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

Kaito Fujii (National Institute of Informatics)

14 October 2025 @ UTokyo Microeconomic Theory Workshop

Overview of today's talk

1. **Various definitions of correlated equilibria in Bayesian games**
   - Definitions and relations of various correlated-equilibrium concepts [Forges'93]
   - Corresponding regret notions and learning dynamics [Fujii'25a]
2. **Welfare guarantees for submodular social welfare** [Fujii'25b]
   - Introduction of Bayesian games with submodular social welfare
   - Gaps in welfare guarantees for Bayes correlated equilibria

Table of Contents

Correlated equilibria

Players’ actions can be correlated via a traffic signal

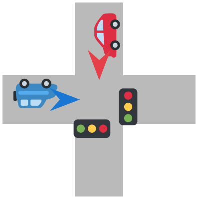

		Go	Stop
	Go	0, 0	4, 3
	Stop	3, 4	1, 1

Correlated equilibria:

Players’ actions can be correlated

infinitely many including Nash eq.

e.g.) (Go, Stop) with prob. $1/2$

(Stop, Go) with prob. $1/2$

Correlated equilibria

$N = \{1, 2, \dots, n\}$ players — $N = \{$🚙, 🚗$\}$

A_i finite set of actions for player $i \in N$ — $A_i = \{\text{Go}, \text{Stop}\}$

$A = A_1 \times A_2 \times \cdots \times A_n$ set of action profiles

$v_i \colon A \to \mathbb{R}$ utility function for player $i \in N$ — $u_{\text{🚙}}(\text{Go}, \text{Stop}) = 4$

Definition

A distribution over action profiles $\pi \in \Delta(A)$ is a correlated equilibrium

$\overset{\triangle}{\Leftrightarrow}$ For any player $i \in N$ and deviation $\phi \colon A_i \to A_i$,

$$\mathop{\mathbb{E}}_{a \sim \pi}[v_i(\phi(a_i), a_{-i})] \leq \mathop{\mathbb{E}}_{a \sim \pi}[v_i(a)].$$

※ If π is a product distribution, this definition coincides with Nash equilibria

Correlated equilibria

Definition

A distribution over action profiles $\pi \in \Delta(A)$ is a correlated equilibrium $\overset{\triangle}{\Leftrightarrow}$ For any player $i \in N$ and deviation $\phi \colon A_i \to A_i$,

$$\mathop{\mathbb{E}}_{a\sim\pi}[v_i(\phi(a_i), a_{-i})] \le \mathop{\mathbb{E}}_{a\sim\pi}[v_i(a)].$$

	Go	Stop
Go	0, 0	4, 3
Stop	3, 4	1, 1

We can define a CE $\pi \in \Delta(A)$ as follows:

$\pi(\text{Go}, \text{Stop}) = 1/2$, $\pi(\text{Stop}, \text{Go}) = 1/2$

Each player cannot increase the payoff by any ϕ

e.g., $\phi(\text{Go}) = \text{Stop}$, $\phi(\text{Stop}) = \text{Stop}$ decreases it

No-regret dynamics

Algorithm Simulate **no-regret dynamics** converging to a CE

Players learn their strategy in repeated play of the same game

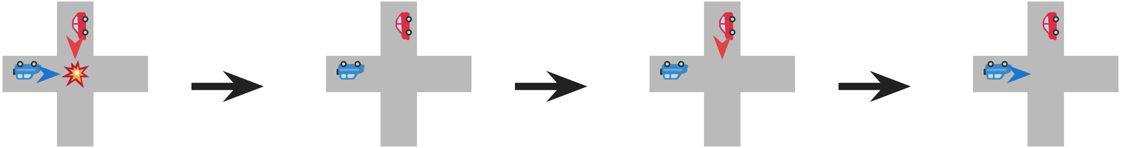

for $t = 1, 2, \ldots, T$ **do**

Each player $i \in N$ decides a (mixed) strategy $\pi_i^t \in \Delta(A_i)$

All players' strategies $(\pi_i^t)_{i \in N}$ are revealed to each other

Each player i obtains reward $\mathbb{E}[v_i(a^t)]$, where $a_i^t \sim \pi_i^t$ independently ($\forall i$)

Swap regret [Blum–Mansour'07]

$$R^T_{\mathrm{swap},i} \triangleq \max_{\phi:\, A_i \to A_i} \sum_{t=1}^{T} \underbrace{\mathbb{E}\left[v_i^t(\phi(a_i^t), a_{-i}^t)\right]}_{\text{reward in round } t \text{ if the actions are replaced according to } \phi} - \sum_{t=1}^{T} \underbrace{\mathbb{E}\left[v_i^t(a^t)\right]}_{\text{reward in round } t}$$

Theorem [Foster–Vohra'97, Hart–Mas-Collel'00, Blum–Mansour'07]

If swap regret of every player grows sublinearly in T,

the empirical distribution converges to a correlated equilibrium

The uniform mixture of action profiles of T rounds

Bayesian games (incomplete info. + common prior) [Harsanyi'67]

Players' types are generated from a common prior distribution

Each of 👮‍♀️ and 🧑 prefers 🌊 and 🌋 with prob. $1/2$ for each

(Each player knows the prior distribution only, not the others' types)

w.p. $1/4$

		🧑 type: 🌊	
		🌊	🌋
👮‍♀️ type: 🌊	🌊	4, 4	1, 0
	🌋	0, 1	3, 3

w.p. $1/4$

		🧑 type: 🌋	
		🌊	🌋
👮‍♀️ type: 🌊	🌊	4, 3	1, 1
	🌋	0, 0	3, 4

Notations for Bayesian games

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

A_i finite set of actions for player $i \in N$ — $A_1 = A_2 = \{$🌊, 🌋$\}$

Θ_i finite set of types for player $i \in N$ — $\Theta_1 = \Theta_2 = \{$type:🌊, type:🌋$\}$

$A = \prod_{i \in N} A_i$ action profiles, $\Theta = \prod_{i \in N} \Theta_i$ type profiles

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

$v_i : \Theta \times A \to \mathbb{R}$ utility function for player $i \in N$ — $v_1($type:🌊, type:🌊; 🌊, 🌋$) = 1$

Various Bayes correlated equilibria [Forges'93]

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

have many variants with various communication protocols

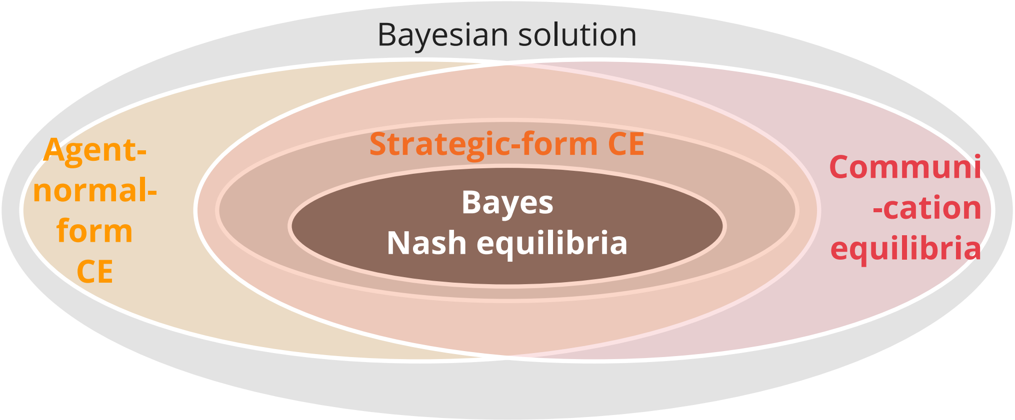

No-regret dynamics in Bayesian games

For $t = 1, 2, \ldots, T$**:**

Each player $i \in N$ decides a (mixed) strategy $\pi_i^t \in \Delta(A_i)^{\Theta_i}$

All players' strategies $(\pi_i^t)_{i \in N}$ are revealed to each other

Each player i obtains reward $\mathbb{E}[v_i(\theta^t; a^t)]$,

where $\theta^t \sim \rho$ and then $a_i^t \sim \pi_i^t(\theta_i^t)$ independently for each i

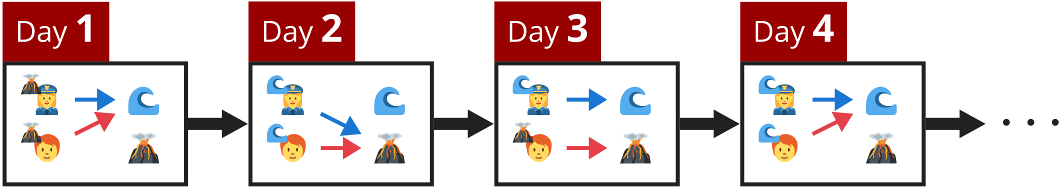

※ We consider the expected value w.r.t. θ and a in each round

Table of Contents

Strategic-form correlated equilibria

Strategic form of Bayesian games

A **strategy** $s_i : \Theta_i \to A_i$ is interpreted as an action

The set of all actions in this interpretation is $S_i := A_i^{\Theta_i}$

🤵‍♀️ **privately recommends an action for each type separately**

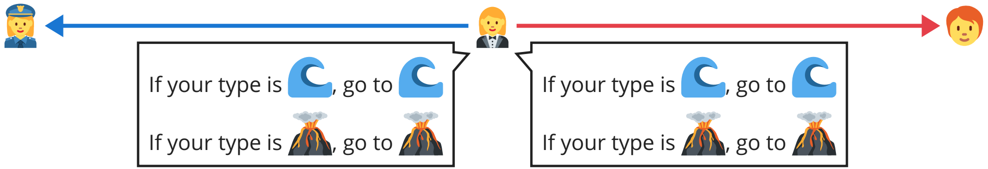

Strategic-form correlated equilibria

Strategic form of Bayesian games

A **strategy** $s_i : \Theta_i \to A_i$ is interpreted as an action

The set of all actions in this interpretation is $S_i := A_i^{\Theta_i}$

privately recommends an action for each type separately

$\leftarrow$ **No incentive to disobey the recommendation**

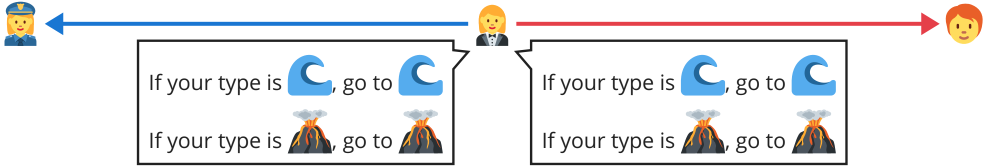

SFCE & Strategy swap regret

Definition

A distribution $\sigma \in \Delta(S_1 \times \cdots \times S_n)$ is an SFCE

$\stackrel{\triangle}{\Leftrightarrow}$ For any player $i \in N$, $\phi_{\mathsf{SF}}: S_i \to S_i$,

$$\mathop{\mathbb{E}}_{\theta\sim\rho}\left[\mathop{\mathbb{E}}_{s\sim\sigma}[v_i(\theta; s_1(\theta_1), \ldots, s_n(\theta_n))]\right] \geq \mathop{\mathbb{E}}_{\theta\sim\rho}\left[\mathop{\mathbb{E}}_{s\sim\sigma}[v_i(\theta; \phi_{\mathsf{SF}}(s_i)(\theta_i), s_{-i}(\theta_{-i}))]\right].$$

$$R^T_{\mathrm{SS},i} \triangleq \max_{\phi_{\mathsf{SF}}: S_i \to S_i} \sum_{t=1}^{T} \underbrace{\mathbb{E}\left[v_i(\phi_{\mathsf{SF}}(s_i^t)(\theta_i^t), a_{-i}^t)\right]}_{\text{reward in round } t \text{ if the strategies are replaced according to } \phi_{\mathsf{SF}}} - \sum_{t=1}^{T} \underbrace{\mathbb{E}\left[v_i(s_i^t(\theta_i^t), a_{-i}^t)\right]}_{\text{reward in round } t}$$

※ Each player chooses $\sigma_i^t \in \Delta(S_i)$ and generates $s_i^t \sim \sigma_i^t$

Definition

A distribution $\sigma \in \Delta(S_1 \times \cdots \times$

Choosing strategy $\phi_{\mathsf{SF}}(s_i)$ instead of recommended s_i

$\triangleq$ For any player $i \in N$, $\phi_{\mathsf{SF}} : S_i \to S_i$,

$$\mathop{\mathbb{E}}_{\theta\sim\rho}\left[\mathop{\mathbb{E}}_{s\sim\sigma}\left[v_i(\theta; s_1(\theta_1), \ldots, s_n(\theta_n))\right]\right] \geq \mathop{\mathbb{E}}_{\theta\sim\rho}\left[\mathop{\mathbb{E}}_{s\sim\sigma}\left[v_i(\theta; \phi_{\mathsf{SF}}(s_i)(\theta_i), s_{-i}(\theta_{-i}))\right]\right].$$

$$R^T_{\mathrm{SS},i} \triangleq \max_{\phi_{\mathsf{SF}} : S_i \to S_i} \sum_{t=1}^{T} \underbrace{\mathbb{E}\left[v_i(\phi_{\mathsf{SF}}(s_i^t)(\theta_i^t), a_{-i}^t)\right]}_{\text{reward in round } t \text{ if the strategies are replaced according to } \phi_{\mathsf{SF}}} - \sum_{t=1}^{T} \underbrace{\mathbb{E}\left[v_i(s_i^t(\theta_i^t), a_{-i}^t)\right]}_{\text{reward in round } t}$$

※ Each player chooses $\sigma_i^t \in \Delta(S_i)$ and generates $s_i^t \sim \sigma_i^t$

Table of Contents

Agent-normal-form correlated equilibria

ANFCE is defined as CE of the agent normal form

Agent normal form of Bayesian games

The same player with different types are regarded as different players

Only (hypothetical) players with realized types play the game

In our example, randomly selected two out of (👮‍♀️, 🌊), (👮‍♀️, 🌋), (🧑‍🦰, 🌊), (🧑‍🦰, 🌋) play the game

Difference from SFCE

Each player cannot observe the recommendation to unrealized types

※ No realistic scenario involving a mediator 🤵‍♀️

ANFCE & Type-wise swap regret

Definition

A distribution $\sigma \in \Delta(S_1 \times \cdots \times S_n)$ is an ANFCE

$\triangleq$ For any player $i \in N$, $\phi\colon \Theta_i \times A_i \to A_i$,

$$\mathop{\mathbb{E}}_{\theta\sim\rho}\left[\mathop{\mathbb{E}}_{s\sim\sigma}[v_i(\theta; s_1(\theta_1), \ldots, s_n(\theta_n))]\right] \geq \mathop{\mathbb{E}}_{\theta\sim\rho}\left[\mathop{\mathbb{E}}_{s\sim\sigma}[v_i(\theta; \phi(\theta_i, s_i(\theta_i)), s_{-i}(\theta_{-i}))]\right].$$

$$R^T_{\mathrm{TS},i} \triangleq \max_{\phi\colon \Theta_i\times A_i\to A_i} \sum_{t=1}^{T} \underbrace{\mathbb{E}\left[v_i^t(\phi(\theta_i, s_i^t(\theta_i)), a_{-i}^t)\right]}_{\text{reward in round } t \text{ if the actions are replaced according to } \phi} - \sum_{t=1}^{T} \underbrace{\mathbb{E}\left[v_i^t(s_i^t(\theta_i), a_{-i}^t)\right]}_{\text{reward in round } t}$$

ANFCE & Type-wise swap regret

Definition

Choosing strategy $\phi(\theta_i, s_i(\theta_i))$ instead of recommended $s_i(\theta_i)$

A distribution $\sigma \in \Delta(S_1 \times \cdots \times$

$\triangleq$ For any player $i \in N$, $\phi\colon \Theta_i \times A_i \to A_i$,

$$\mathop{\mathbb{E}}_{\theta\sim\rho}\left[\mathop{\mathbb{E}}_{s\sim\sigma}\left[v_i(\theta; s_1(\theta_1), \ldots, s_n(\theta_n))\right]\right] \geq \mathop{\mathbb{E}}_{\theta\sim\rho}\left[\mathop{\mathbb{E}}_{s\sim\sigma}\left[v_i(\theta; \phi(\theta_i, s_i(\theta_i)), s_{-i}(\theta_{-i}))\right]\right].$$

$$R^T_{\mathrm{TS},i} \triangleq \max_{\phi\colon \Theta_i\times A_i\to A_i} \underbrace{\sum_{t=1}^{T} \mathbb{E}\left[v_i^t(\phi(\theta_i, s_i^t(\theta_i)), a_{-i}^t)\right]}_{\text{reward in round } t \text{ if the actions are replaced according to } \phi} - \underbrace{\sum_{t=1}^{T} \mathbb{E}\left[v_i^t(s_i^t(\theta_i), a_{-i}^t)\right]}_{\text{reward in round } t}$$

Table of Contents

Bayesian solutions [Forges'93]

Mediator knows the true types in advance

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

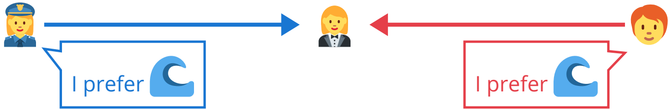

2. The mediator privately sends a recommendation to each player

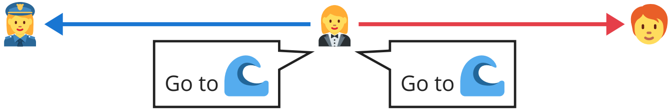

Bayesian solutions [Forges'93]

Mediator knows the true types in advance

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

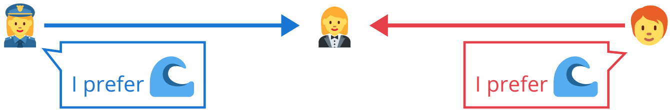

2. The mediator privately sends a recommendation to each player

← **No incentive to disobey the recommendation**

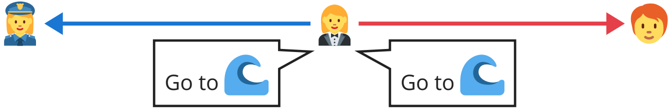

Bayesian solutions

Definition

A distribution $\pi \in \Delta(A)^\Theta$ is a Bayesian solution

$\overset{\triangle}{\Leftrightarrow}$ For any player $i \in N$, $\phi\colon \Theta_i \times A_i \to A_i$,

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

Difference from ANFCE

$\pi \in \Delta(A)^\Theta$ can express broader distributions than $\sigma \in \Delta(S)$,

which we call **strategy representability** (e.g., π in the previous page)

Table of Contents

Communication equilibria [Myerson'82, Forges'86]

Equilibria realized by with bidirectional communication

1 Each player privately tells their types to the mediator

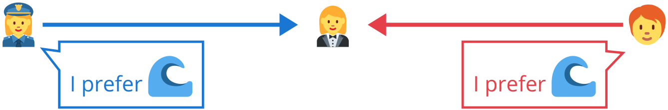

2 The mediator privately sends a recommendation to each player

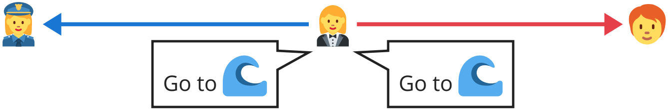

Communication equilibria [Myerson'82, Forges'86]

Equilibria realized by with bidirectional communication

1. Each player privately tells their types to the mediator

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

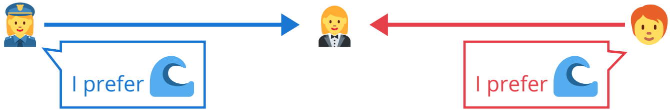

2. The mediator privately sends a recommendation to each player

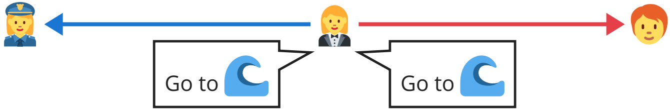

Communication equilibria [Myerson'82, Forges'86]

Equilibria realized by with bidirectional communication

1. Each player privately tells their types to the mediator

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

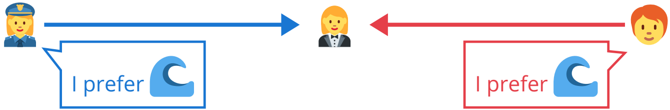

2. The mediator privately sends a recommendation to each player

   ← **No incentive to disobey the recommendation**

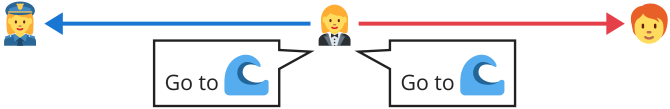

Communication equilibria [Myerson'82, Forges'86]

Definition

A distribution $\pi \in \Delta(A)^{\Theta}$ is a communication equilibrium

$\overset{\triangle}{\Leftrightarrow}$ For any player $i \in N$, $\psi \colon \Theta_i \to \Theta_i$, and $\phi \colon \Theta_i \times A_i \to A_i$,

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

1. Each player $i \in N$ privately tells θ_i (possibly $\psi(\theta_i)$) to 🤵‍♀️
2. 🤵‍♀️ privately sends recommendations $a \sim \pi(\theta)$ to each player
3. Each player i chooses their action a_i (possibly deviates to $\phi(\theta_i, a_i)$)

Communication equilibria [Myerson'82, Forges'86]

Definition

A distribution $\pi \in \Delta(A)^{\Theta}$ i

Misreporting $\psi(\theta_i)$ instead of true type θ_i

Choosing action $\phi(\theta_i, a_i)$ instead of recommended a_i

$\overset{\Delta}{\Leftrightarrow}$ For any player $i \in N$, $\psi\colon \Theta_i \to \Theta_i$, and $\phi\colon \Theta_i \times A_i \to A_i$,

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

1. Each player $i \in N$ privately tells θ_i (possibly $\psi(\theta_i)$) to 🤵‍♀️
2. 🤵‍♀️ privately sends recommendations $a \sim \pi(\theta)$ to each player
3. Each player i chooses their action a_i (possibly deviates to $\phi(\theta_i, a_i)$)

Communication equilibrium combines ...

Mechanism design

1. Each player tells their types
   ← **No incentive to lie**

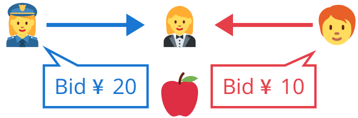

2. 🤵‍♀️ decides the outcome
   ← This decision is binding

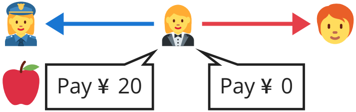

Correlated equilibria

1. No type (complete info.)

2. 🤵‍♀️ recommends actions
   ← **No incentive to deviate**

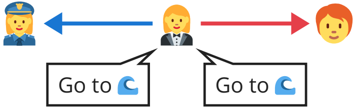

Untruthful swap regret [Fujii'25a]

Untruthful swap regret for player $i \in N$

$$R^T_{\mathsf{US},i} = \max_{\substack{\psi\colon \Theta_i \to \Theta_i \\ \phi\colon \Theta_i \times A_i \to A_i}} \sum_{t=1}^{T} \mathbb{E}_{\theta \sim \rho} \left[\mathbb{E}_{\substack{a_i \sim \pi_i^t(\psi(\theta_i)), \\ a_{-i} \sim \pi_{-i}^t(\theta_{-i})}} [v_i(\theta; \phi(\theta_i, a_i), a_{-i})] \right] - \sum_{t=1}^{T} \mathbb{E}_{\theta \sim \rho} \left[\mathbb{E}_{\substack{a_i \sim \pi_i^t(\theta_i), \\ a_{-i} \sim \pi_{-i}^t(\theta_{-i})}} [v_i(\theta; a_i, a_{-i})] \right]$$

- Dynamics minimizing this regret converge to **communication equilibria with strategy representability**
- An efficient learning algorithm and lower bound

Relations among BCE concepts

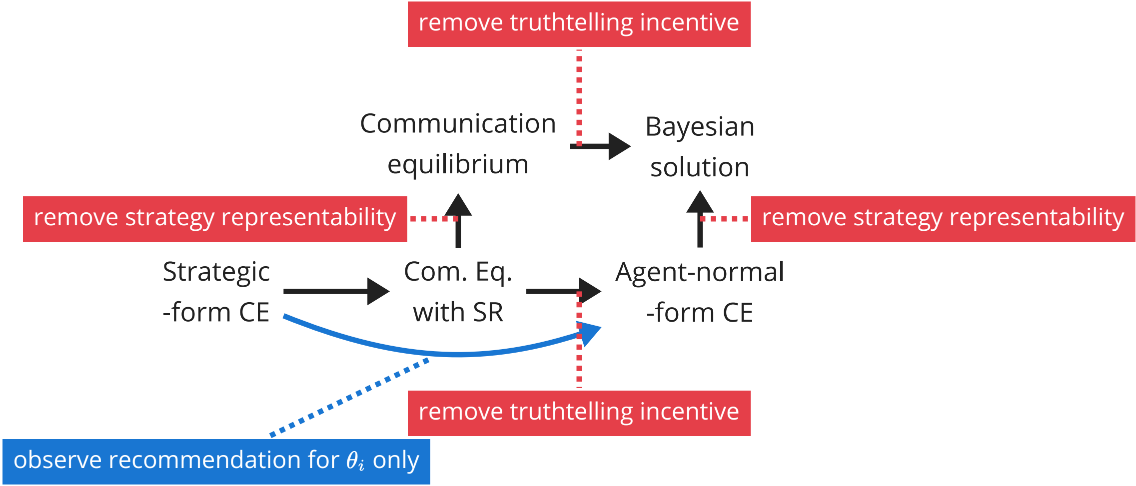

Table of Contents

Example of valid utility games

Players simultaneously choose a resource to share

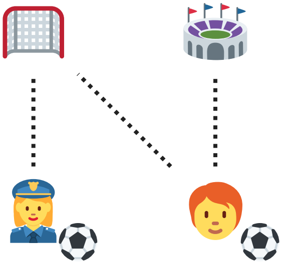

Resources chosen by multiple players are partitioned in a prespecified way

Example:  is prioritized over 

Example of valid utility games

Players simultaneously choose a resource to share

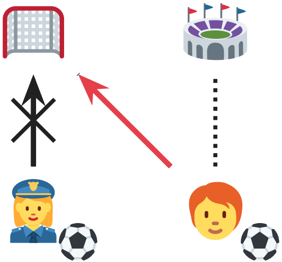

Resources chosen by multiple players are partitioned in a prespecified way

Example: 🧑 is prioritized over 👮‍♀️

No player can benefit from deviations

◁⇕

Worst **Nash equilibrium** $= 1$

Example of valid utility games

Players simultaneously choose a resource to share

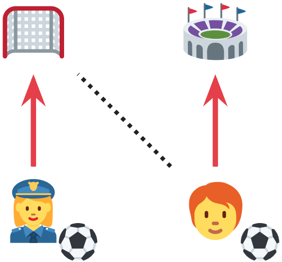

Resources chosen by multiple players are partitioned in a prespecified way

Example: 🧑 is prioritized over 👮‍♀️

Optimal social welfare $= 2$

Example of valid utility games

Players simultaneously choose a resource to share

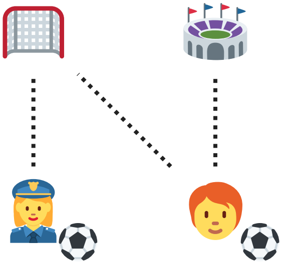

Resources chosen by multiple players are partitioned in a prespecified way

Example:  is prioritized over 

No player can benefit from deviations

$\lhd\Updownarrow$

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

(price of anarchy)

Example of valid utility games

Players simultaneously choose a resource to share

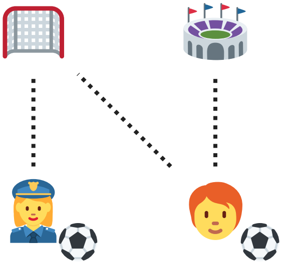

Resources chosen by multiple players are partitioned in a prespecified way

Example: is prioritized over

No player can benefit from deviations

$\Updownarrow$

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

(price of anarchy)

Theorem [Vetta'02]

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**

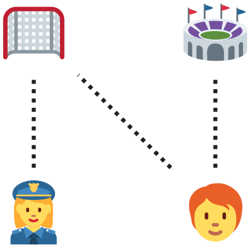

A mediator 👩‍💼 sends recommendations

(👩‍💼 realizes correlated equilibrium)

Example of valid utility games

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

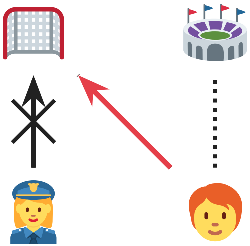

A mediator sends recommendations

(realizes correlated equilibrium)

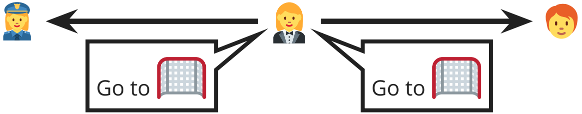

Example of valid utility games

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

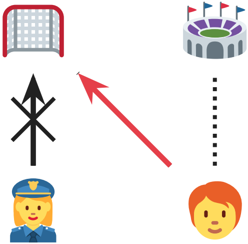

A mediator sends recommendations

(realizes correlated equilibrium)

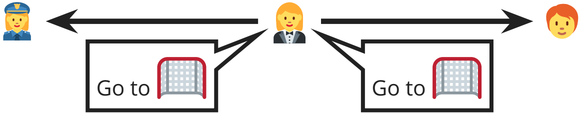

Theorem [Roughgarden'15a]

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

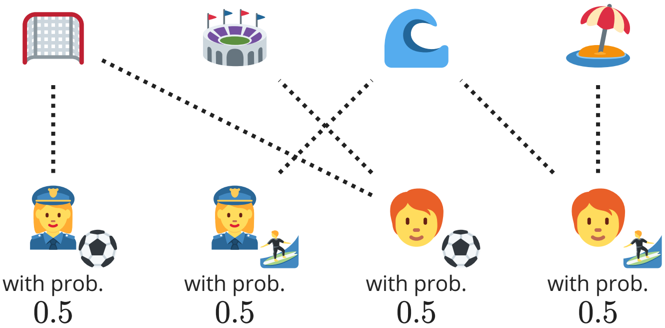

Q **How do mediators work in Bayesian games?**

Example of Bayesian valid utility games

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

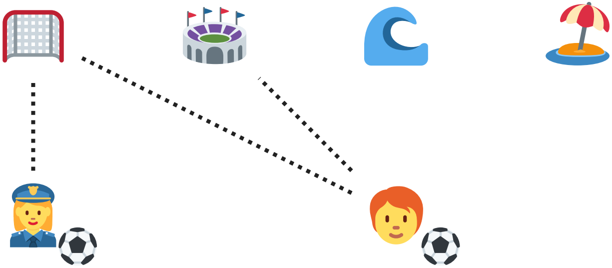

with probability $1/4$

👮‍♀️ 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

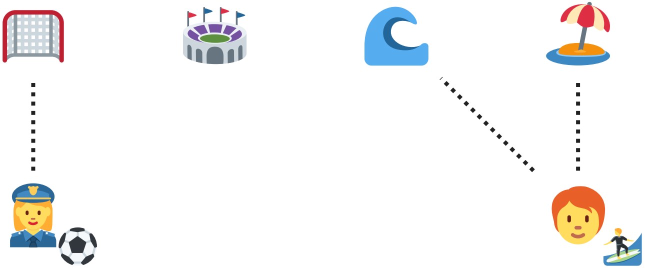

with probability $1/4$

👮‍♀️ does not know 🧑‍🦰's type

🧑‍🦰 does not know 👮‍♀️'s type

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

Our results

For various equilibrium concepts, we provide PoA and PoS bounds

PoA bounds for independent priors

PoA $\in [0.316, 0.441]$

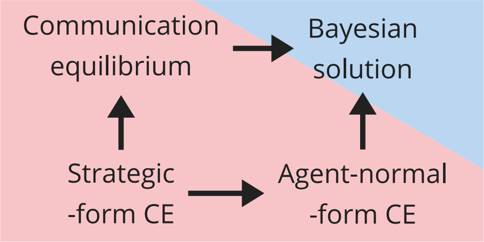

PoA $= 0.5$

PoS bounds for independent priors

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

PoS $= 1$

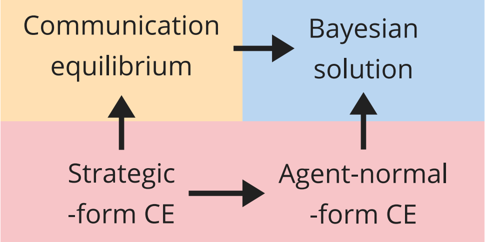

PoS $= 1 - 1/e$
under the basic utility assumption

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

Table of Contents

Notations for Bayesian games

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

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

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

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

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

$v_i : A \to \mathbb{R}_{\geq 0}$ 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

Equivalence of two formulations

Original formulation

A_i finite set of actions for player $i \in N$

$v_i : \Theta \times A \to \mathbb{R}$ utility function for player $i \in N$

(θ_i, a_i) as an action ↓ ↑ $A_i \coloneqq \bigcup_{\theta_i} A_i^{\theta_i}$ and ignore actions for $\forall \theta_i' \neq \theta_i$

Type-dependent-action formulation

$A_i^{\theta_i}$ finite set of actions for player $i \in N$ with type $\theta_i \in \Theta_i$

$v_i : A \to \mathbb{R}_{\geq 0}$ 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

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

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'02]

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

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

Submodular social welfare functions

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

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

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'02]

- $\sum_{i \in N} v_i(a) \leq f(\{a_1, \ldots, a_n\})$ 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

$v_i \colon A \to \mathbb{R}_{\geq 0}$ 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'02]

- $\sum_{i \in N} v_i(a) \leq f(\{a_1, \ldots, a_n\})$ 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

Price of anarchy (PoA) in Bayesian games

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

$$\mathsf{PoA}_{\Pi} \triangleq \frac{\underbrace{\inf_{\pi \in \Pi} \mathop{\mathbb{E}}_{\theta \sim \rho} \left[\mathop{\mathbb{E}}_{a \sim \pi(\theta)} [v_{\mathsf{SW}}(a)] \right]}_{\text{the social welfare achieved by the worst equilibrium}}}{\underbrace{\mathop{\mathbb{E}}_{\theta \sim \rho} \left[\max_{a^* \in A^{\theta}} v_{\mathsf{SW}}(a^*) \right]}_{\text{the optimal social welfare}}},$$

where $v_{\mathsf{SW}}(a) \triangleq f(\{a_1, \ldots, a_n\})$

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

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

the optimal social welfare

Strategy-representability gap (SR gap)

the social welfare achieved
by the worst equilibrium

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

the social welfare achieved
by the optimal strategy profile

the optimal social welfare

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

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 $s(\theta) \triangleq (s_1(\theta_1), \ldots, s_n(\theta_n))$

the social welfare achieved by the worst equilibrium

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

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)

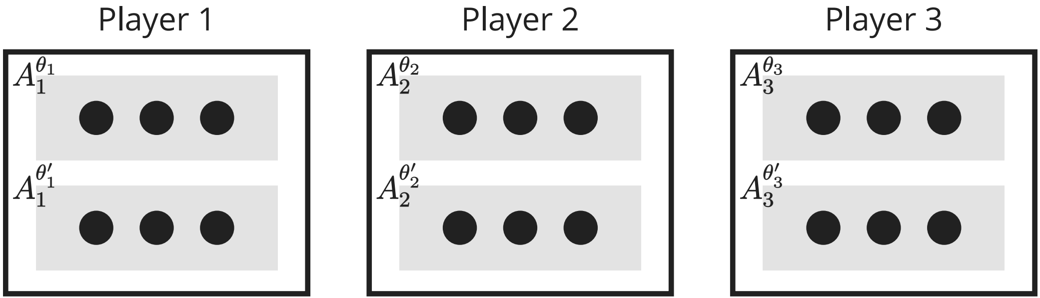

For each player, one block is chosen according to a known distribution

Strategy-representability gap (SR gap)

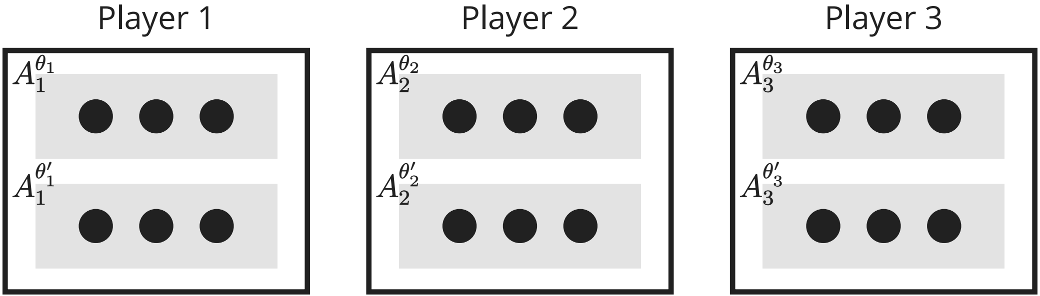

For each player, one block is chosen according to a known distribution

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

Strategy-representability gap (SR gap)

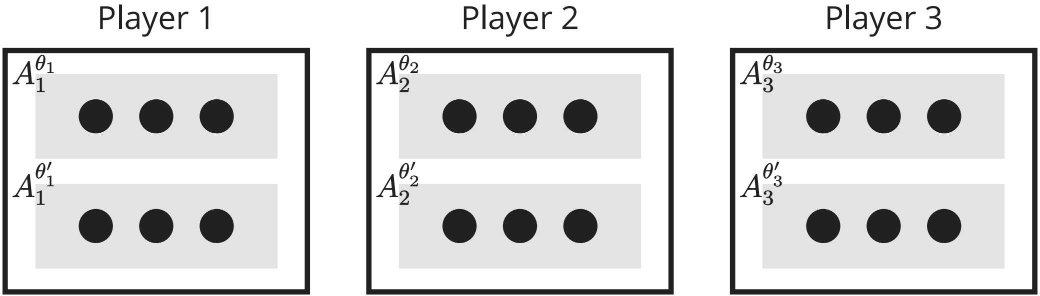

For each player, one block is chosen according to a known distribution

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

Strategy-representability gap (SR gap)

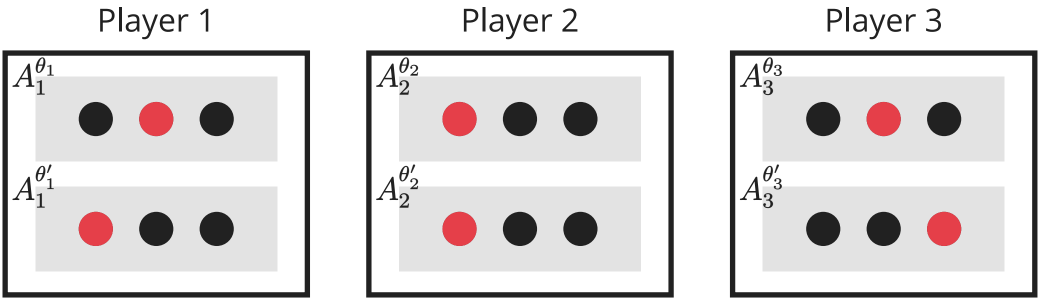

For each player, one block is chosen according to a known distribution

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

Strategy-representability gap (SR gap)

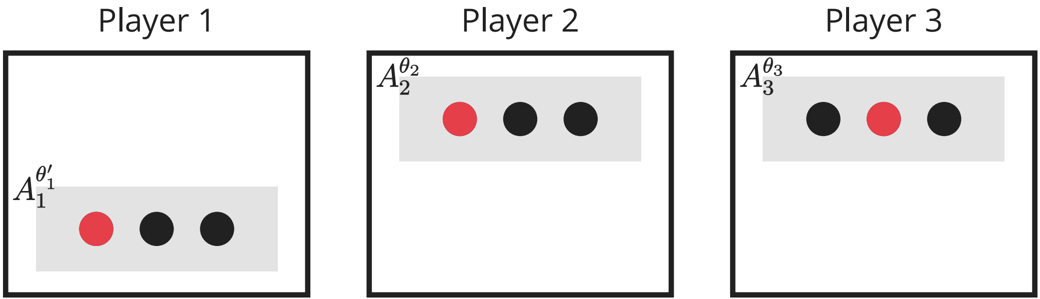

For each player, one block is chosen according to a known distribution

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

Strategy-representability gap (SR gap)

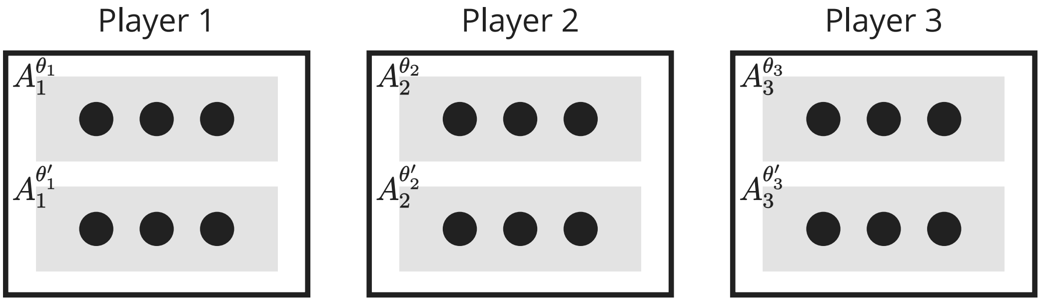

For each player, one block is chosen according to a known distribution

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

Strategy-representability gap (SR gap)

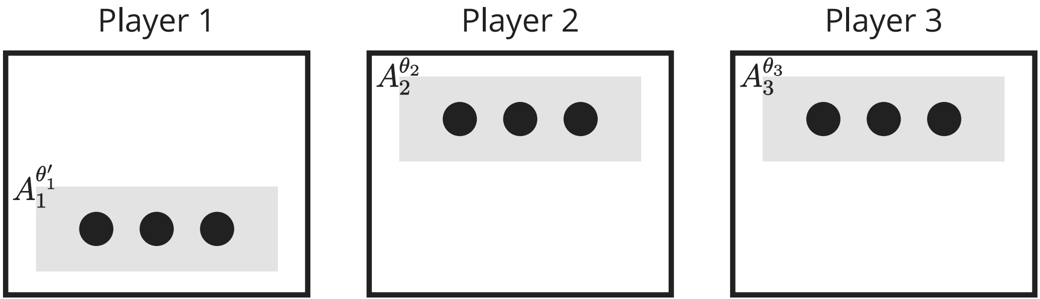

For each player, one block is chosen according to a known distribution

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

Strategy-representability gap (SR gap)

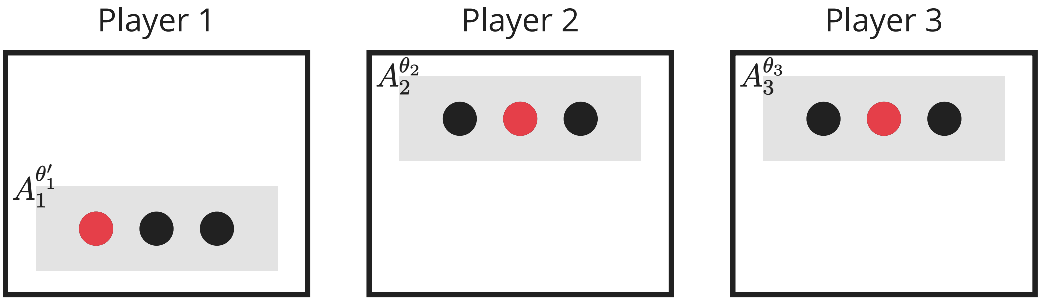

For each player, one block is chosen according to a known distribution

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

SR gap analyses for two prior assumptions

Q What is the worst-case value of the SR gap?

$$\mathrm{SRgap} = \frac{\max_{s^* \in S} \mathbb{E}_{\theta \sim \rho}\left[v_{\mathsf{SW}}(s^*(\theta))\right]}{\mathbb{E}_{\theta \sim \rho}\left[\max_{a \in A^\theta} v_{\mathsf{SW}}(a)\right]}$$

- 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

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

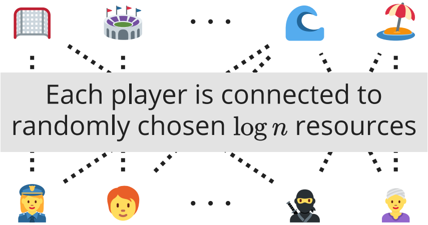

Optimal social welfare: n

$\because$ 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

$\mathrm{SRgap} = \Omega(1/\sqrt{n})$, 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

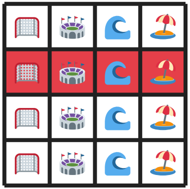

1st action

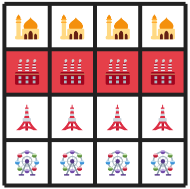

2nd action

$E = [k] \times [n]$ 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

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'15b, 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

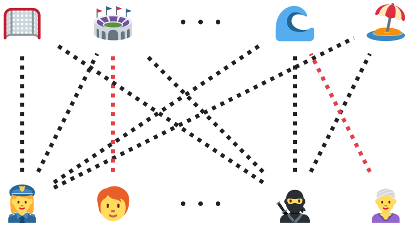

Odd players are connected to all resources
Even players are connected to random one
Odd players are prioritized over even ones

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

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

Our results

For various equilibrium concepts, we provide PoA and PoS bounds

PoA bounds for independent priors

PoA $\in [0.316, 0.441]$

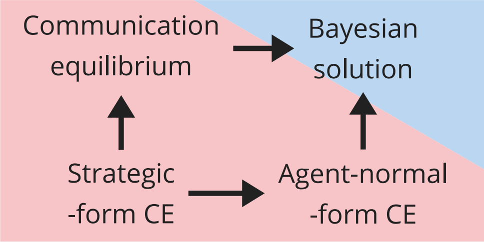

PoA $= 0.5$

PoS bounds for independent priors

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

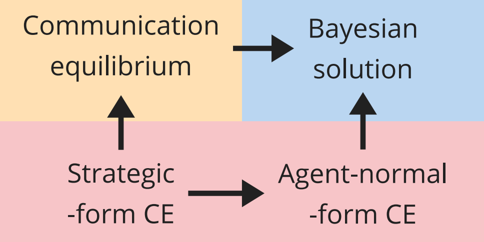

PoS $= 1 - 1/e$
under the basic utility assumption

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

Reference (1/2)

- Dirk Bergemann, & Stephen Morris. 2016. Information design, Bayesian persuasion, and Bayes correlated equilibrium. *American Economic Review*, 106(5), 586–591.
- Avrim Blum and Yishay Mansour. 2007. From External to Internal Regret. *Journal of Machine Learning Research* 8, 1307–1324.
- 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.
- Kaito Fujii. 2025a. Bayes correlated equilibria, no-regret dynamics in Bayesian games, and the price of anarchy. Proceedings of Machine Learning Research vol, 291, 1-2.
- Kaito Fujii. 2025b. The Power of Mediators: Price of Anarchy and Stability in Bayesian Games with Submodular Social Welfare. In Proceedings of the 26th ACM Conference on Economics and Computation, pp. 251-251.
- Dean P Foster and Rakesh V Vohra. 1997. Calibrated learning and correlated equilibrium. *Games and Economic Behavior* 21(1-2), 40–55.

Reference (2/2)

- 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.
- Sergiu Hart and Andreu Mas-Colell. 2000. A simple adaptive procedure leading to correlated equilibrium. *Econometrica* 68(5), 1127–1150.
- Roger B. Myerson. 1982. Optimal coordination mechanisms in generalized principal – agent problems. *Journal of Mathematical Economics*, 10(1), 67–81.
- 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.