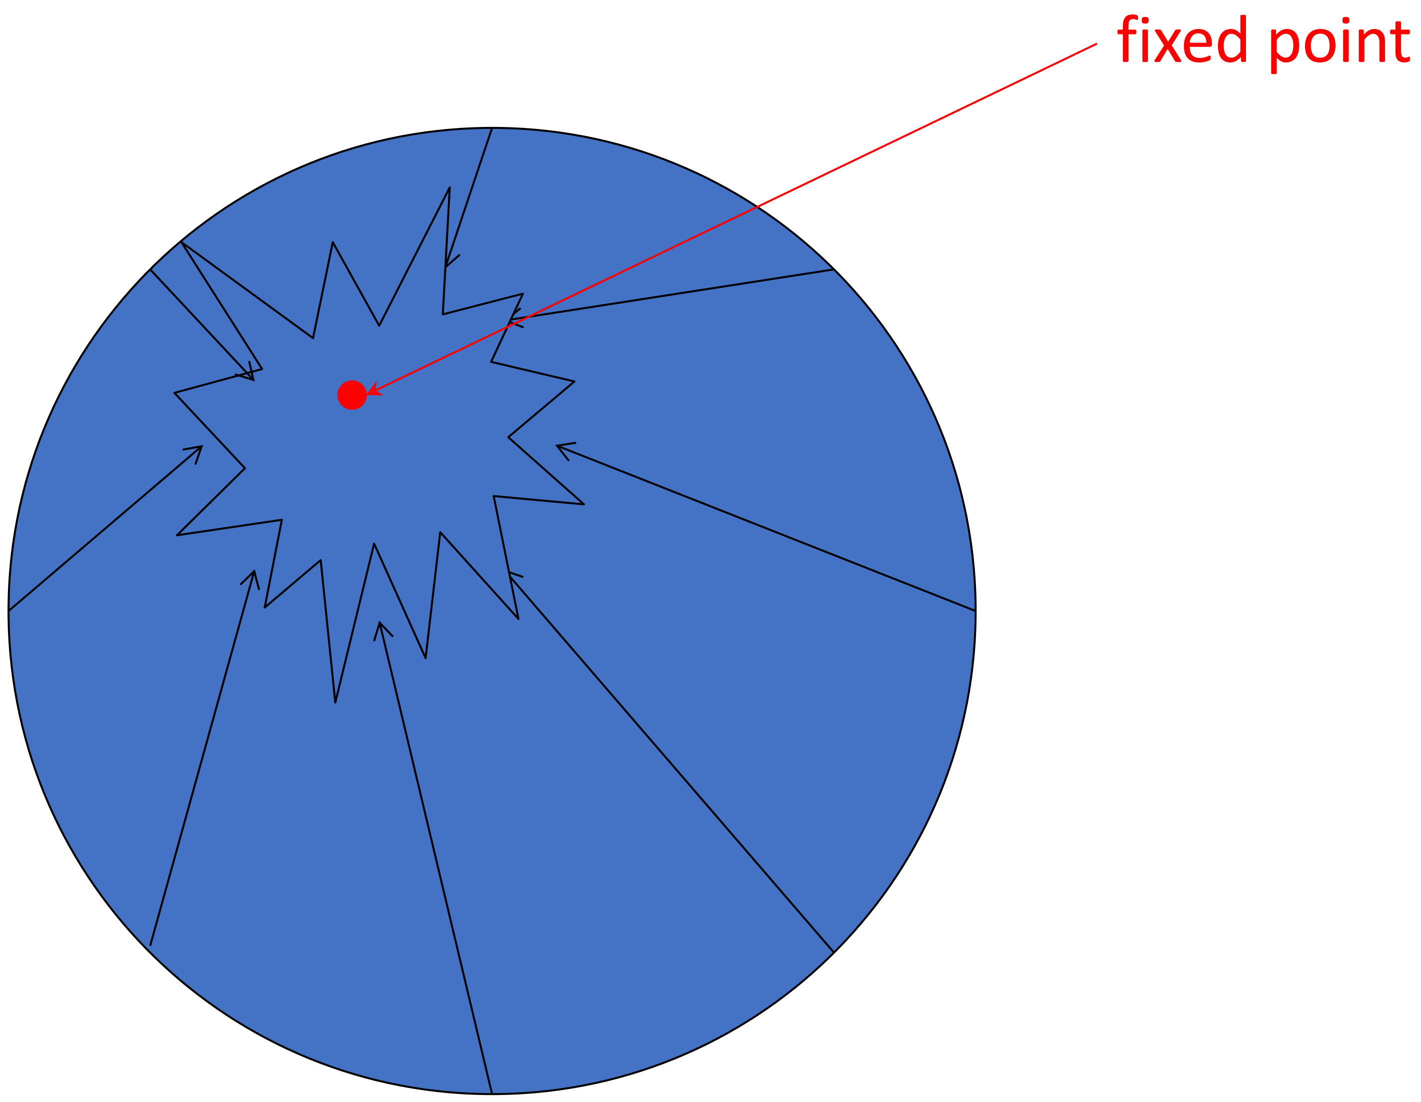


Brouwer ⇒ Nash

Visualizing Nash's Proof

Kick / Dive	Left	Right
Left	1 , -1	-1 , 1
Right	-1 , 1	1, -1

Penalty Shot Game

$\mapsto$

f: $[0,1]^2 \rightarrow [0,1]^2$, continuous
such that
fixed points $\equiv$ Nash eq.

Visualizing Nash's Proof

Kick / Dive	Left	Right
Left	1 , -1	-1 , 1
Right	-1 , 1	1, -1

Penalty Shot Game

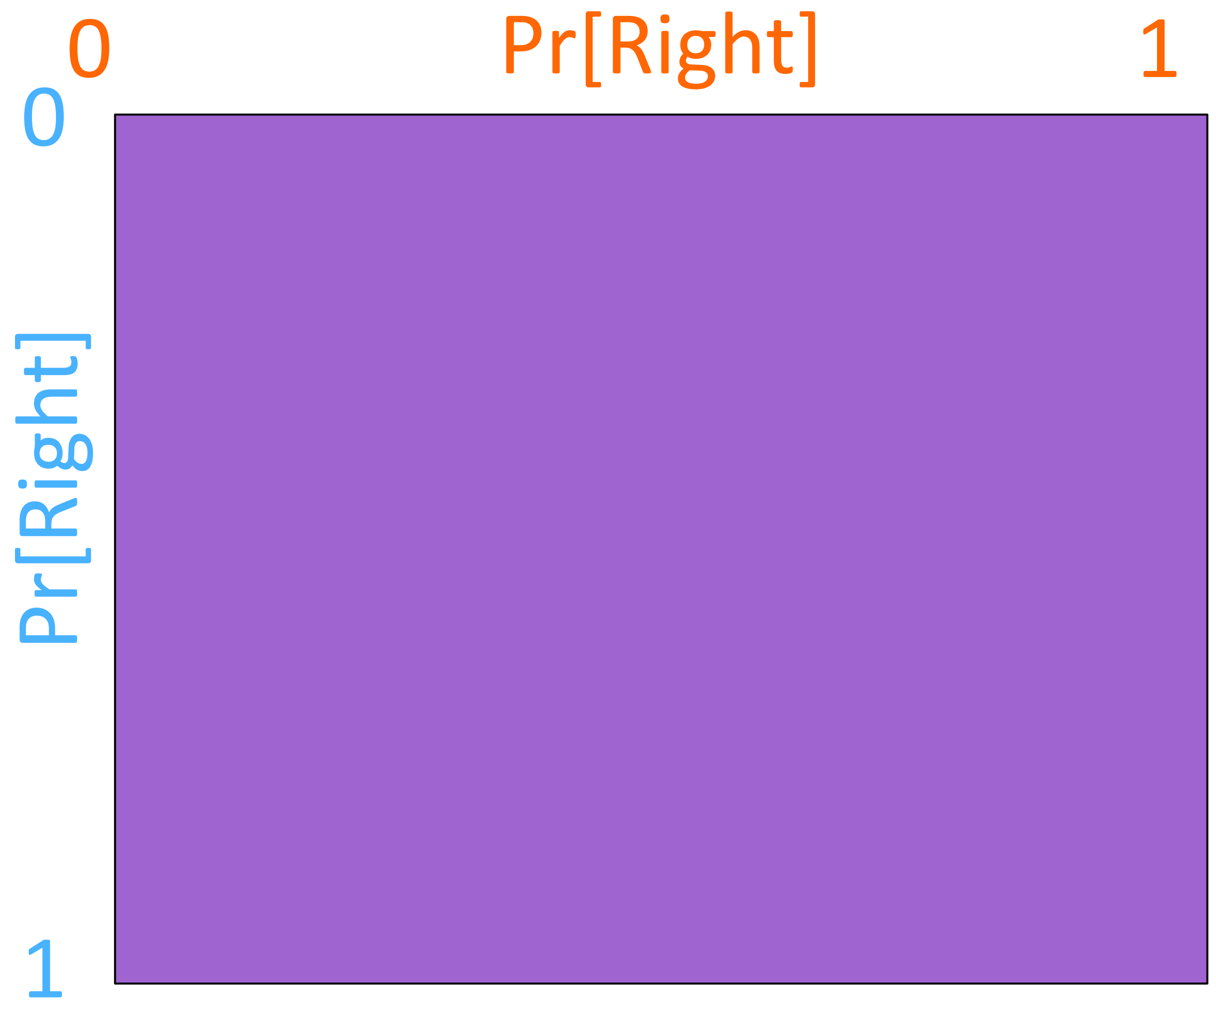

Visualizing Nash's Proof

Dive \ Kick	Left	Right
Left	1 , -1	-1 , 1
Right	-1 , 1	1, -1

Penalty Shot Game

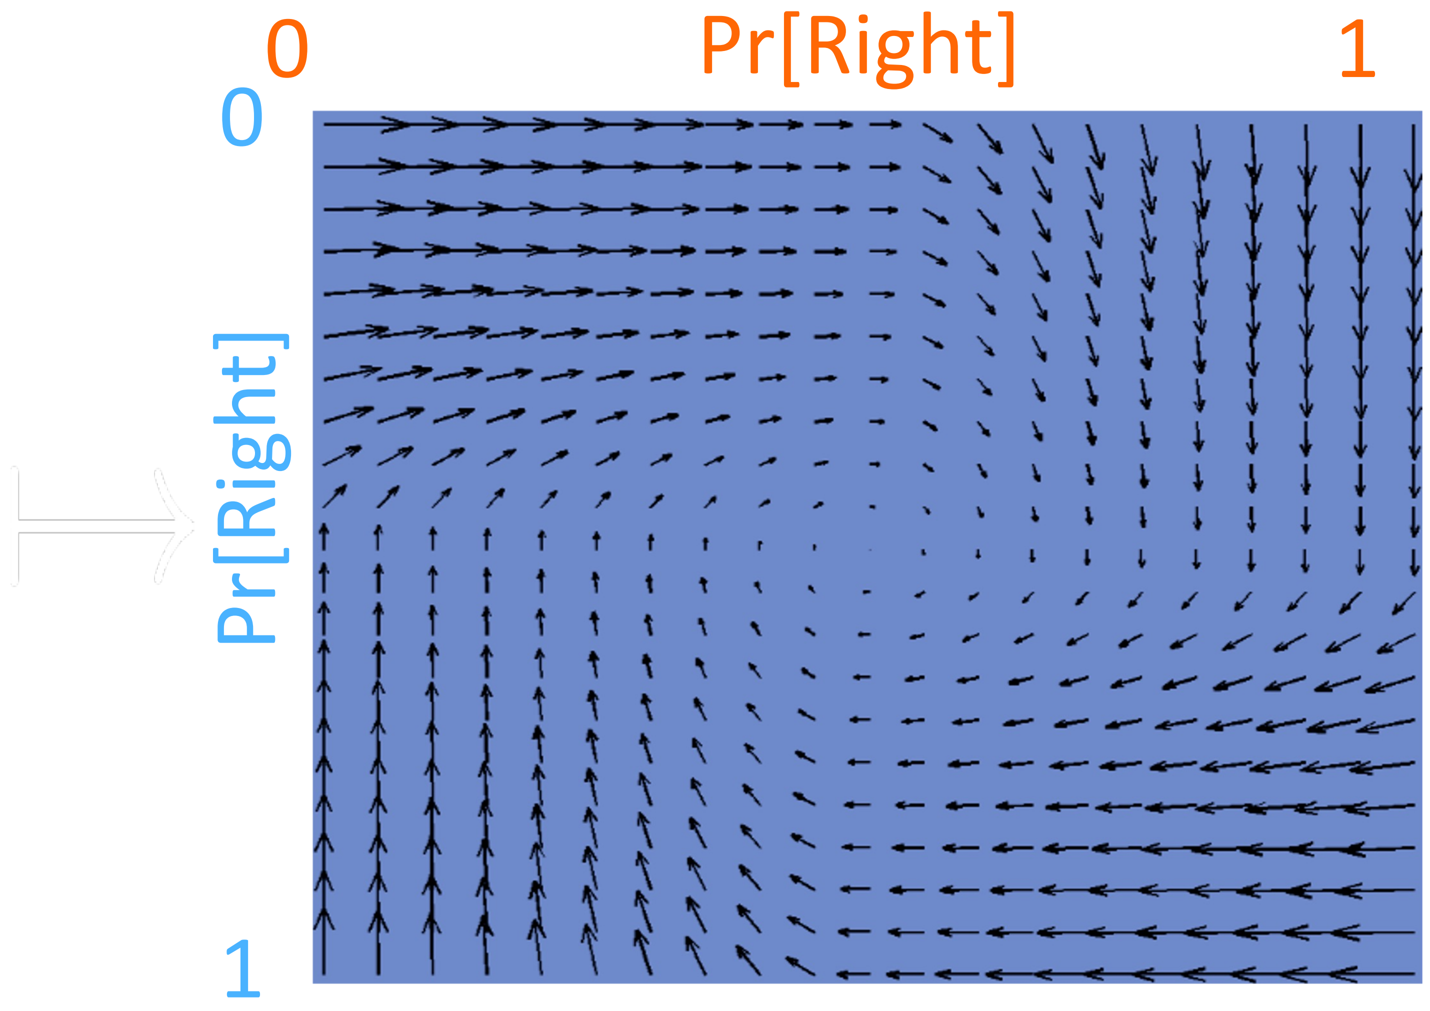

Visualizing Nash's Proof

Kick / Dive	Left	Right
Left	1 , -1	-1 , 1
Right	-1 , 1	1, -1

Penalty Shot Game

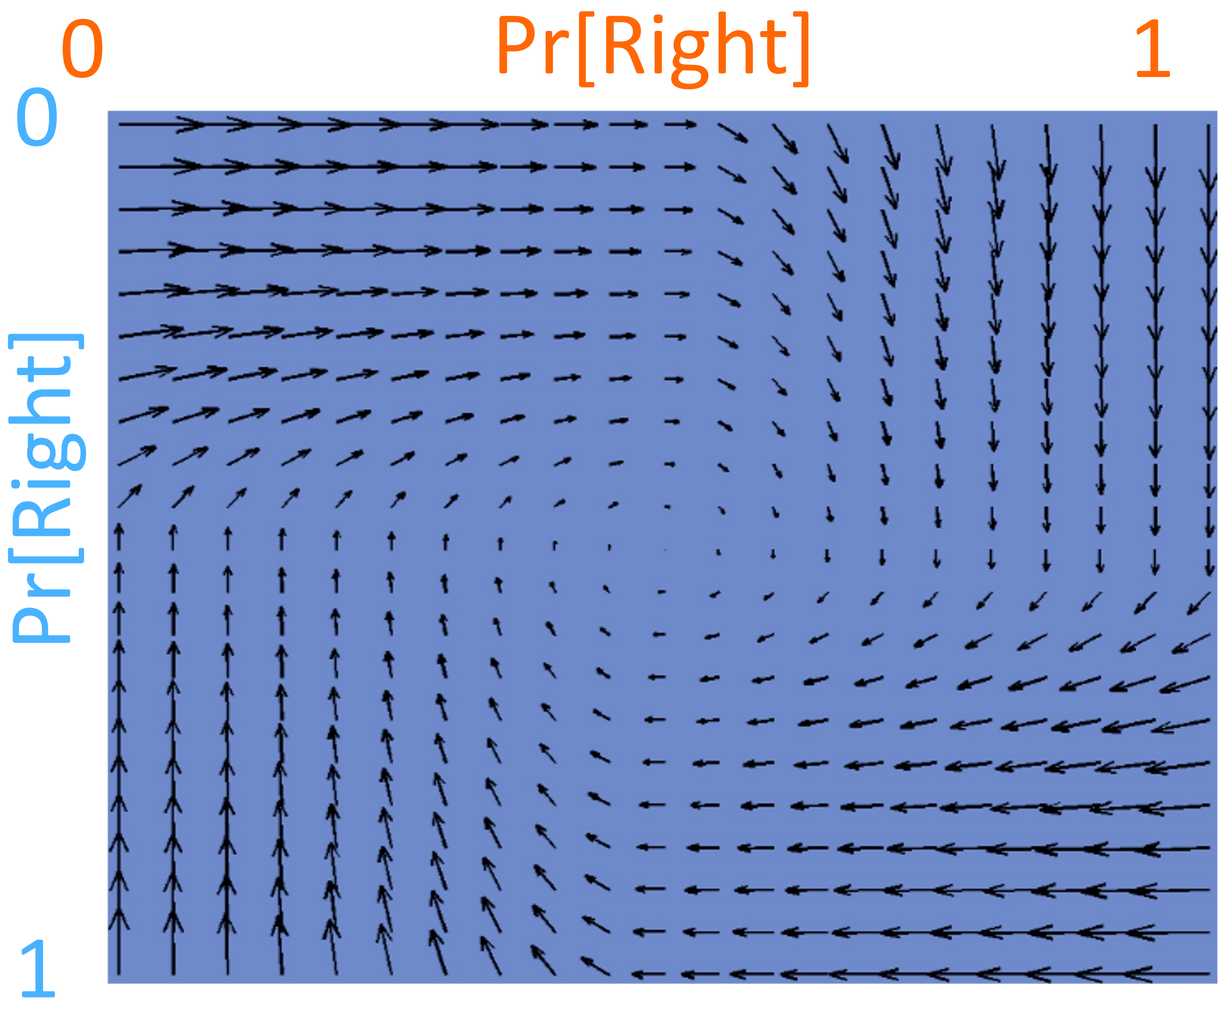

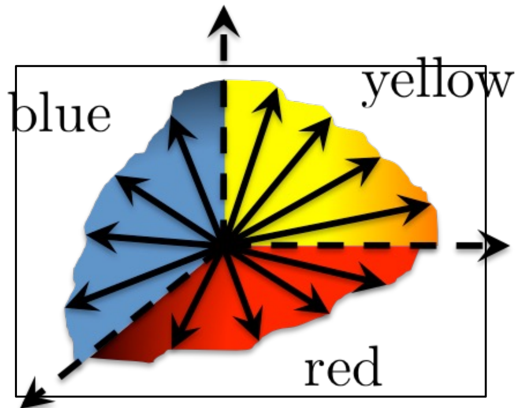

Visualizing Nash’s Proof

		½	½
	Kick / Dive	Left	Right
½	Left	1 , -1	-1 , 1
½	Right	-1 , 1	1, -1

Penalty Shot Game

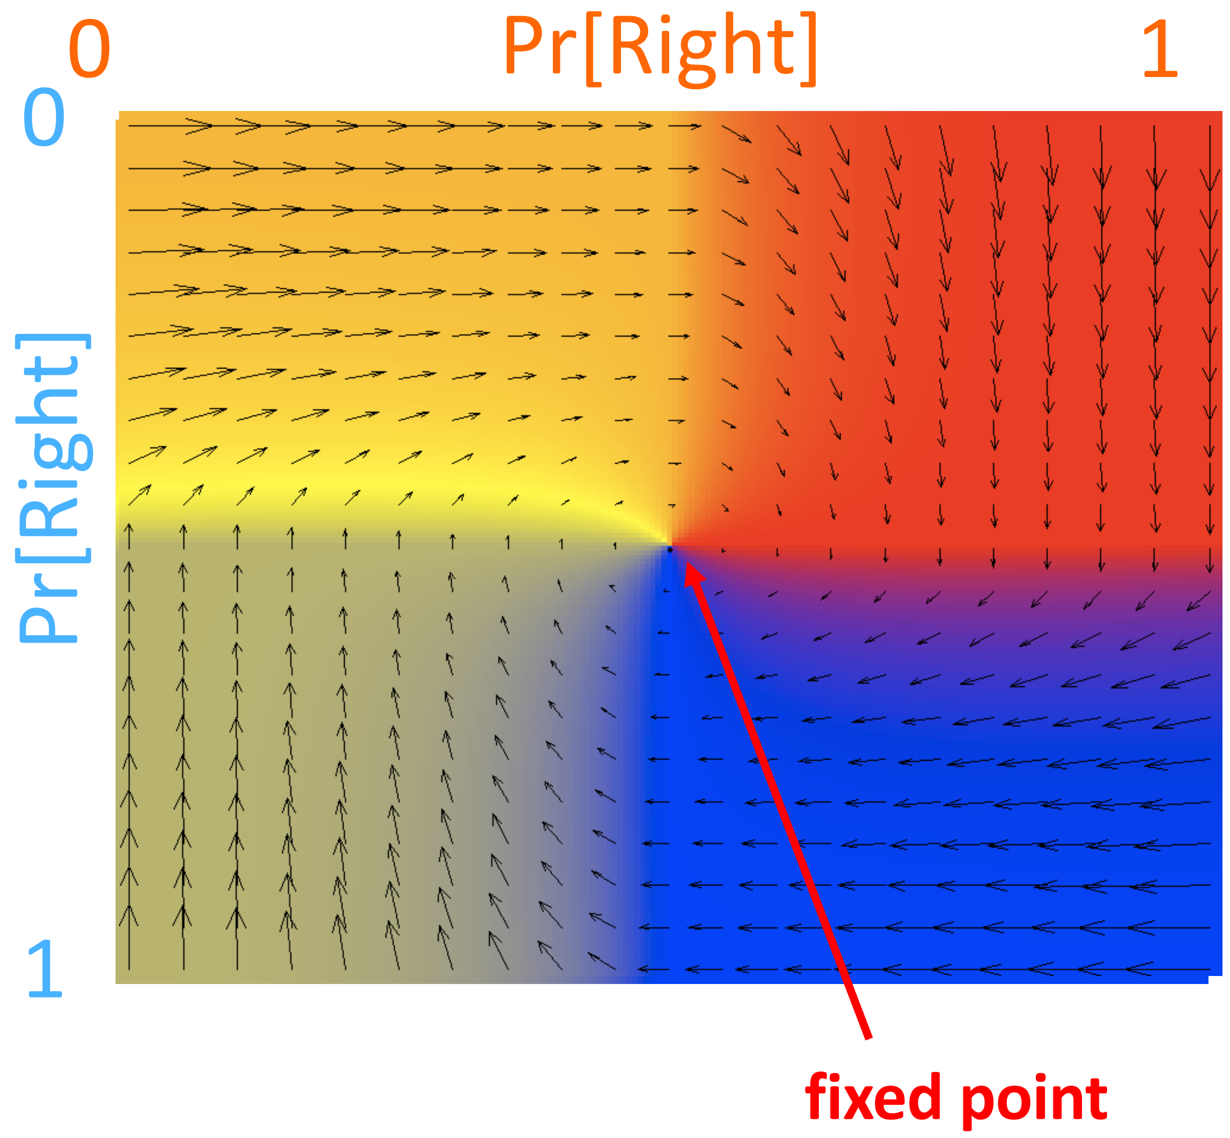

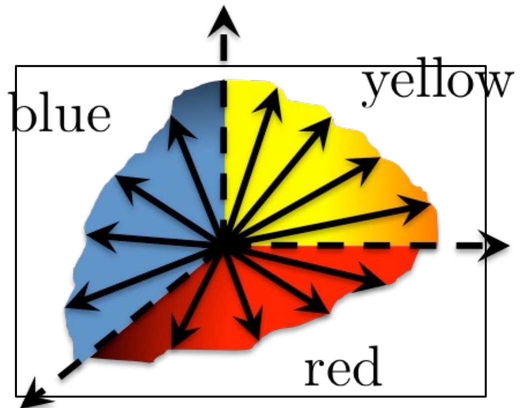

Real proof: on the board

[Nash '50]: Every finite game has a Nash Equilibrium

Proof: ① Define a function $f: \Delta^{S_1}\times\Delta^{S_2}\times\dots\times\Delta^{S_n} \to \Delta^{S_1}\times\Delta^{S_2}\times\dots\times\Delta^{S_n} = \Delta$

$$(x_1, x_2, \dots, x_n) \overset{f}{\longmapsto} \underbrace{(y_1, y_2, \dots, y_n)}_{\vec{y} = f(\vec{x})}$$

where $\forall i$: y_i is "soft best response of player i to x_{-i}"

namely $\forall s_i \in S_i$: $y_i(s_i) = \dfrac{x_i(s_i) + \mathrm{Gain}_{i,s_i}(x)}{1 + \sum_{s_i'} \mathrm{Gain}_{i,s_i'}(x)}$ where $\mathrm{Gain}_{i,s_i}(x) = \max(u_i(s_i; x_{-i}) - u_i(x), 0)$

[Notice $\sum_{s_i} y_i(s_i) = 1$ so y_i is valid mixed strategy]

② f continuous, Δ: convex + compact $\overset{\text{Brouwer}}{\Longrightarrow}$ $\exists (x_1^*, x_2^*, \dots, x_n^*) = f(x_1^*, \dots, x_n^*)$ (fixed point)

③ Claim: fixed point of f is Nash Equilibrium

x^*: Nash Eq. $\Leftrightarrow$ $\forall i, \forall x_i'$: $u_i(x_i'; x_{-i}^*) \le u_i(x^*)$ $\Leftrightarrow$ $\forall i, \forall s_i$

Proof: • Suffices to show $\forall i, s_i \in S_i$: $\mathrm{Gain}_{i,s_i}(x^*) = 0$ $\Leftrightarrow$ $u_i(s_i; x_{-i}^*) \le u_i(x^*)$ $\forall i, \forall s_i$

• Suppose this is not true

$\Rightarrow$ then $\exists i, s_i$ s.t. $\mathrm{Gain}_{i,s_i}(x^*) > 0$ $\left(\Leftrightarrow u_i(s_i; x_{-i}^*) - u_i(x^*) > 0\right)$

$\Rightarrow$ $x_i^*(s_i) > 0$ (o.w. can't be that $x_i^*(s_i) = \dfrac{x_i^*(s_i) + \overbrace{\mathrm{Gain}_{i,s_i}(x^*)}^{>0}}{1 + \sum_{s_i'} \mathrm{Gain}_{i,s_i'}(x^*)}$)

$\Rightarrow$ $\exists s_i''$ s.t. $x_i^*(s_i'') > 0$

and $u_i(s_i''; x_{-i}^*) - u_i(x^*) < 0$

(b.c. $u_i(x^*) = \sum_{s_i'} u_i(s_i'; x_{-i}^*) \cdot x_i^*(s_i')$

and b.c. $u_i(s_i; x_{-i}^*) > u_i(x^*)$

there must exist s_i'' s.t. $u_i(s_i''; x_{-i}^*) < u_i(x^*)$ & $x_i^*(s_i'') > 0$)

in particular

$\mathrm{Gain}_{i,s_i''}(x^*) = 0$ (b.c. ↑)

But $x_i^*(s_i'') = \dfrac{x_i^*(s_i'') + \overbrace{\mathrm{Gain}_{i,s_i''}(x^*)}^{=0}}{1 + \underbrace{\sum_{s_i'} \mathrm{Gain}_{i,s_i'}(x^*)}_{>0}} < x_i^*(s_i'')$

⇓

contradiction!!! ☒

Menu

- **Refresher and game-theoretic formalism**
- **Nash's theorem**
    - **Brouwer's theorem**
    - **Nash from Brouwer**
- von Neumann's theorem

Menu

- Refresher and game-theoretic formalism
- Nash's theorem
    - Brouwer's theorem
    - Nash from Brouwer
- von Neumann's theorem

Two-player games

- **Def:** A *finite n-player game* is described by:
  - a set of pure strategies/actions per player: S_p
  - a utility/payoff function per player: $u_p : \times_q S_q \to \mathbb{R}$
- A 2-player can be summarized by two matrices $(R\ , C)_{m\times n}$
  - rows : indexed by pure strategies of "row player"
  - columns : indexed by pure strategies of "column player"
- Mixed strategy for row player: $x \in \Delta^m$
- Mixed strategy for column player: $y \in \Delta^n$
- Expected utility of row player: $u_{\text{row}}(x, y) = x^T R y = \sum_{ij} R_{ij} x_i y_j$
- Expected utility of column player: $u_{\text{column}}(x, y) = x^T C y = \sum_{ij} C_{ij} x_i y_j$
- (x, y) is Nash equilibrium iff

$$\forall x' : x^T R y \geq x'^T R y$$
$$\forall y' : x^T C y \geq x^T C y'$$

Two-player *Zero-Sum* games

Minimax Theorem **[von Neumann'28]:** Consider a two-player game zero-sum game $(R, C)_{m\times n}$ i.e. $R + C = 0$. Then $\min_{x\in\Delta^m} \max_{y\in\Delta^n} x^T C y = \max_{y\in\Delta^n} \min_{x\in\Delta^m} x^T C y$ (*)

Interpretation:

- (*) says: "If $\forall y, \exists x$ s.t. $x^T C y \le v^* \Rightarrow \exists x, \forall y$ s.t. $x^T C y \le v^*$"
- If x^* is argmin of LHS, y^* argmax of RHS, v^* optimal value of (*), then (x^*, y^*) is a *Nash equilibrium*, i.e. if *min* and *max* adopt x^* and y^* then (i) *min* pays v^* to *max* and (ii) no player can improve by unilaterally deviating
- why? Because
  - under (x^*, y^*) *min* pays *max* at most v^* (since v^* optimum of LHS and x^* is argmin)
  - under (x^*, y^*) *max* receives from *min* at least v^* (since v^* optimum of RHS and y^* is argmax)
  - by the above two: under (x^*, y^*) min pays exactly v^* to max, hence (i) is proven
  - to prove (ii), suppose $\exists x$ that is a better response for min to y^*i.e. $x^T C y^* < {x^*}^T C y^* = v^*$
    - the existence of such x violates the fact that the optimum of RHS is v^* and y^* is an argmax for RHS
    - similarly the existence of a better response to x^* by max violates that the optimum of LHS is v^* and x^* is an argmin for the LHS