The power of mediators: Price of anarchy and stability in Bayesian games with submodular social welfare

Kaito Fujii (National Institute of Informatics)

10 July 2025 @ EC 2025

Example of valid utility games

Players simultaneously choose a resource to share

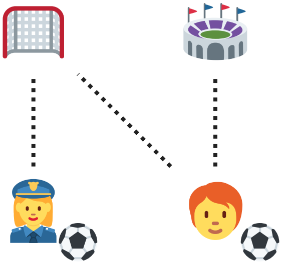

Resources chosen by multiple players are partitioned in a prespecified way

Example: 🧑 is prioritized over 👮‍♀️

Example of valid utility games

Players simultaneously choose a resource to share

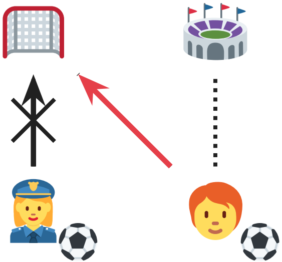

Resources chosen by multiple players are partitioned in a prespecified way

Example: 🧑 is prioritized over 👮‍♀️

No player can benefit from deviations

$\Updownarrow$

Worst **Nash equilibrium** $= 1$

Example of valid utility games

Players simultaneously choose a resource to share

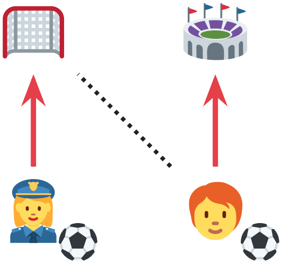

Resources chosen by multiple players are partitioned in a prespecified way

Example: is prioritized over

Optimal social welfare $= 2$

Example of valid utility games

Players simultaneously choose a resource to share

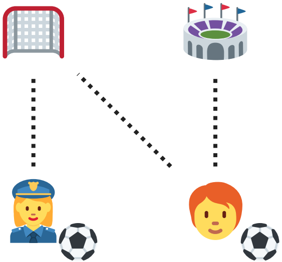

Resources chosen by multiple players are partitioned in a prespecified way

Example: 🧑 is prioritized over 👮‍♀️

No player can benefit from deviations

⇕

$$\text{PoA} := \frac{\text{Worst } \textbf{Nash equilibrium}}{\textbf{Optimal social welfare}} = \frac{1}{2}$$

(price of anarchy)

Example of valid utility games

Players simultaneously choose a resource to share

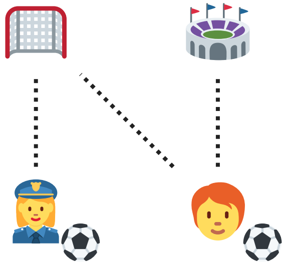

Resources chosen by multiple players are partitioned in a prespecified way

Example: is prioritized over

No player can benefit from deviations

$$\Updownarrow$$

PoA (price of anarchy) $:= \dfrac{\text{Worst } \textbf{Nash equilibrium}}{\textbf{Optimal social welfare}} = \dfrac{1}{2}$

Theorem [Vetta 2002]

PoA ≥ 0.5 in any valid utility game

Example of valid utility games

Q **How good or bad social welfare can be achieved by mediators**

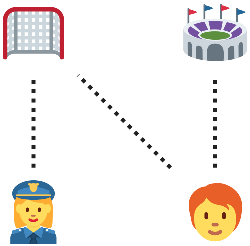

A mediator sends recommendations
(realizes correlated equilibrium)

Example of valid utility games

Q **How good or bad social welfare can be achieved by mediators**

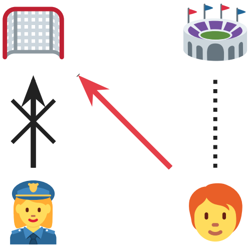

A mediator sends recommendations

(realizes correlated equilibrium)

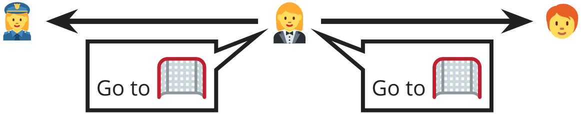

Example of valid utility games

Q **How good or bad social welfare can be achieved by mediators**

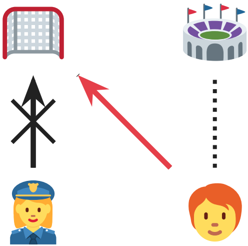

A mediator sends recommendations

(realizes correlated equilibrium)

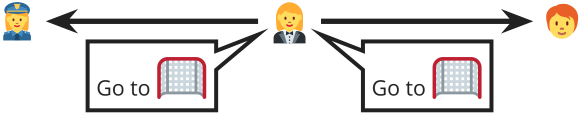

Theorem **[Roughgarden 2015]**

PoA ≥ 0.5 in any valid utility game for **correlated equilibria**

Example of Bayesian valid utility games

The set of actions for each player changes depending on their type

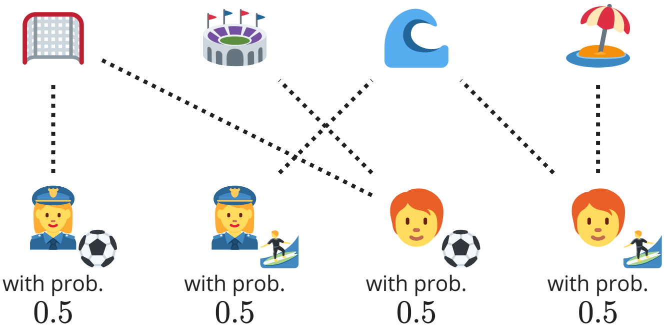

Q **How do mediators work in Bayesian games?**

Example of Bayesian valid utility games

The set of actions for each player changes depending on their type

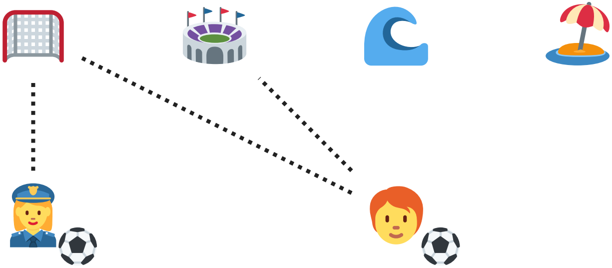

with probability $1/4$

👮‍♀️ does not know 🧑's type

🧑 does not know 👮‍♀️'s type

Q **How do mediators 🤵‍♀️ work in Bayesian games?**

Example of Bayesian valid utility games

The set of actions for each player changes depending on their type

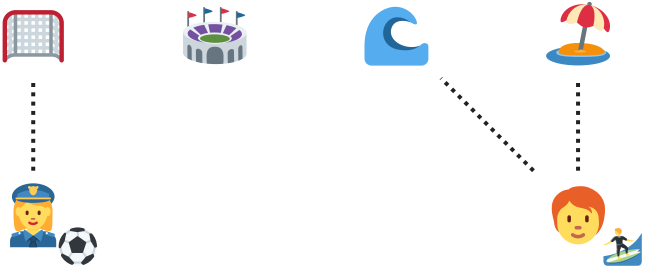

with probability $1/4$

👮‍♀️ does not know 🧑's type

🧑 does not know 👮‍♀️'s type

Q **How do mediators 🤵‍♀️ work in Bayesian games?**

Various Bayes correlated equilibria [Forges'93]

Bayes correlated equilibria (= correlated eq. in Bayesian games)

have many variants with various communication protocols

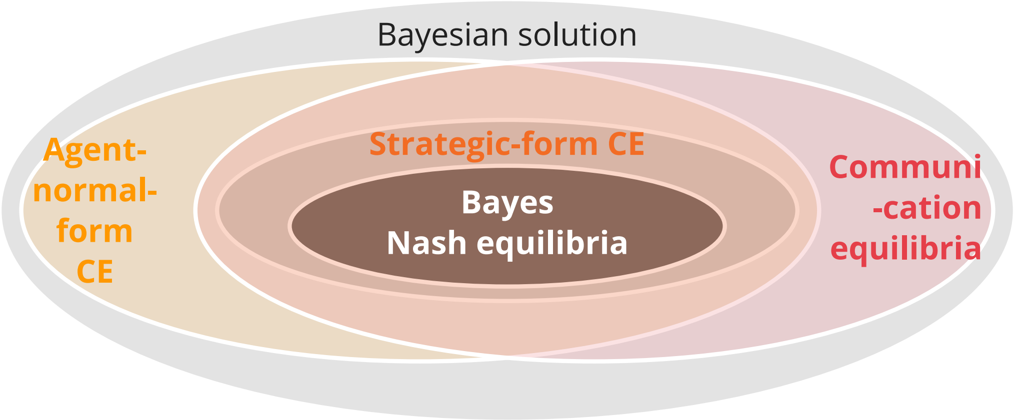

Equilibria realized by with bidirectional communication

1. Each player privately tells their types to the mediator

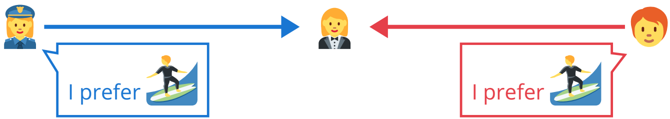

2. The mediator privately sends a recommendation to each player

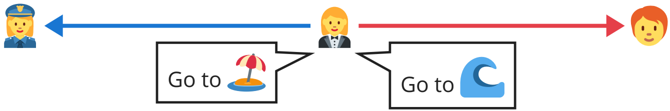

Communication equilibria [Myerson'82, Forges'86]

Equilibria realized by 🤵‍♀️ with bidirectional communication

1. Each player privately tells their types to the mediator 🤵‍♀️

   ← **No incentive to tell an untrue type**

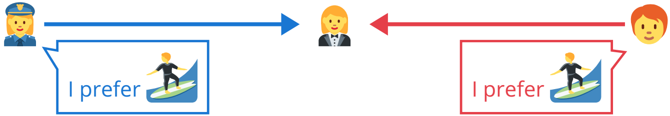

2. The mediator 🤵‍♀️ privately sends a recommendation to each player

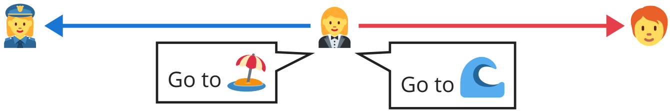

Communication equilibria [Myerson'82, Forges'86]

Equilibria realized by 🤵‍♀️ with bidirectional communication

1. Each player privately tells their types to the mediator 🤵‍♀️

← **No incentive to tell an untrue type**

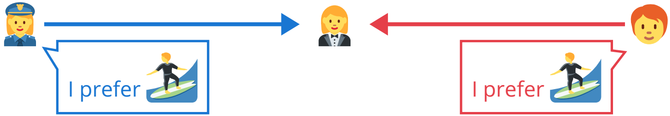

2. The mediator 🤵‍♀️ privately sends a recommendation to each player

← **No incentive to disobey the recommendation**

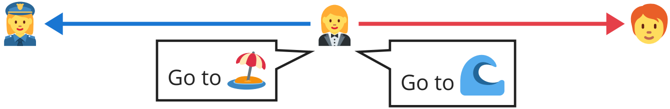

Bayesian solutions [Forges'93]

Mediator knows the true types in advance

1. Each player privately tells their **true** types to the mediator

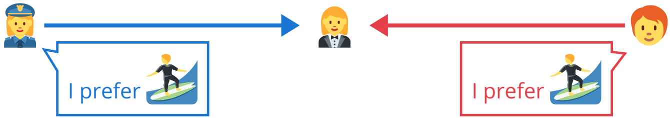

2. The mediator privately sends a recommendation to each player

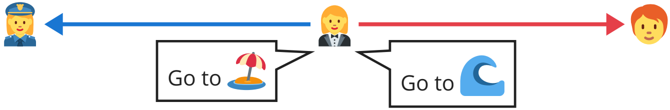

Bayesian solutions [Forges'93]

Mediator knows the true types in advance

1. Each player privately tells their **true** types to the mediator

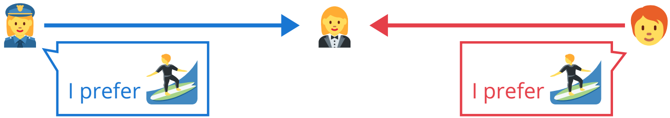

2. The mediator privately sends a recommendation to each player

← **No incentive to disobey the recommendation**

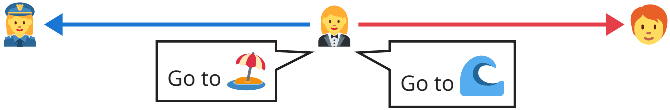

Our results

For various equilibrium concepts, we provide PoA and PoS bounds

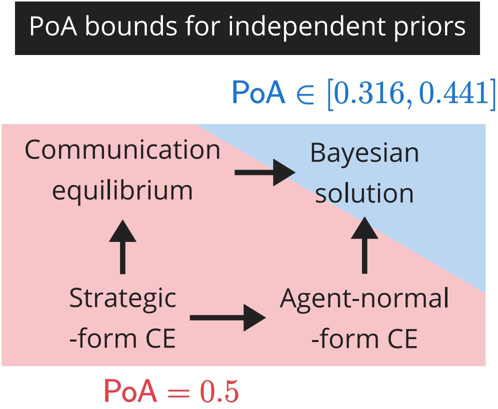

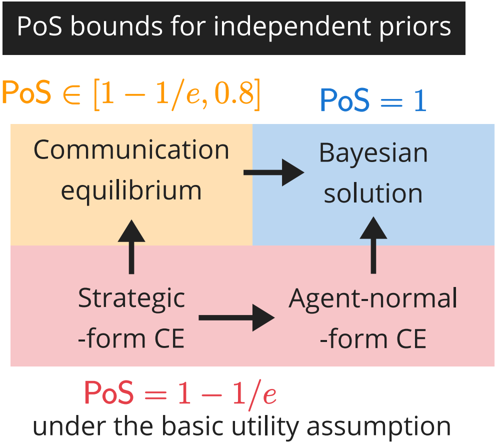

We also obtained PoA and PoS bounds for the correlated prior case

Table of Contents

Our setting: Bayesian valid utility games

Our technique: Strategy-representability gap

Other results

Notations for Bayesian games

$N = \{1, 2, \ldots, n\}$ players — $N = \{$👮‍♀️, 🧑‍🦰$\}$

Θ_i finite set of types for player $i \in N$ — $\Theta_{👮‍♀️} = \Theta_{🧑‍🦰} = \{$⚽, 🏄$\}$

$A_i^{\theta_i}$ finite set of actions for player $i \in N$ with type $\theta_i \in \Theta_i$ — $A_{🧑‍🦰}^{⚽} = \{$🥅, 🏟️$\}$

$\Theta = \prod_{i \in N} \Theta_i$ type profiles

$\rho \in \Delta(\Theta)$ prior distribution over type profiles — $\rho($⚽, ⚽$) = 1/4$

Notations for Bayesian games

$N = \{1, 2, \ldots, n\}$ players — $N = \{$👮‍♀️, 🧑$\}$

Θ_i finite set of types for player $i \in N$ — $\Theta_{👮‍♀️} = \Theta_{🧑} = \{$⚽, 🏄$\}$

$A_i^{\theta_i}$ finite set of actions for player $i \in N$ with type $\theta_i \in \Theta_i$ — $A_{🧑}^{⚽} = \{$🥅, 🏟️$\}$

$\Theta = \prod_{i \in N} \Theta_i$ type profiles

$\rho \in \Delta(\Theta)$ prior distribution over type profiles — $\rho($⚽, ⚽$) = 1/4$

Two settings in this study

- ρ is **independent** ($\exists \theta_i \in \Delta(\Theta_i)$ for each $i \in N$ s.t. $\rho(\theta) = \prod_{i \in N} \rho_i(\theta_i)$ for all $\theta \in \Theta$)
- ρ is **correlated** (no assumption on ρ)

Let $E = \bigcup_{i \in N} \bigcup_{\theta_i \in \Theta_i} A_i^{\theta_i}$ be the set of all possible actions

Submodular social welfare functions

Let $E = \bigcup_{i \in N} \bigcup_{\theta_i \in \Theta_i} A_i^{\theta_i}$ be the set of all possible actions

Assumption [Vetta 2002]

The social welfare function $f\colon 2^E \to \mathbb{R}$ is assumed to be

- **non-negative**: $f(X) \geq 0$ for any $X \subseteq E$
- **monotone**: $f(X \cup \{v\}) \geq f(X)$ for any $X \subseteq E$ and $v \in E$
- **submodular**: $f(X \cup \{v\}) - f(X) \geq f(Y \cup \{v\}) - f(Y)$ for any $X \subseteq Y \subseteq E$ and $v \in E \setminus Y$

Submodular social welfare functions

The marginal contribution to social welfare of each action decreases as other actions are added

Submodular social welfare functions

The marginal contribution to social welfare of each action decreases as other actions are added

$$f(\{🥅👮‍♀️\}) - f(\{\})$$

The increase in social welfare when no one attended yet

Submodular social welfare functions

The marginal contribution to social welfare of each action decreases as other actions are added

$$f(\{🥅_{👮‍♀️}\}) - f(\{\})$$

The increase in social welfare when no one attended yet

$$f(\{🥅_{👮‍♀️}, 🥅_{🧑}\}) - f(\{🥅_{🧑}\})$$

The increase in social welfare when other players already attended

Submodular social welfare functions

The marginal contribution to social welfare of each action decreases as other actions are added

$$\underbrace{f(\{\text{🥅}_{\text{👮‍♀️}}\}) - f(\{\})}_{\text{The increase in social welfare when no one attended yet}} \geq \underbrace{f(\{\text{🥅}_{\text{👮‍♀️}}, \text{🥅}_{\text{🧒}}\}) - f(\{\text{🥅}_{\text{🧒}}\})}_{\text{The increase in social welfare when other players already attended}}$$

Submodular social welfare functions

The marginal contribution to social welfare of each action decreases as other actions are added

$$\underbrace{f(\{🥅_{👮}\}) - f(\{\})}_{\text{The increase in social welfare when no one attended yet}} \geq \underbrace{f(\{🥅_{👮}, 🥅_{🧑}\}) - f(\{🥅_{🧑}\})}_{\text{The increase in social welfare when other players already attended}}$$

Intuitively, this assumption is **substitutability** among players' actions

- Note that we assume this property even among the same player's actions

Bayesian valid utility games

$v_i \colon A \to \mathbb{R}_{\geq 0}$ utility function for each player $i \in N$,

where $A = \prod_{i \in N} \left(\bigcup_{\theta_i \in \Theta_i} A_i^{\theta_i} \right)$ is the set of all possible action profiles

Bayesian valid utility games

$v_i \colon A \to \mathbb{R}_{\geq 0}$ utility function for each player $i \in N$,

where $A = \prod_{i \in N} \left(\bigcup_{\theta_i \in \Theta_i} A_i^{\theta_i} \right)$ is the set of all possible action profiles

Assumption [Vetta 2002]

- $\sum_{i \in N} v_i(a) \leq f(\{a_1, \ldots, a_n\})$ for any $a \in A$ (total utility condition)
- $v_i(a) \geq f(\{a_1, \ldots, a_n\}) - f(\{a_j \mid j \in N \setminus \{i\}\})$ for any $i \in N$ and $a \in A$ (marginal contribution condition)

Bayesian valid utility games

$v_i : A \to \mathbb{R}_{\geq 0}$ utility function for each player $i \in N$,

where $A = \prod_{i \in N} \left(\bigcup_{\theta_i \in \Theta_i} A_i^{\theta_i} \right)$ is the set of all possible action profiles

Assumption [Vetta 2002]

- $\sum_{i \in N} v_i(a) \leq f(\{a_1, \ldots, a_n\})$ for any $a \in A$ (total utility condition)
- $v_i(a) \geq f(\{a_1, \ldots, a_n\}) - f(\{a_j \mid j \in N \setminus \{i\}\})$ for any $i \in N$ and $a \in A$ (marginal contribution condition)

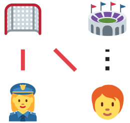

The sum of utility values is at most $f($🥅$)$

The contribution of 👮‍♀️ is at least $f($🥅$) - f($🥅$) = 0$

Example: 🧑 gets all, 👮‍♀️ gets all, two players share equally, or both get 0

Table of Contents

Our setting: Bayesian valid utility games

Our technique: Strategy-representability gap

Other results

For an equilibrium class $\Pi \subseteq \Delta(A)^{\Theta}$, PoA is defined as

the social welfare achieved by the worst equilibrium

$$\text{PoA}_{\Pi} \triangleq \frac{\inf_{\pi \in \Pi} \mathbb{E}_{\theta \sim \rho}\left[\mathbb{E}_{a \sim \pi(\theta)}[v_{\text{SW}}(a)]\right]}{\mathbb{E}_{\theta \sim \rho}\left[\max_{a^* \in A^{\theta}} v_{\text{SW}}(a^*)\right]}$$

the optimal social welfare

Challenge optimal action a_i^* depends on the other players' types θ_{-i}

Strategy-representability gap (SR gap)

the social welfare achieved by the worst equilibrium

$$\mathrm{PoA}_{\Pi} \triangleq \frac{\inf_{\pi \in \Pi} \mathbb{E}_{\theta \sim \rho}\left[\mathbb{E}_{a \sim \pi(\theta)}\left[v_{\mathrm{SW}}(a)\right]\right]}{\mathbb{E}_{\theta \sim \rho}\left[\max_{a^* \in A^\theta} v_{\mathrm{SW}}(a^*)\right]}$$

the optimal social welfare

Strategy-representability gap (SR gap)

the social welfare achieved by the worst equilibrium

$$\mathrm{PoA}_{\Pi} \triangleq \frac{\inf_{\pi\in\Pi} \mathbb{E}_{\theta\sim\rho}\left[\mathbb{E}_{a\sim\pi(\theta)}[v_{\mathsf{SW}}(a)]\right]}{\mathbb{E}_{\theta\sim\rho}\left[\max_{a^*\in A^\theta} v_{\mathsf{SW}}(a^*)\right]}$$

the optimal social welfare

the social welfare achieved by the optimal strategy profile

$$= \underbrace{\frac{\inf_{\pi\in\Pi} \mathbb{E}_{\theta\sim\rho}\left[\mathbb{E}_{a\sim\pi(\theta)}[v_{\mathsf{SW}}(a)]\right]}{\max_{s^*\in S} \mathbb{E}_{\theta\sim\rho}[v_{\mathsf{SW}}(s^*(\theta))]}}_{\text{reduced to the non-Bayesian case}} \cdot \underbrace{\frac{\max_{s^*\in S} \mathbb{E}_{\theta\sim\rho}[v_{\mathsf{SW}}(s^*(\theta))]}{\mathbb{E}_{\theta\sim\rho}\left[\max_{a\in A^\theta} v_{\mathsf{SW}}(a)\right]}}_{\textbf{SR gap}}$$

Strategy-representability gap (SR gap)

$S_i = \prod_{\theta_i \in \Theta_i} A_i^{\theta_i}$ the set of strategies for $i \in N$

$s_i \in S_i$ determines an action $s_i(\theta_i)$ for θ_i

$S \triangleq \prod_{i \in N} S_i$ and $s(\theta) \triangleq (s_1(\theta_1), \dots, s_n(\theta_n))$

the social welfare achieved by the worst equilibrium

$$\mathrm{PoA}_\Pi \triangleq \frac{\inf_{\pi \in \Pi} \mathbb{E}_{\theta \sim \rho}\left[\mathbb{E}_{a \sim \pi(\theta)}[v_{\mathrm{SW}}(a)]\right]}{\mathbb{E}_{\theta \sim \rho}\left[\max_{a^* \in A^\theta} v_{\mathrm{SW}}(a^*)\right]}$$

the optimal social welfare

the social welfare achieved by the optimal strategy profile

$$= \underbrace{\frac{\inf_{\pi \in \Pi} \mathbb{E}_{\theta \sim \rho}\left[\mathbb{E}_{a \sim \pi(\theta)}[v_{\mathrm{SW}}(a)]\right]}{\max_{s^* \in S} \mathbb{E}_{\theta \sim \rho}[v_{\mathrm{SW}}(s^*(\theta))]}}_{\text{reduced to the non-Bayesian case}} \cdot \underbrace{\frac{\max_{s^* \in S} \mathbb{E}_{\theta \sim \rho}[v_{\mathrm{SW}}(s^*(\theta))]}{\mathbb{E}_{\theta \sim \rho}\left[\max_{a \in A^\theta} v_{\mathrm{SW}}(a)\right]}}_{\textbf{SR gap}}$$

Strategy-representability gap (SR gap)

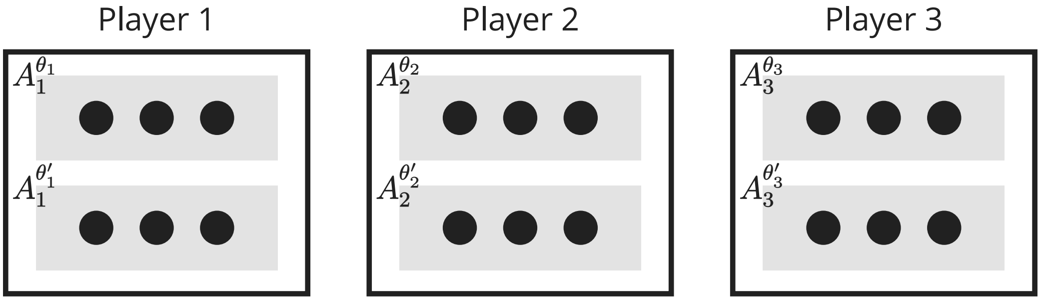

For each player, one block is chosen according to a known distribution

Strategy-representability gap (SR gap)

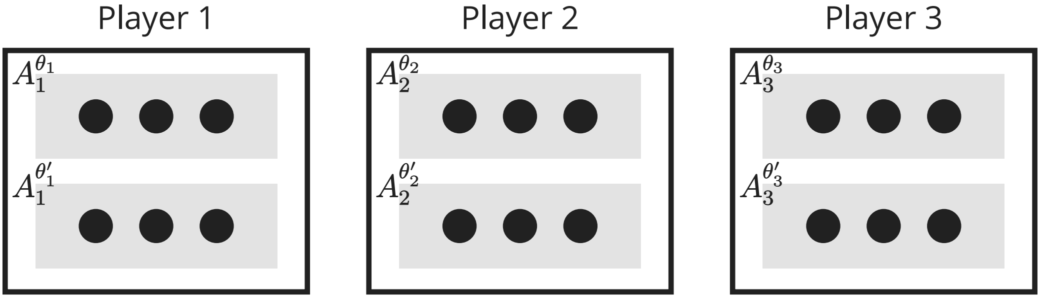

For each player, one block is chosen according to a known distribution

$$\mathrm{SRgap} \triangleq \frac{\text{Choose one element from each block, and then blocks are selected}}{\text{Blocks are selected, and then choose one element from each block}}$$

Strategy-representability gap (SR gap)

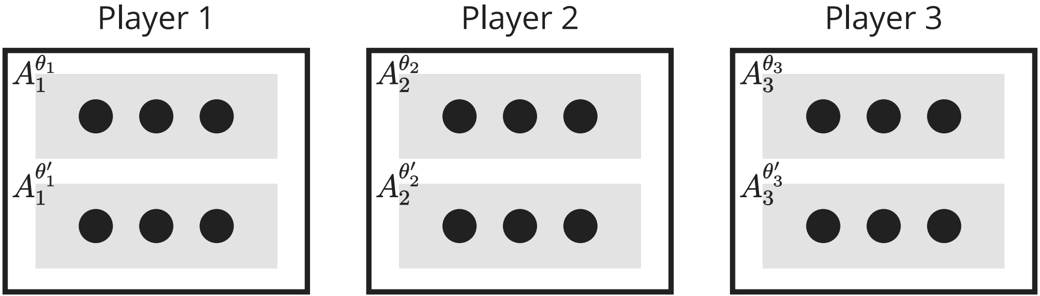

For each player, one block is chosen according to a known distribution

$$\mathrm{SRgap} \triangleq \frac{\text{Choose one element from each block, and then blocks are selected}}{\text{Blocks are selected, and then choose one element from each block}}$$

Strategy-representability gap (SR gap)

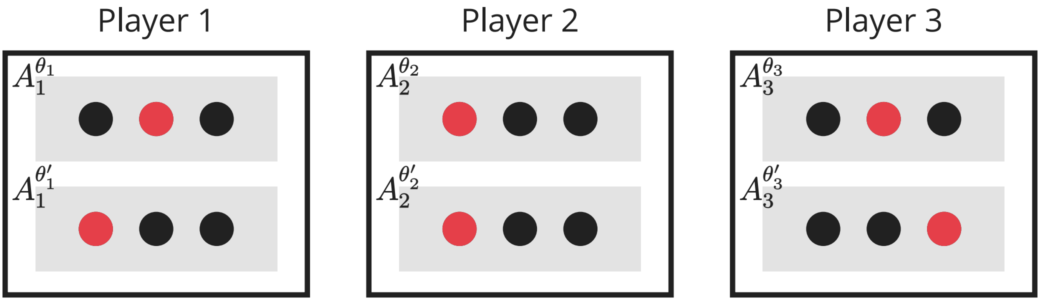

For each player, one block is chosen according to a known distribution

$$\text{SRgap} \triangleq \frac{\text{Choose one element from each block, and then blocks are selected}}{\text{Blocks are selected, and then choose one element from each block}}$$

Strategy-representability gap (SR gap)

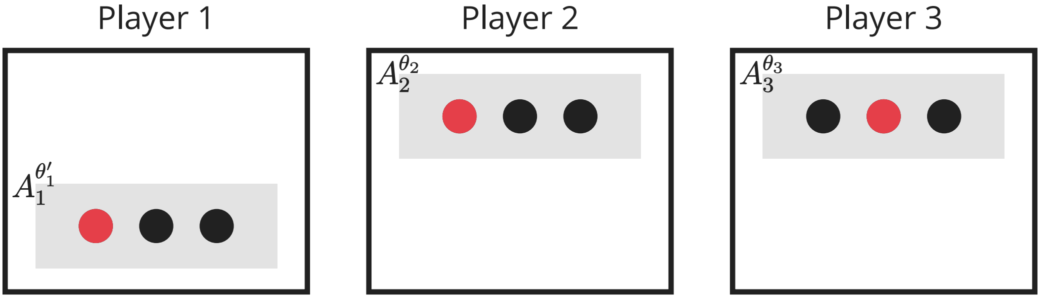

For each player, one block is chosen according to a known distribution

$$\mathrm{SRgap} \triangleq \frac{\text{Choose one element from each block, and then blocks are selected}}{\text{Blocks are selected, and then choose one element from each block}}$$

Strategy-representability gap (SR gap)

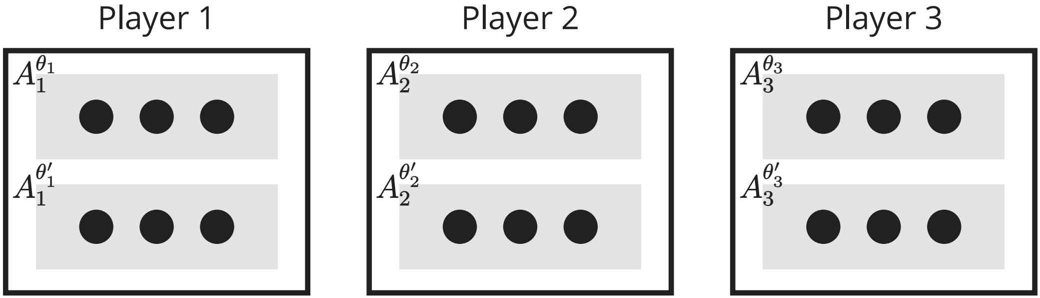

For each player, one block is chosen according to a known distribution

$$\mathrm{SRgap} \triangleq \frac{\text{Choose one element from each block, and then blocks are selected}}{\text{Blocks are selected, and then choose one element from each block}}$$

Strategy-representability gap (SR gap)

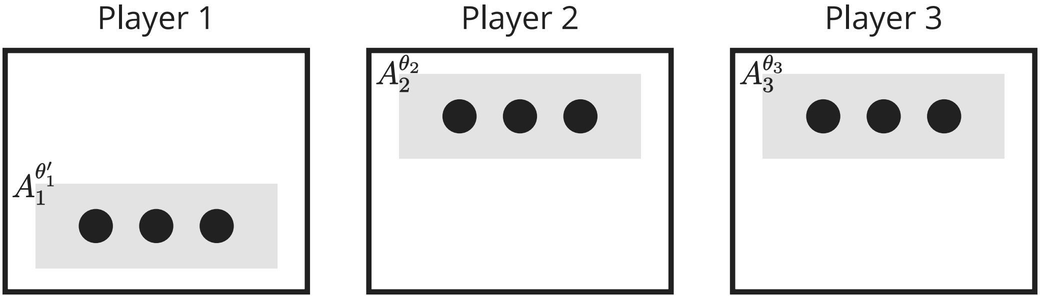

For each player, one block is chosen according to a known distribution

$$\mathrm{SRgap} \triangleq \frac{\text{Choose one element from each block, and then blocks are selected}}{\text{Blocks are selected, and then choose one element from each block}}$$

Strategy-representability gap (SR gap)

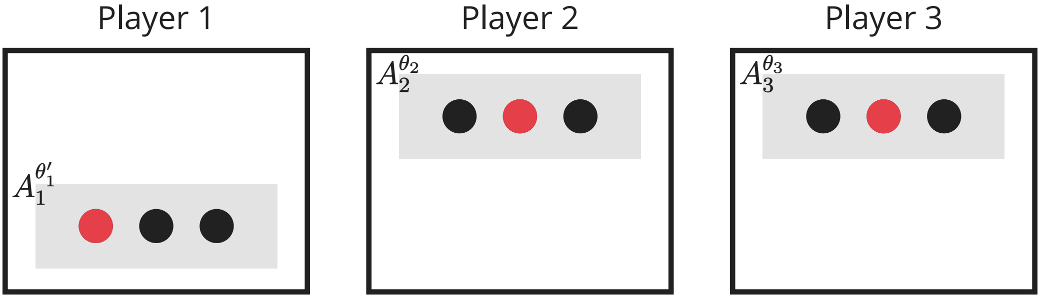

For each player, one block is chosen according to a known distribution

$$\text{SRgap} \triangleq \frac{\text{Choose one element from each block, and then blocks are selected}}{\text{Blocks are selected, and then choose one element from each block}}$$

SR gap lower bound (independent case)

Theorem

If ρ is independent, $\mathrm{SRgap} \geq 1 - 1/e$, and this bound is tight

Lower bound based on the correlation gap bound [Vondrák'07]

Upper bound

Optimal social welfare: n

$\because$ There exists a perfect matching w.h.p.

Optimal strategy profile: $\approx (1 - 1/e)n$

$\because$ The expected probability that each resource is chosen can be upper-bounded

SR gap lower bound (correlated case)

Theorem

$\text{SRgap} = \Omega(1/\sqrt{n})$, and this bound is tight

Lower bound (complicated)

Upper bound $\Theta_1 = \cdots = \Theta_n = [n]^k$, where $k = \sqrt{n}$ $\quad j \sim [k]$ and $\ell_1, \ldots, \ell_k \sim [n]$

Types $\{(\ell_1, \ldots, \ell_{j-1}, t, \ell_{j+1}, \ldots, \ell_k) \mid t \in [n]\}$ are randomly assigned to n players

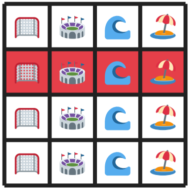

1st action

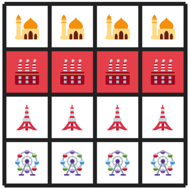

2nd action

$E = [n] \times [k]$ set of resources

The hth action of type ℓ is to choose $(h, \ell_h) \in E$

Optimal social welfare: n

Optimal strategy profile: $\leq k + n/k = 2\sqrt{n}$

Table of Contents

Our setting: Bayesian valid utility games

Our technique: Strategy-representability gap

Other results

Improved PoA lower bound for com. eq.

Proposition

If ρ is independent, $\mathrm{PoA}_{\mathrm{Com.Eq.}} \geq 0.5$, which improves on the SR-gap approach

Based on the smoothness arguments for Bayes–Nash equilibria

[Roughgarden'15, Syrgkanis'12]

The key step of their proof:

Swapping θ_i and θ'_i in $\theta \sim \rho$ and $\theta' \sim \rho$ using the independence of ρ

$\leftarrow$ Incentive constraints for misreporting θ'_i instead of θ_i can be used

Remark The same result also holds for agent-normal-form CE

PoA upper bound for Bayesian solutions

Proposition

$\mathrm{PoA}_{\mathrm{BS}} \leq \dfrac{1 - 1/\sqrt{e}}{3/2 - 1/\sqrt{e}} \approx 0.4403$ for some example with independent ρ

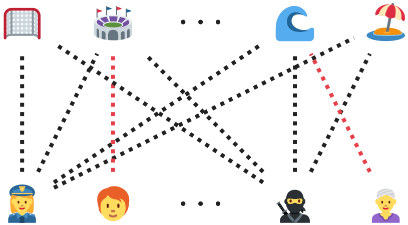

Odd players are connected to all resources
Even players are connected to random one
Odd players are prioritized over even ones

Bad Bayesian solution:
Each $(2k-1)$th player is recommended to choose the $(2k)$th player's action

Optimal: $\approx \underbrace{n/2}_{\text{even}} + \underbrace{(1 - 1/\sqrt{e})n}_{\text{odd}},$ **Bayesian solution**: $\approx (1 - 1/\sqrt{e})n$

For various equilibrium concepts, we provide PoA and PoS bounds

PoA bounds for independent priors

PoA $\in [0.316, 0.441]$

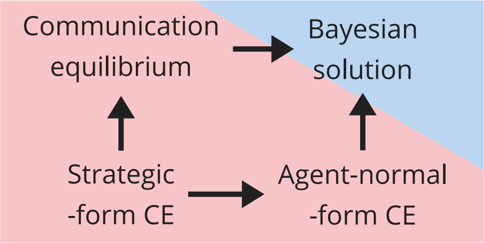

PoA $= 0.5$

PoS bounds for independent priors

PoS $\in [1 - 1/e, 0.8]$

PoS $= 1$

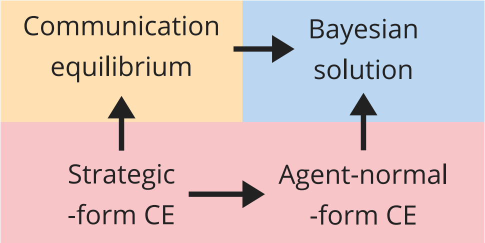

PoS $= 1 - 1/e$
under the basic utility assumption

We also obtained PoA and PoS bounds for the correlated prior case

Reference

- Françoise Forges. 1986. An approach to communication equilibria. *Econometrica*, 1375–1385.
- Françoise Forges. 1993. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision* 35, 277–310.
- John C. Harsanyi. 1967. Games with Incomplete Information Played by “Bayesian” Players, I–III. *Management Science* 14(3):159–182, 14(5):320–334, 14(7):486–502.
- Tim Roughgarden. 2015a. Intrinsic Robustness of the Price of Anarchy. *Journal of the ACM* 62(5), 32:1–32:42.
- Tim Roughgarden. 2015b. The Price of Anarchy in Games of Incomplete Information. *ACM Transactions on Economics and Computation* 3(1), 6:1–6:20.
- Adrian Vetta. 2002. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *FOCS 2002*. 416–425.
- Jan Vondrák. 2007. *Submodularity in combinatorial optimization*. Ph.D. Dissertation.

- Illustrations: ”Twemoji” by Twitter, Inc and other contributors is licensed under CC BY 4.0

Type prior distribution

Types $\theta = (\theta_1, \ldots, \theta_n)$ are generated from prior distribution $\rho \in \Delta(\Theta)$

- All players know ρ as common knowledge
- Each player $i \in N$ knows their own type θ_i but not the others' $\theta_{-i} = (\theta_j)_{j \in N \setminus \{i\}}$

Two settings in this study

- ρ is **independent** ($\exists \theta_i \in \Delta(\Theta_i)$ for each $i \in N$ s.t. $\rho(\theta) = \prod_{i \in N} \rho_i(\theta_i)$ for all $\theta \in \Theta$)
  - Types represent each player's preferences or attributes
- ρ is **correlated** (no assumption on ρ)
  - Types represent each player's weather or traffic conditions

Submodularity

Marginal gain of each elem. decreases as elements are obtained

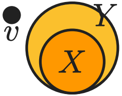

$f(v|X) \triangleq f(X \cup \{v\}) - f(X)$

marginal gain of adding $v \in E$ to $X \subseteq E$

$f : 2^E \to \mathbb{R}$ is submodular

$\overset{\triangle}{\Leftrightarrow}$ For all $X \subseteq Y$ with $X, Y \subseteq E$ and $v \in E \setminus Y$, we have $f(v|X) \geq f(v|Y)$

e.g.) $X = \{$🍌$\}$, $Y = \{$🍌, 🍱$\}$ and $v =$ 🍕

$f(\{$🍕, 🍌$\}) - f(\{$🍌$\}) \geq f(\{$🍕, 🍌, 🍱$\}) - f(\{$🍌, 🍱$\})$

Monotonicity

The function value is non-decreasing when elements are added

$f(v|X) \triangleq f(X \cup \{v\}) - f(X)$

marginal gain of adding $v \in E$ to $X \subseteq E$

$f : 2^E \to \mathbb{R}$ is monotone

$\overset{\triangle}{\Leftrightarrow}$ For all $X \subseteq E$ and $v \in E$, we have $f(v|X) \geq 0$

e.g.) $X = \{$🍌, 🍱$\}$ and $v =$ 🍕 $f(\{$🍕, 🍌, 🍱$\}) - f(\{$🍌, 🍱$\}) \geq 0$

Summary

For various equilibrium concepts, we provide PoA and PoS bounds

	BNE, SF/ANFCE, ANF/SFCCE	Com.Eq.	BS, ANF/SFCBS
PoA (v, i)	$1/2$	$1/2$	$\in \left[\frac{1-1/e}{2}, 0.441\right]$
PoA (v, c)	$\Theta\left(\frac{1}{\sqrt{n}}\right)$	$\Theta\left(\frac{1}{\sqrt{n}}\right)$	$\Theta\left(\frac{1}{\sqrt{n}}\right)$
PoS (b, i)	$1 - 1/e$	$\leq 4/5$	1
PoS (b, c)	$\Theta\left(\frac{1}{\sqrt{n}}\right)$	$\Omega\left(\frac{1}{\sqrt{n}}\right)$	1

"v"=valid utility games, "b"=basic utility games,
"i"=type prior distribution ρ is independent, "c"=ρ can be correlated

Communication equilibria

$N = \{1, 2, \ldots, n\}$ players — $N = \{$👮‍♀️, 🧑‍🦰$\}$

Θ_i finite set of types for player $i \in N$ — $\Theta_{👮‍♀️} = \Theta_{🧑‍🦰} = \{$⚽, 🏄$\}$

$A_i^{\theta_i}$ finite set of actions for player $i \in N$ with type $\theta_i \in \Theta_i$ — $A_{🧑‍🦰}^{⚽} = \{$🥅, 🏟️$\}$

$\Theta = \prod_{i \in N} \Theta_i$ type profiles

$\rho \in \Delta(\Theta)$ prior distribution over type profiles — $\rho($⚽, ⚽$) = 1/4$

A distribution $\pi \in \prod_{\theta \in \Theta} \Delta(A^\theta)$ is a communication equilibrium if for any $i \in N$, $\theta_i, \theta_i' \in \Theta_i$, and $\phi\colon A_i^{\theta_i'} \to A_i^{\theta_i}$, it holds that

$$\mathop{\mathbb{E}}_{\theta_{-i} \sim \rho|_{\theta_i}} \left[\mathop{\mathbb{E}}_{a \sim \pi(\theta)} [v_i(a)] \right] \geq \mathop{\mathbb{E}}_{\theta_{-i} \sim \rho|_{\theta_i}} \left[\mathop{\mathbb{E}}_{a \sim \pi(\theta_i', \theta_{-i})} [v_i(\phi(a_i), a_{-i})] \right].$$