Nash equilibria for normal-form games

Lecture 2

Tomáš Kroupa

Normal-form games

1. *Player set* $N = \{1, \ldots, n\}$
2. *Finite strategy set* S_i for each $i \in N$, $\mathbf{S} := S_1 \times \cdots \times S_n$
3. *Utility function* $u_i \colon \mathbf{S} \to \mathbb{R}$ for each $i \in N$

- Player $i \in N$ can use a mixed strategy $p_i \in \boldsymbol{\Delta}_i$
- Mixed strategy profile $\mathbf{p} = (p_1, \ldots, p_n) \in \boldsymbol{\Delta}$ yields the expected utility $U_i(\mathbf{p})$ of player i
- If player i uses pure strategy $s_i \in S_i$ and the rest plays $\mathbf{p}_{-i}$, the expected utility of i is $U_i(s_i, \mathbf{p}_{-i}) := U_i(\delta_{s_i}, \mathbf{p}_{-i})$

Nash equilibria

The following are equivalent for a NE $\mathbf{p}^* = (p_1^*, \ldots, p_n^*) \in \Delta$.

1. $U_i(p_i, \mathbf{p}^*_{-i}) \leq U_i(\mathbf{p}^*)$, for each $i \in N$ and every $p_i \in \Delta_i$.
2. $U_i(s_i, \mathbf{p}^*_{-i}) \leq U_i(\mathbf{p}^*)$, for each $i \in N$ and every $s_i \in S_i$.

- The second condition says that a candidate for NE can be tested using only (finitely) many pure strategies
- We can use it to frame a NE computation as an optimization problem

Computing NE: Opt. problem (1)

- Vector variable $p_i \in \Delta_i$ for each $i \in N$
- Auxiliary variables $e_i \in \mathbb{R}$ representing the equilibrium utility of player $i \in N$

The *objective* is to minimize $\sum_{i \in N} (e_i - U_i(p_1, \ldots, p_n))$ subject to the *constraints*

1. $e_i \geq U_i(s_i, \mathbf{p}_{-i})$ for all $i \in N$ and each $s_i \in S_i$
2. $p_i \in \Delta_i$ for each $i \in N$

Computing NE: Opt. problem (2)

The following are equivalent for $\mathbf{p}^* = (p_1^*, \ldots, p_n^*)$.

1. $\mathbf{p}^*$ is a NE with $e_i^* = U_i(\mathbf{p}^*)$
2. $\mathbf{p}^*$ is the minimizer of the opt. problem with optimal value 0

But

- this optimization problem is nonconvex
- it has typically many local minima

Supports of mixed strategies

The *support* of a mixed strategy $p_i \in \Delta_i$ is the set

$$S(p_i) := \{s_i \in S_i \mid p_i(s_i) > 0\}.$$

The indifference principle

For every NE $\mathbf{p}^*$, each player $i \in N$, and every $s_i, t_i \in S(p_i^*)$,

$$U_i(s_i, \mathbf{p}^*_{-i}) = U_i(t_i, \mathbf{p}^*_{-i}) = U_i(\mathbf{p}^*).$$

Supports of equilibrium strategies

The following are equivalent for $\mathbf{p}^* = (p_1^*, \ldots, p_n^*) \in \Delta$.

1. $\mathbf{p}^*$ is a Nash equilibrium.
2. $S(p_i^*) \subseteq \mathbf{BR}_i(\mathbf{p}_{-i}^*)$, for each $i \in N$, where the best-response map here is

$$\mathbf{BR}_i(\mathbf{p}_{-i}^*) = \left\{ s_i \in S_i \mid U_i(s_i, \mathbf{p}_{-i}^*) = \max_{s_i' \in S_i} U_i(s_i', \mathbf{p}_{-i}^*) \right\}$$

Computing NE: Testing supports (1)

- Assumption: two players, $n = 2$
- Take $\Sigma_1 \subseteq S_1, \Sigma_2 \subseteq S_2$ and consider the *linear feasibility problem*

$$
\begin{aligned}
U_i(s_i, p_{-i}) &= e_i, \quad i = 1, 2,\ \forall s_i \in \Sigma_i, \\
U_i(s_i, p_{-i}) &\leq e_i, \quad i = 1, 2,\ \forall s_i \notin \Sigma_i, \\
p_i &\in \Delta_i, \\
p_i(s_i) &= 0, \quad i = 1, 2,\ \forall s_i \notin \Sigma_i.
\end{aligned}
$$

Computing NE: Testing supports (2)

- If there is a NE with supports Σ_1 and Σ_2, then it is a solution to the linear feasibility problem
- Any solution to the linear feasibility problem is a NE (p_1^*, p_2^*) such that $S(p_1^*) \subseteq \Sigma_1$ and $S(p_2^*) \subseteq \Sigma_2$ (caveat)

This leads to the *simple enumerating algorithm* to find one NE.

Computing NE: Support enumeration

This method will find at least one NE for a two-player game:

1. Generate subsets $\Sigma_1 \subseteq S_1$ and $\Sigma_2 \subseteq S_2$
2. Solve the linear feasibility problem for Σ_1 and Σ_2
   i. If *solvable*, then end.
   ii. If *unsolvable*, try 1.

The performance is increased with heuristics to search the space of supports (for example, preference of small supports).

Two-player zero-sum games (TPZS)

This class of games is computationally tractable:

1. *Player set* $N = \{1, 2\}$
2. *Finite strategy sets* S_1 and S_2
3. *Utility functions*: $u_1 \colon S_1 \times S_2 \to \mathbb{R}$ and $u_2 = -u_1$

- $u_1 + u_2 = 0$
- We write $u := u_1$
- Player 1 is *maximizing* (row player) and player 2 is *minimizing* (column player)

Solving TPZS games

- Nash equilibrium applies to TPZS games
- We will refine the equilibrium concept using the zero-sum assumption
- We use the well-known decision rule *minimax* to recover the worst-case *minimum gain*/*maximum loss* and show that this principle is equivalent to NE

Maximin strategy and lower value

1. If player 1 uses a mixed strategy p_1, then player 2 can achieve the loss

$$\min_{p_2 \in \Delta_2} U(p_1, p_2)$$

2. Player 1 then employs the mixed strategy p_1^* achieving the worst-case minimum gain

$$\underline{v} := \min_{p_2 \in \Delta_2} U(p_1^*, p_2) = \max_{p_1 \in \Delta_1} \min_{p_2 \in \Delta_2} U(p_1, p_2)$$

Minimax strategy and upper value

1. If player 2 uses a mixed strategy p_2, then player 1 can achieve the utility

$$\max_{p_1 \in \Delta_1} U(p_1, p_2)$$

2. Player 2 then employs the mixed strategy p_2^* achieving the worst-case maximum loss

$$\overline{v} := \max_{p_1 \in \Delta_1} U(p_1, p_2^*) = \min_{p_2 \in \Delta_2} \max_{p_1 \in \Delta_1} U(p_1, p_2)$$

Value of the TPZS game

It is easy to see that $\underline{v} \leq \overline{v}$. The converse is non-trivial:

Minimax theorem of von Neumann (1928)

$\underline{v} = \overline{v}$, for every TPZS game.

- The common value v is called the *value* of the game
- If players use maximin/minimax strategies (p_1^*, p_2^*), the resulting utility of player 1 is

$$\underline{v} = U(p_1^*, p_2^*) = \overline{v}$$

Nash equilibria in TPZS games

The following are equivalent for a mixed strategy profile $(p_1^*, p_2^*) \in \Delta$ in a TPZS.

1. For every $p_1 \in \Delta_1$ and every $p_2 \in \Delta_2$,

$$U(p_1, p_2^*) \leq U(p_1^*, p_2^*) \leq U(p_1^*, p_2).$$

2. For every $s_1 \in S_1$ and every $s_2 \in S_2$,

$$U(s_1, p_2^*) \leq U(p_1^*, p_2^*) \leq U(p_1^*, s_2).$$

3. $\underline{v} = U(p_1^*, p_2^*) = \overline{v}$

Solving TPZS games by LP (1)

- We know that

$$\max_{p_1 \in \Delta_1} \min_{p_2 \in \Delta_2} U(p_1, p_2) = \max_{p_1 \in \Delta_1} \min_{s_2 \in S_2} U(p_1, s_2)$$

  since linear function $U(p_1, .\,)$ on the convex polyhedron Δ_2 achieves minima over the extreme points
- The maximization problem for player 1 then becomes a *linear program* since we are maximizing a piecewise-linear concave function under linear constraints

Solving TPZS games by LP (Player 1)

- The LP for player 1 has variables $v_1 \in \mathbb{R}$ and $p_1 \in \Delta_1$
- *Maximize* v_1 subject to the constraints

$$U(p_1, s_2) \geq v_1, \quad \forall s_2 \in S_2$$
$$p_1 \in \Delta_1$$

- The optimal solution is the maximin strategy p_1^* and the lover value of the game $v_1^* = \underline{v}$

Solving TPZS games by LP (Player 2)

- The LP for player 2 has variables $v_2 \in \mathbb{R}$ and $p_2 \in \Delta_2$
- *Minimize* v_2 subject to the constraints

$$\begin{aligned} U(s_1, p_2) \leq v_2, &\quad \forall s_1 \in S_1 \\ p_2 \in \Delta_2 & \end{aligned}$$

- The optimal solution is the minimax strategy p_2^* and the upper value of the game $v_2^* = \overline{v}$

Solving TPZS games: Stocktaking

- The two LPs are in fact *dual* to each other
- This means that in the optimum (p_1^*, v_1^*) of the first LP and the second LP (p_2^*, v_2^*), respectively, we obtain $v_1^* = v_2^*$
- The profile of maximin/minmax strategies (p_1^*, p_2^*) is a NE

Computational experiments

- Two-player games
- Randomly generated (normal distribution) general-sum vs. zero-sum games with different sizes of strategy spaces
- SCIP solver
- The state-of-the-art multilinear formulation vs. standard LP

Computational experiments (1)

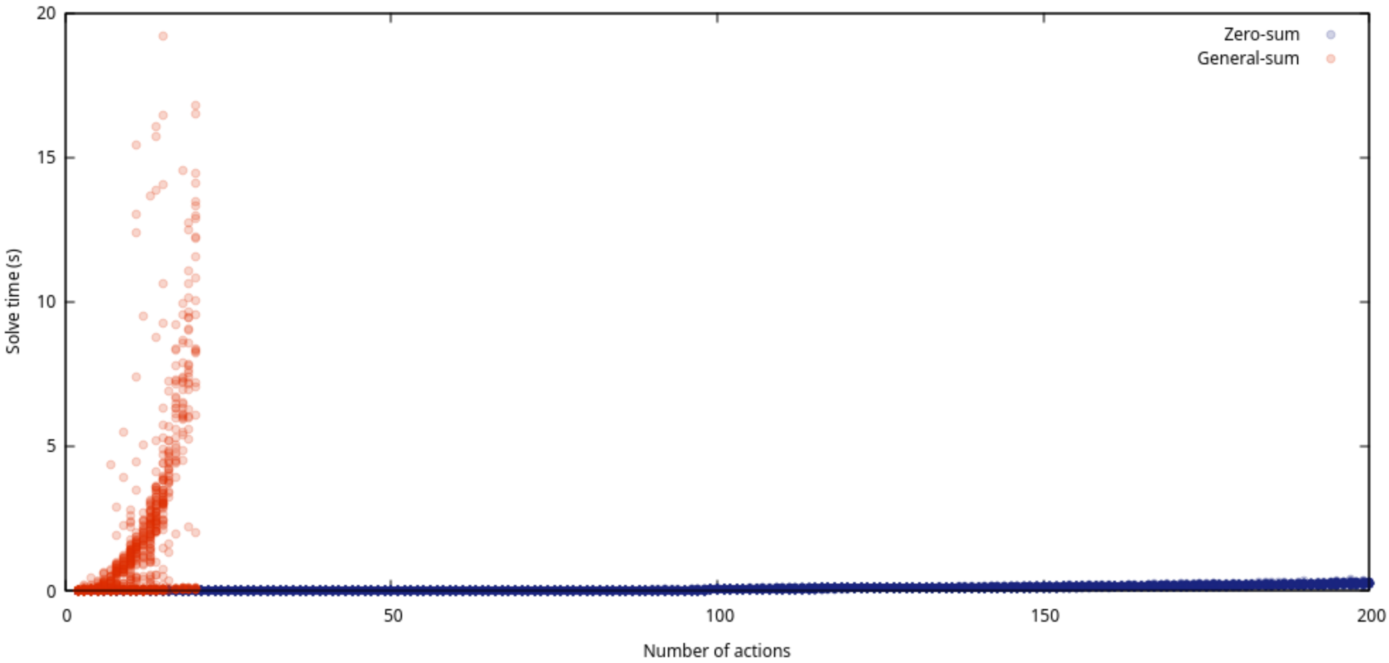

T. Votroubek - General-sum games vs. zero-sum games

Computational experiments (2)

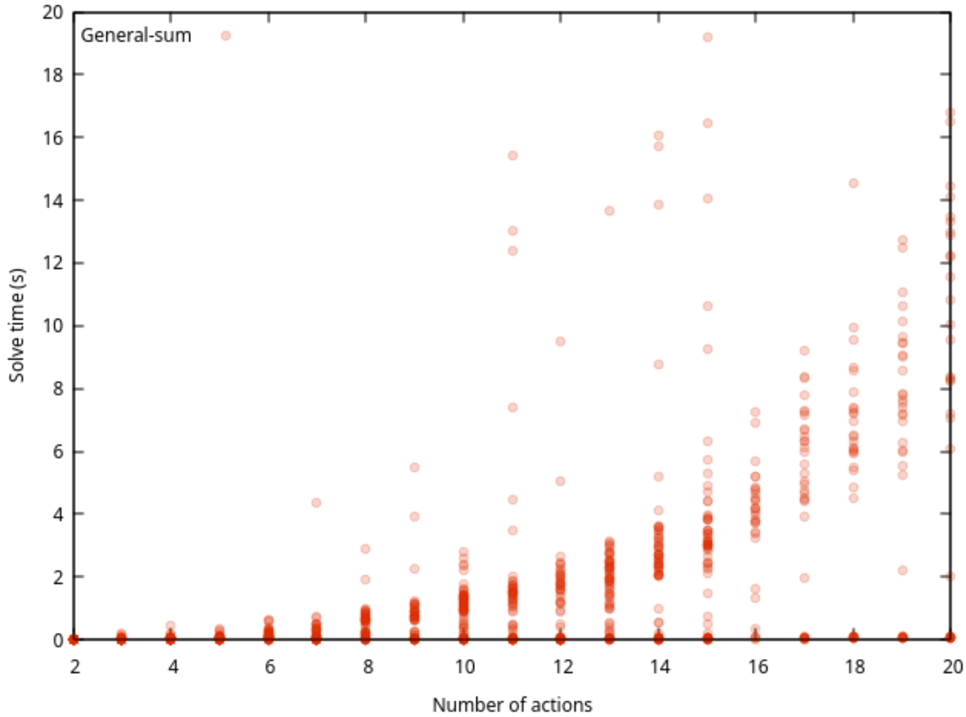

Computational experiments (3)

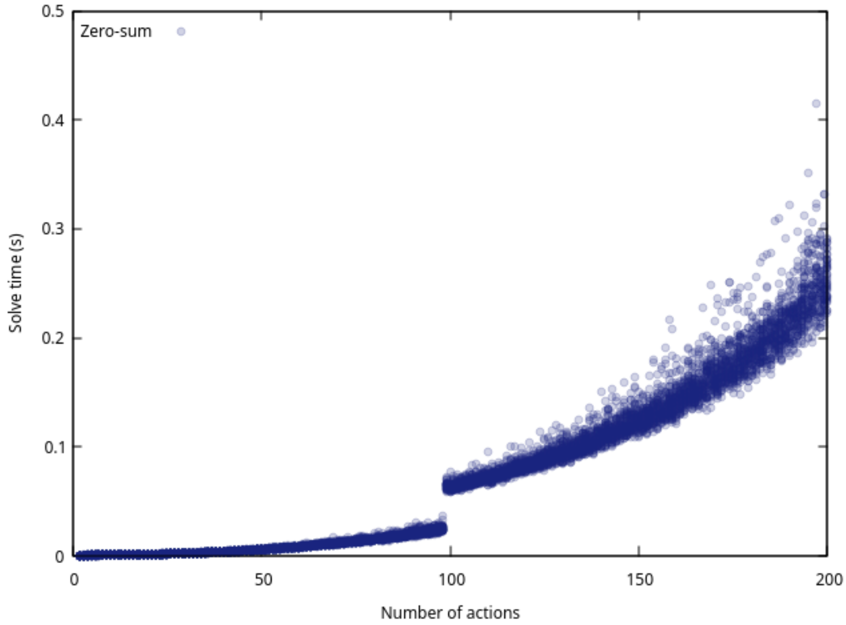

T. Votroubek - Zero-sum games (note the artefact at 100)