TECHNISCHE UNIVERSITÄT DRESDEN

Computational Logic ∴ Group

Hannes Strass

Faculty of Computer Science, Institute of Artificial Intelligence, Computational Logic Group

Normal-Form Games: Mixed Strategies

Lecture 2, 22th April 2024 // Algorithmic Game Theory, SS 2024

Previously ...

- Games can model real-life situations, but model fidelity is important.
- Noncooperative (strategic) games in **normal form** comprise players, **strategies** for the players, and **gain functions** for all **strategy profiles**.
- Various concepts can help predict/analyse the outcome of a game:
  - **Dominant strategies**
  - **Pareto optimality**
  - (pure) **Nash equilibria**
- We have analysed a number of example games: prisoner's dilemma, battle of the partners, chicken, penalties, and guessing numbers.
- Pure Nash equilibria need not always exist.

Chicken

Two people, Eli and Fyn, are racing towards each other in cars. Whoever swerves ("chickens out") loses face. If neither swerves, both get seriously injured.

(Eli, Fyn)	Swerve	RaceOn
Swerve	(2,2)	(1,3)
RaceOn	(3,1)	(0,0)

Overview

Motivation

Mixed Strategies and Equilibria

Computation of Nash Equilibria

Motivation

Motivation

- So far we considered (Nash equilibria in) pure strategies.
- For some games, such pure equilibria did not exist . . . , e.g. penalties:

Penalties

Two football players face off at a (simplified) single penalty kick. The kicker can kick left or right; the goal keeper can jump left or right. The kicker scores a goal iff they choose a different side than the keeper.

(Kicker,Keeper)	JumpL	JumpR
KickL	(-1,1)	(1,-1)
KickR	(1,-1)	(-1,1)

What does that predict about how the game will be played?

Background

Recall

A (discrete) **probability distribution** on a countable set S is a function

$$\pi\colon S \to [0,1] \quad \text{such that} \quad \sum_{s\in S} \pi(s) = 1$$

where $[0,1] \subseteq \mathbb{R}$ is the real unit interval.

$\pi(s)$ represents the probability of a random variable taking the value $s \in S$.

Definition

A **lottery** consists of a countable set S and a probability distribution on S.

Assumption: Expected Utility Hypothesis

If (S,π) is a lottery and there is a function $u\colon S \to \mathbb{R}$ assigning a utility to each outcome $s \in S$, then the expected utility of the lottery is given by

$$U(S,\pi) := \sum_{s\in S} u(s)\pi(s)$$

Lotteries and Risk Neutrality

Example

Consider the following events: and lotteries:

- e_0: "You get 0€."
- e_1: "You get 100€."
- e_2: "You get 200€."

- $L_1 = (\{e_1\}, \{e_1 \mapsto 1.0\})$
- $L_2 = (\{e_0, e_2\}, \{e_0 \mapsto 0.5, e_2 \mapsto 0.5\})$

Which of these lotteries would you prefer?

Terminology

- A **risk neutral** player is one who is indifferent between L_1 and L_2;
- a **risk averse** player is one who prefers L_1 over L_2;
- a **risk seeking** player is one who prefers L_2 over L_1.

We assume throughout this course that players are risk neutral.

Mixed Strategies and Equilibria

Nash Equilibrium in Mixed Strategies

Definition

Let $(P, \mathbf{S}, \mathbf{u})$ be a game in normal form and assume that all S_i are finite.

1. A **mixed strategy for player** $i \in P$ is a probability distribution π_i on S_i.
   - $\pi_i(s_j)$ is the probability of the event that player i chooses strategy $s_j \in S_i$.
   - Π_i denotes the set of all probability distributions on S_i, for each $1 \leq i \leq n$.
   - Denote $\boldsymbol{\Pi} := \Pi_1 \times \cdots \times \Pi_n$ and $\boldsymbol{\Pi}_{-i} := \Pi_1 \times \cdots \times \Pi_{i-1} \times \Pi_{i+1} \times \cdots \times \Pi_n$.
2. The **expected utility of a mixed-strategy profile** $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_n)$ for i is

$$U_i(\boldsymbol{\pi}) := \sum_{\mathbf{s}=(s_1,\ldots,s_n)\in\mathcal{S}} u_i(\mathbf{s}) \cdot \prod_{j=1}^{n} \pi_j(s_j)$$

- $\boldsymbol{\Pi}$ is the set of all mixed-strategy profiles for all players.
- $\boldsymbol{\Pi}_{-i}$ is the set of all mixed-strategy profiles for all players except i.
- Likewise, for $\boldsymbol{\pi} \in \boldsymbol{\Pi}$ and $i \in P$, we have $\boldsymbol{\pi}_{-i} \in \boldsymbol{\Pi}_{-i}$.

Mixed Strategies: Examples

Notation

Let $(P, \mathbf{S}, \mathbf{u})$ be a game with $P = \{1, \ldots, n\}$ and $S_i = \{s_1, \ldots, s_{k_i}\}$ for $1 \leq i \leq n$.
We denote a mixed strategy π_i for player i as a k_i-tuple $\pi_i = (\pi_i(s_1), \ldots, \pi_i(s_{k_i}))$.

Examples

- In penalties, a mixed strategy for Kicker is

$$\pi_{\texttt{Kicker}} = (\pi_{\texttt{Kicker}}(\texttt{KickL}), \pi_{\texttt{Kicker}}(\texttt{KickR})) = \left(\frac{1}{2}, \frac{1}{2}\right)$$

  Another mixed strategy for Kicker is $\pi_{\texttt{Kicker}} = \left(\frac{2}{3}, \frac{1}{3}\right)$.
- In rock-paper-scissors, a mixed strategy for player Ann is

$$\pi_{\texttt{Ann}} = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$$

Best Responses and Mixed Nash Equilibria

Definition

Let $(P, \mathbf{S}, \mathbf{u})$ be a game in normal form and assume that all S_i are finite.

1. A mixed strategy $\pi_i \in \Pi_i$ is player i's **best response** to the mixed-strategy profile $\boldsymbol{\pi}_{-i} \in \boldsymbol{\Pi}_{-i}$ iff for all mixed strategies $\pi_i' \in \Pi_i$, we have

$$U_i(\pi_1, \ldots, \pi_{i-1}, \pi_i, \pi_{i+1}, \ldots, \pi_n) \geq U_i(\pi_1, \ldots, \pi_{i-1}, \pi_i', \pi_{i+1}, \ldots, \pi_n)$$

2. A profile $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_n)$ is (in) a **Nash equilibrium in mixed strategies** iff mixed strategy π_i is a best response to $\boldsymbol{\pi}_{-i}$ for all players $1 \leq i \leq n$.

Examples

- In penalties, assume `Kicker` plays $\left(\frac{1}{2}, \frac{1}{2}\right)$. A best response of `Keeper` to this is $\left(\frac{1}{2}, \frac{1}{2}\right)$; other best responses are $(1, 0)$ and $(0, 1)$.
- In Rock-Paper-Scissors, a best response to $\left(\frac{2}{3}, \frac{1}{6}, \frac{1}{6}\right)$ is $(0, 1, 0)$.

Mixed Nash Equilibria: Characterisation (1)

Observation

A pure strategy $s_j \in S_i$ is a special case of a mixed strategy π_i with $\pi_i(s_j) = 1$ and $\pi_i(s_k) = 0$ for all $s_k \in S_i$ with $k \neq j$. We conveniently denote such π_i by s_j.

Definition

The **support** of a mixed strategy π_i for player i is the set $\{s_j \mid \pi_i(s_j) > 0\}$.

Theorem

1. Let $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_n)$ be a mixed-strategy profile in a game in normal form. A mixed strategy π_i is a best response to mixed-strategy profile $\boldsymbol{\pi}_{-i}$ if and only if all pure strategies in the support of π_i are best responses to $\boldsymbol{\pi}_{-i}$.
2. Every pure Nash equilibrium is also a mixed Nash equilibrium.

The converse of 2. is not the case: there are “proper” mixed Nash equilibria.

Mixed Nash Equilibria: Characterisation (2)

By definition, a mixed-strategy profile $\boldsymbol{\pi}$ is a Nash equilibrium iff for all i,

$$U_i(\boldsymbol{\pi}) = \max_{\pi'_i \in \Pi_i} U_i(\boldsymbol{\pi}_{-i}, \pi'_i)$$

where $U_i(\boldsymbol{\pi}_{-i}, \pi'_i)$ denotes $U_i(\pi_1, \ldots, \pi_{i-1}, \pi'_i, \pi_{i+1}, \ldots, \pi_n)$.

By the previous theorem it is enough to focus on the pure strategies, thus

$$\max_{\pi'_i \in \Pi_i} U_i(\boldsymbol{\pi}_{-i}, \pi'_i) = \max_{s_j \in \Pi_i} U_i(\boldsymbol{\pi}_{-i}, s_j)$$

Altogether, $\boldsymbol{\pi}$ is a mixed Nash equilibrium if and only if for all players i:

$$U_i(\boldsymbol{\pi}) = \max_{s_j \in \Pi_i} U(\boldsymbol{\pi}_{-i}, s_j)$$

Thus only best responses occur in the support of each player's π_i in $\boldsymbol{\pi}$.

Mixed Nash Equilibria: Examples (1)

Battle of the Partners

- By the previous theorem, both pure Nash equilibria are also mixed N. e.
- A third mixed Nash equilibrium $\boldsymbol{\pi}$ is obtained by reasoning as follows:

$$U_{\texttt{Cat}}(\texttt{Cinema}) = 10 \cdot \pi_{\texttt{Dee}}(\texttt{Cinema}) + 2 \cdot \pi_{\texttt{Dee}}(\texttt{Dancing})$$
$$U_{\texttt{Cat}}(\texttt{Dancing}) = 0 \cdot \pi_{\texttt{Dee}}(\texttt{Cinema}) + 7 \cdot \pi_{\texttt{Dee}}(\texttt{Dancing})$$

To make `Cat` indifferent between the two choices, `Dee` must choose the values for $\pi_{\texttt{Dee}}(\texttt{Cinema})$ and $\pi_{\texttt{Dee}}(\texttt{Dancing})$ such that

$$10 \cdot \pi_{\texttt{Dee}}(\texttt{Cinema}) + 2 \cdot \pi_{\texttt{Dee}}(\texttt{Dancing}) = 7 \cdot \pi_{\texttt{Dee}}(\texttt{Dancing})$$

With $\pi_{\texttt{Dee}}(\texttt{Cinema}) + \pi_{\texttt{Dee}}(\texttt{Dancing}) = 1$, we obtain $\pi_{\texttt{Dee}} = \left(\frac{1}{3}, \frac{2}{3}\right)$.

By symmetry, $\boldsymbol{\pi} = (\pi_{\texttt{Cat}}, \pi_{\texttt{Dee}}) = \left(\left(\frac{2}{3}, \frac{1}{3}\right), \left(\frac{1}{3}, \frac{2}{3}\right)\right)$ with $U_{\texttt{Cat}}(\boldsymbol{\pi}) = 4\frac{2}{3}$.

Mixed Nash Equilibria: Examples (2)

Penalties

- For the mixed strategies $\pi_{\texttt{Kicker}} = \left(\frac{1}{2}, \frac{1}{2}\right)$ and $\pi_{\texttt{Keeper}} = \left(\frac{1}{2}, \frac{1}{2}\right)$, the mixed-strategy profile $\boldsymbol{\pi} = (\pi_{\texttt{Kicker}}, \pi_{\texttt{Keeper}})$ is a (strict) Nash equilibrium with expected utilities $U_{\texttt{Kicker}}(\boldsymbol{\pi}) = 2 \cdot \frac{1}{4} \cdot (-1) + 2 \cdot \frac{1}{4} \cdot 1 = 0 = U_{\texttt{Keeper}}(\boldsymbol{\pi})$:
- If (e.g.) Kicker were to deviate by (e.g.) playing $\boldsymbol{\pi}' = \left(\frac{2}{3}, \frac{1}{3}\right)$, then Keeper would best-respond by playing $(1, 0)$, that is, playing JumpL, leading to expected utilities $U_{\texttt{Kicker}}(\boldsymbol{\pi}') = \frac{2}{3} \cdot (-1) + \frac{1}{3} \cdot 1 = -\frac{1}{3} = -U_{\texttt{Keeper}}(\boldsymbol{\pi}')$.

Rock-Paper-Scissors

- Similarly, for $\pi_{\text{Ann}} = \pi_{\text{Bob}} = \left(\frac{1}{3}, \frac{1}{3}, \frac{1}{3}\right)$, the mixed-strategy profile $\boldsymbol{\pi} = (\pi_{\text{Ann}}, \pi_{\text{Bob}})$ is a (strict) Nash equilibrium in mixed strategies.
- Note that every pure strategy is a best response to $\boldsymbol{\pi}$.

Mixed Strategies: Discussion

Question

What does it mean to play a mixed strategy?

There are a (at least) four answers:

- Players may randomise to confuse their opponents.
- Players randomise because they are uncertain about actions of others.
- Mixed strategies describe what might happen in repeated play.
- Mixed strategies describe population dynamics: they describe the probability of choosing a specific pure strategy out of a population of pure strategies.

Nash's Theorem

Theorem (Nash, 1950)

Let $G = (P, \mathbf{S}, \mathbf{u})$ be a noncooperative game in normal form.
If $P = \{1, \ldots, n\}$ is finite and for $\mathbf{S} = (S_1, \ldots, S_n)$ each S_i is finite, then there exists a Nash equilibrium in mixed strategies.

Proof Sketch.

- View pure strategies $s_i \in S_i$ as unit vectors in $\mathbb{R}^{|S_i|}$; mixed strategies $\pi_i \in \Pi_i$ are then points of a simplex, a convex subset of $\mathbb{R}^{|S_i|}$; $\boldsymbol{\Pi}$ is a simplotope.
- Define functions $\varphi_{ij}(\boldsymbol{\pi}) = \max \{0, U_i(\boldsymbol{\pi}_{-i}, s_j) - U_i(\boldsymbol{\pi})\}$ for $i \in P$, $1 \leq j \leq |S_i|$.
- Define (continuous) function $f \colon \boldsymbol{\Pi} \to \boldsymbol{\Pi}$ with $\boldsymbol{\pi} \mapsto \boldsymbol{\pi}' = (\pi'_1, \ldots, \pi'_n)$, where

$$\pi'_i(s_j) := \frac{\pi_i(s_j) + \varphi_{ij}(\boldsymbol{\pi})}{\sum_{s_k \in S_i} \left(\pi_i(s_k) + \varphi_{ik(\boldsymbol{\pi})}\right)}$$

- Use Brouwer's fixpoint theorem to deduce that f has at least one fixpoint.
- Show that $f(\boldsymbol{\pi}) = \boldsymbol{\pi}$ if and only if $\boldsymbol{\pi}$ is a mixed Nash equilibrium for G. $\square$

Computation of Nash Equilibria

Approaches to Find Equilibria for Two Players

Lemke-Howson algorithm (1964)

- Path-finding approach with geometrical interpretation
- Needs exponentially many steps in the worst case

Porter-Nudelman-Shoham (2004)

- Enumerates possible supports of mixed strategies, checks for equilibria
- Dominance checks and search bias for optimisation

Mixed Integer Programming (Sandholm, Gilpin, and Conitzer, 2005)

- Encode equilibrium property for a given game as a mixed integer program, i.e., as a mathematical (numerical) feasibility problem
- "Mixed" expresses that values for some variables may be real numbers

Mixed Integer Programming (in a Nutshell)

Definition

- A **mixed integer (linear) program** is of the form

$$\begin{aligned} \text{maximise} \quad & \mathbf{c}^T\mathbf{x} \\ \text{subject to} \quad & \mathbf{A}\mathbf{x} \leq \mathbf{b}, \\ & \mathbf{x} \geq 0, \\ \text{and} \quad & \mathbf{x} \in \mathbb{Z}^k \times \mathbb{R}^\ell \end{aligned}$$

  where $\mathbf{x}$ is a vector of **decision variables**, and $\mathbf{A}$, $\mathbf{b}$, $\mathbf{c}$ are a matrix and two vectors of real values; the expression $\mathbf{c}^T\mathbf{x}$ is the **objective function**.
- If there is no objective function the program is a **feasibility problem**.
- A **solution** is a variable-value assignment that satisfies all constraints.

E.g.: maximise $2x_1 - 3x_2$ subject to $x_1 + x_2 \leq 7, \quad 2x_1 - x_2 \leq 12, \quad x_1, x_2 \geq 0,$ and $(x_1, x_2) \in \mathbb{Z} \times \mathbb{R}$

Area of active research; used in industrial applications; solvers exist, e.g. SCIP.

Regret

Definition

Let $(P, \mathbf{S}, \mathbf{u})$ be a noncooperative game in normal form, $i \in P$, and $s_j \in S_i$.
The **regret** of i playing s_j w.r.t. opponent profile $\boldsymbol{\pi}_{-i}$ is

$$r_{\boldsymbol{\pi}_{-i}, s_j} := \left(\max_{\pi_k \in \Pi_i} U_i(\boldsymbol{\pi}_{-i}, \pi_k) \right) - U_i(\boldsymbol{\pi}_{-i}, s_j)$$

Example: Prisoner's Dilemma

- The regret of playing `Silent` in response to `Confess` is 1.
- The regret of `Confess` in response to any opponent strategy (profile) is 0.

More generally: The regret of any best response is zero.

Observation

A mixed strategy profile $\boldsymbol{\pi}$ is a Nash equilibrium if and only if every pure strategy is either played with probability zero or has zero regret.

Computing Nash Equilibria via MIP (1)

For every player $i \in P$ and pure strategy $s_j \in S_i$, introduce variables

- b_{s_j} ... to express that s_j is *not* played by i, i.e.
  - $\mathrm{b}_{s_j} = 1$ expresses that $\boldsymbol{\pi}(s_j) = 0$, and
  - $\mathrm{b}_{s_j} = 0$ expresses that $\boldsymbol{\pi}(s_j) > 0$ and $r_{\boldsymbol{\pi}_{-i},s_j} = 0$;
- u_i ... to express the maximal utility of i given $\boldsymbol{\pi}_{-i}$;
- p_{s_j} ... to express the probability with which s_j is played;
- u_{s_j} ... to express the expected utility from playing s_j;
- r_{s_j} ... to express the regret from playing s_j.

The formulation also uses the constants v_i, denoting the maximally possible difference between two payoffs for player i:

$$v_i := \max_{\substack{s_u^{(i)}, s_\ell^{(i)} \in S_i, \\ s_u^{(3-i)}, s_\ell^{(3-i)} \in S_{3-i}}} \left\{ u_i(s_u^{(1)}, s_u^{(2)}) - u_i(s_\ell^{(1)}, s_\ell^{(2)}) \right\}$$

(Note that $3 - i$ for $i \in P = \{1, 2\}$ just refers to the player other than i.)

Computing Nash Equilibria via MIP (2)

Definition

Let $G = (P, \mathbf{S}, \mathbf{u})$ be a strategic game with $P = \{1, 2\}$. Its MIP formulation is

$$\text{find} \quad \mathrm{p}_{s_j}, \mathrm{u}_i, \mathrm{u}_{s_j}, \mathrm{r}_{s_j}, \mathrm{b}_{s_j} \quad \text{such that}$$

$$\forall i: \quad \sum_{s_j \in S_i} \mathrm{p}_{s_j} = 1 \qquad (1)$$

$$\forall i: \forall s_j \in S_i: \quad \mathrm{u}_{s_j} = \sum_{s_k \in S_{3-i}} u_i(s_k, s_j)\mathrm{p}_{s_k} \qquad (2)$$

$$\forall i: \forall s_j \in S_i: \quad \mathrm{r}_{s_j} = \mathrm{u}_i - \mathrm{u}_{s_j} \qquad (3)$$

$$\forall i: \forall s_j \in S_i: \quad \mathrm{p}_{s_j} \leq 1 - \mathrm{b}_{s_j} \qquad (4)$$

$$\forall i: \forall s_j \in S_i: \quad \mathrm{r}_{s_j} \leq v_i \mathrm{b}_{s_j} \qquad (5)$$

$$\text{where} \quad \mathrm{p}_{s_j}, \mathrm{u}_i, \mathrm{u}_{s_j}, \mathrm{r}_{s_j} \geq 0 \quad \text{and} \quad \mathrm{b}_{s_j} \in \{0, 1\}$$

Computing Nash equilibria via MIP (3)

The intuition behind the constraints is as follows:

(1) Guarantees that the values of the p_{s_j} constitute a valid probability distribution for each player.

(2) Guarantees the correct utility value for playing s_j (where $u_i(s_k, s_j)$ denotes the constant $u_1(s_j, s_k)$ for $i = 1$ and the constant $u_2(s_k, s_j)$ otherwise).

(3) Guarantees the correct regret value for playing s_j.

(4) Guarantees that the probability of playing s_j is indeed zero whenever the strategy is chosen not to be played (otherwise the constraint is vacuous).

(5) Guarantees that the regret for playing s_j is indeed zero whenever the strategy is chosen to be played (otherwise the constraint is vacuous).

Proposition

For any two-player strategic game, the solutions of the MIP formulation correspond one-to-one to the mixed Nash equilibria of the game.

Recall: Computational Complexity

Recall

- Complexity class P contains all languages (decision problems) that can be decided by polynomial-time bounded deterministic Turing machines.
- Complexity class NP contains all languages $L \subseteq \Sigma^*$ for which there exists a **polynomial verifier**, that is, a polynomial-time bounded deterministic Turing machine M (and a polynomial p) such that:
  - M accepts only pairs (x, y) of words such that
    - $x \in L$, and
    - the length of y is at most polynomial in the length of x (i.e. $|y| \leq p(|x|)$);
  - for every $x \in L$ there is such a pair (x, y).
- A (polynomial-time) **(many-one) reduction** from A to B is a (polynomial-time) computable function f such that $w \in A$ iff $f(w) \in B$.
- A language L is **NP-hard** iff all languages in NP can be reduced to L.
- A language is **NP-complete** iff it is NP-hard and in NP.

Function Complexity Classes

- For (mixed) Nash equilibria, the question is not whether they exist.
- "Function" complexity classes contain "function" problems $F \subseteq \Sigma^* \times \Sigma^*$:

  Input: A word $x \in \Sigma^*$.

  Output: Any one $y \in \Sigma^*$ such that $(x, y) \in F$, if such a y exists; "no", otherwise.
- Output y can be thought of as solution to a search problem instance x.
- Solution y need not be unique for x (relation F need not be functional).
- Complexity class FP contains all search problems F where any y with $(x, y) \in F$ can be computed from x in deterministic polynomial time.
- Class FNP contains all F that are accepted by a polynomial verifier.

Examples

- Given a propositional formula φ, find a satisfying assignment if one exists.
- Given an undirected graph G and a $k \in \mathbb{N}$, find a k-clique in G if one exists.

Equilibria and Computational Complexity

Note

Finding a solution for a mixed integer feasibility problem with binary decision variables is FNP-complete.

Is finding Nash equilibria of noncooperative games also FNP-complete?
This is unlikely, as every game has at least one equilibrium.

Consider however the following variant:

Next-NE

Input: A strategic game G in normal form and a Nash equilibrium for G.

Output: Another Nash equilibrium of G, if one exists; “no” otherwise.

Proposition

Next-NE is FNP-complete.

Intuitively: Computing (mixed) Nash equilibria is computationally hard.

Conclusion

Summary

- A **mixed strategy** is a probability distribution on pure strategies.
- In a **mixed Nash equilibrium**, all players play best responses.
- Nash's Theorem: Mixed Nash equilibria always exist (for finite games).
- Nash equilibria for concrete games can be obtained via a translation to a **mixed integer program**:
  - Binary variables model the choices of pure strategies to put in the support;
  - real-valued variables model probabilities, utilities, and regret.
- Given a game and an equilibrium, it is FNP-complete to find *another* equilibrium for the game.

Action Points

- Obtain all (mixed) Nash equilibria for chicken and interpret them.