Competing Mechanisms: Revelation Principle and Extensions

Guillaume Pommey
University of Rome Tor Vergata

April 2022

Environment

Multi-principal, (multi)-agent settings.

Each principal can **fully design a mechanism**.

An agent has **private information** and can **communicate with principals.**

I will define this formally later on.

Challenges and goals

Private information of an agent **is not just a "type".**

An agent also have **market information**.

- The principals' mechanisms.

The usual **revelation principle** does not apply.

- We will see an example.

Our goal is to find a way to **restore/replace** the revelation principle and to **characterize** equilibrium allocations.

A simple principal-agent problem

First, let me state the revelation principle in a basic framework.

One principal, one agent.

The agent has **private information.**

- Let $\theta \in \Theta$ denote the agent's type.
- The principal only knows that $\theta \sim F(\theta)$.

The principal must take **an action** $x \in X$.

Utilities are

- Principal. $v : X \times \Theta \to \mathbb{R}$
- Agent. $u : X \times \Theta \to \mathbb{R}$

Communication and messages

The **agent's type** matters to the principal as $v(x, \theta)$.

We assume that the agent and the principal **can communicate.**

Formally, we say that the agent can send a **message** $m \in M$ to the principal.

We do not make **any assumption** of the set of messages M.

- Possibly a very complex object.

Mechanisms: Definition

Before communication takes place, the principal can **commit** to a **mechanism.**

A mechanism associates to each message an action of the principal.

Formally, **a mechanism is a function** $\gamma : M \to X$.

- Let $\Gamma(M) := X^M$ denote the set of all mechanisms from M to X.

Timing.

1. The principal commits to a γ.
2. The agent observes γ and sends a message $m \in M$.
3. The mechanism implements the action $\gamma(m) \in X$.

Generality

Notice that X can be very general.

It could be a "simple" set like $X = \{\text{Hire}, \text{Do not hire}\}$.

Or something much more complex like the set of all possible pairs of quantity-price $(q, p) \in \mathbb{R}_+^2$.

And we could define for instance:

- $v(x, \theta) = v(q, p, \theta) = p - c(q)$.
- $u(x, \theta) = u(q, p, \theta) = \theta q - p$.

The agent's problem

For now, **fix the set of messages** M.

When the agent of type θ faces γ, they **communicate a message** such that

$$m(\theta) \in \arg\max_{m \in M} u(\gamma(m), \theta),$$

i.e. the agent simply reports a message implementing the action $x \in X$ that maximize their payoff.

The principal's problem

The principal must find the best $\gamma \in \Gamma(M)$ to maximize their expected payoff.

The principal expects type θ to report $m(\theta)$.

- Hence that type θ will implement $\gamma(m(\theta))$.

The expected payoff of the principal under γ is therefore

$$V(\gamma) := \int_{\Theta} v(\gamma(m(\theta)), \theta) dF(\theta).$$

Challenges

Finding the optimal γ for the principal requires to also manipulate the agent's reporting strategy $m(\cdot)$.

- Not (very) tractable given that γ is already a complex object.

Much worse, there is **no reason to fix** M.

- Maybe very rich/empty communication spaces are good for the principal.

Hence, the principal should optimize over M and $\gamma \in \Gamma(M)$.

Direct mechanisms

Let $M = \Theta$ and $\gamma^D : \Theta \to X$.

Mechanisms γ^D are called **direct mechanisms.**

- The agent reports directly a type $\theta \in \Theta$.

Also assume that γ^D is such that for type θ

$$\theta \in \arg\max_{\theta' \in \Theta} u(\gamma(\theta'), \theta),$$

that is, type θ's reporting strategy is to report θ.

We say that γ^D is **truthful.**

The revelation principle

Allocation rule: Given M and γ, let $x^R(\theta) := \gamma(m(\theta))$

- $x^R : \Theta \to X$ is a mapping from types into actions.

Theorem. Any allocation rule x^R obtained with a mechanism (M, γ) can also be implemented with a truthful direct mechanism γ^D (and so $M = \Theta$).

In other words, **there is no loss of generality** in restricting attention to truthful direct mechanisms.

The revelation principle: Proof

By definition, $x^R(\theta) = \gamma(m(\theta))$ is a mapping from Θ to X.

Hence $\gamma(m(\cdot)) = \gamma^D$ **is a direct mechanism.**

$$
\begin{tikzcd}
\Theta \arrow[r, "m(\cdot)"] \arrow[rr, bend right, "\gamma^D = \gamma \circ m"'] & M \arrow[r, "\gamma(\cdot)"] & x^R(\theta) = \gamma(m(\theta))
\end{tikzcd}
$$

It must also be shown that γ^D is truthful.

The revelation principle: Intuition

It means that the *language* Θ is sufficient to represent all possible (indirect) mechanisms.

Communication has no value in itself for the agent, it is just a mean to ask for a particular x to be implemented.

The revelation principle easily extends to **multi-agent environments.**

"Failure" of the RP

Consider now an environment with **two principals, and one agent.**

Define X^i, v^i, M^i and γ^i naturally for each principal $i = 1, 2$.

Assume that $|\Theta| = 1$ and $X^i = \{A, B, C\}$ for $i = 1, 2$.

If we consider direct mechanisms, then $\gamma^i \in X^i$ (or possibly $\in \Delta(X^i)$).

- The agent has no role.
- Each principal directly chooses an action in X^i.

"Failure" of the RP

Assume that players' payoffs $\{v^1, v^2, u\}$ are as follows.

$1 \setminus 2$	A	B	C
A	$1, 1, 1$	$2, 0, 2$	$-1, 5, 10$
B	$0, 2, 2$	$1, 1, 1$	$0, 0, 0$
C	$5, -1, 10$	$0, 0, 0$	$0, 0, 0$

(Example from Martimort and Stole, 2002).

"Failure" of the RP

1 \ 2	A	B	C
A	$1, 1, 1$	$2, 0, 2$	$-1, 5, 10$
B	$0, 2, 2$	$1, 1, 1$	$0, 0, 0$
C	$5, -1, 10$	$0, 0, 0$	$0, 0, 0$

(C, C) can be implemented through direct mechanisms $\gamma^i = C$.

(B, B) cannot be implemented through direct mechanisms $\gamma^i = B$.

- If $\gamma^i = B$, then j wants to deviate to $\gamma^j = A$.

Indirect mechanisms

Consider now $M^i = \{m^i_B, m^i_C\}$.

Let $\gamma^i(m^i_B) = B$ and $\gamma^i(m^i_C) = C$ for $i = 1, 2$.

- The agent *recommends* an action to the principal.

Then, the outcome (B, B) can be implemented.

Indirect mechanisms

$1 \setminus 2$	A	B	C
A	$1, 1, 1$	$2, 0, 2$	$-1, 5, 10$
B	$0, 2, 2$	$1, 1, 1$	$0, 0, 0$
C	$5, -1, 10$	$0, 0, 0$	$0, 0, 0$

For the agent, when the two principals offer contracts including only $\{B, C\}$, it is clear that recommending B to both is optimal.

But why offering also C in the contracts?

Indirect mechanisms

1 \ 2	A	B	C
A	$1, 1, 1$	$2, 0, 2$	$-1, 5, 10$
B	$0, 2, 2$	$1, 1, 1$	$0, 0, 0$
C	$5, -1, 10$	$0, 0, 0$	$0, 0, 0$

Assume 2 offers only $\{B\}$, then 1 best-response is to offer $\{A\}$.

But if 2 offers $\{B, C\}$, then 1 does not want to offer $\{A\}$ anymore.

- If the agent faces $\{A\}$ and $\{B, C\}$.
- They will make 2 play C and get their maximal payoff (10).
- 1 will get their lowest payoff (-1).

A new role for the agent

(B, B) can be implemented with indirect mechanisms but not with direct mechanisms.

Here the agent has no type to report as $|\Theta| = 1$.

But the agent has **market information**.

- The mechanisms offered by each principal.

(B, B) can be implemented because the agent is used as a way to **coordinate principals.**

- The presence of C in the equilibrium mechanisms $\{B, C\}$ is a credible threat toward a deviation $\{A\}$.

Issues and remedies

In multi-principal settings, some equilibrium allocations **cannot be implemented with direct mechanisms.**

We come back to the initial tractability issues.

We will investigate two remedies.

- Restore the revelation principle by **incorporating market information into the agent's "type"** (Epstein and Peters, 1999).
- Replace the revelation principle by the **delegation principle** (Martimort and Stole, 2002).

Delegation principle: Martimort and Stole (2002)

Martimort and Stole (2002) approach the problem by relying on what they call the **delegation principle.**

The idea is to restrict attention to **decentralized menus.**

- A principal directly offers an agent a **menu of actions**.
- From which the agent directly **picks an action.**

The environment

N principals, one agent.

- X^i, v^i, and M^i defined as before.

The agent's type space Θ is finite.

Define $\gamma^i : M^i \to \Delta(X^i)$ and $\Gamma^i := (\Delta(X^i))^{M^i}$.

- $\gamma^i(x^i \mid m^i)$: Prob of x_i given m^i.
- Let $\gamma := (\gamma^1, \ldots, \gamma^N)$, $\Gamma := \times_{i=1}^N \Gamma^i$.

Timing.

1. Each principal chooses M^i and γ^i.
2. The agent reports $m = (m^1, \ldots, m^N) \in M$.
3. Actions are implemented according to m and the γ^i's.

Strategies

A principal's strategy is $\sigma_i(\gamma^i) \in \Delta(\Gamma^i)$.

The agent's strategy is $\sigma_0 : \Theta \times \Gamma \to \Delta(M)$.

- Where $\sigma_0(m \mid \theta, \gamma)$ is prob of m given θ and γ.

Let $\sigma := (\sigma_0, \sigma_1, \ldots, \sigma_N)$.

Finally define supp $\sigma_0(\theta, \gamma)$ the support of σ_0.

- i.e. the set of messages such that $\sigma_0(m \mid \theta, \gamma) > 0$.
- supp σ_i is defined similarly.

Equilibrium

A strategy profile σ^* is a PBE of $\Gamma(M)$ iff

(a) For all $\theta \in \Theta$ and all $\gamma \in \Gamma(M)$,

$$m \in \text{supp } \sigma_0^*(\theta, \gamma) \Rightarrow m \in \arg\max_{\hat{m} \in M} \sum_{x \in X} u(x, \theta) \Pi_{i=1}^N \gamma^i(x^i \mid \hat{m}^i).$$

(b) For all i, $\gamma^i \in \text{supp } \sigma_i^* \Rightarrow$

$$\gamma^i \in \arg\max_{\hat{\gamma}^i \in \Gamma^i} \int_{\gamma^{-i} \in \Gamma^{-i}} \int_{m \in M} \sum_{\theta \in \Theta} \sum_{x \in X} v^i(x, \theta) f(\theta) \times d\sigma_0^*(m \mid \theta, \hat{\gamma}^i, \gamma^{-i}) \hat{\gamma}^i(x_i \mid m^i) \Pi_{j \neq i} d\sigma_j^*(\gamma^j).$$

Allocation and direct mechanisms

We can denote by $\mu_\sigma(x \mid \theta)$ the **allocation induced by the strategy profile** σ for each $\theta \in \Theta$.

A **direct mechanism** would be a γ with $M^i = \Theta$ and $M = \Theta^N$.

If it were true, the revelation principle would state that for any $\sigma^* \in \mathrm{PBE}(\Gamma(M))$ there exists a $\tilde{\sigma}^* \in \mathrm{PBE}(\Gamma(\Theta^N))$ such that

(a) $\mu_{\sigma^*}(x \mid \theta) = \mu_{\tilde{\sigma}^*}(x \mid \theta)$, for all x, θ.

(b) $\tilde{\sigma}_0^*(\theta \mid \theta, \hat{\gamma}) = 1$ for all θ and γ^i such that $\tilde{\sigma}_i^*(\gamma^i) > 0$.

But we know from the example that **this does not hold.**

The taxation principle

The **taxation principle** is useful to get an idea of the **delegation principle.**

One principal, one agent

- Assume mechanisms are DRM $\{p(\theta), q(\theta)\}_{\theta \in \Theta}$.

The taxation principle says that we can replace these DRM by a nonlinear tariffs.

- $P(q) := p(\theta^{-1}(q))$.

The agent simply picks a quantity q and pays the price $P(q)$.

Distribution induced by mechanisms

Take a given (γ^i, M^i).

Let $\gamma^i(\cdot \mid M^i) \subseteq \Delta(X^i)$ be the image of this mapping.

Notice that if $M^i = \{m^i\}$, $\gamma^i(\cdot \mid M^i)$ is a single prob. dist. over X^i.

If $M^i = \{m^i_1, m^i_2\}$, $\gamma^i(\cdot \mid M^i)$ is at most two prob. dist. over X^i.

Hence $\gamma^i(\cdot \mid M^i)$ describes the possible probability distributions over X^i induced by (γ^i, M^i).

Menus of distributions (1)

We say that an arbitrary menu of distributions, T^i, is **consistent with message space** M^i if there exists a γ^i defined on M^i such that

$$T^i = \gamma^i(\cdot \mid M^i).$$

And let $\mathcal{T}^i(M^i)$ denote the set of all menus consistent with M^i,

$$\mathcal{T}^i := \{\gamma^i(\cdot \mid M^i) \mid \gamma^i \in \Gamma^i(M^i)\}.$$

Menus of distributions (2)

Notice that the cardinality of T^i is at most that of M^i.

- As (γ_i, M^i) can generate at most $|M^i|$ distributions over X^i.

For each (γ_i, M^i), $\gamma^i(\cdot \mid M^i)$ is uniquely defined.

- But $\gamma^i(\cdot \mid M^i)$ can be generated by many (γ^i, M^i).

Hence, we can define an **equivalence class** of mechanisms.

- (γ^i, M^i) and $(\gamma^{i'}, M^{i'})$ belong to the same class if

$$\gamma^i(\cdot \mid M^i) = \gamma^{i'}(\cdot \mid M^{i'}).$$

Delegation principle (1)

The idea of the delegation principle is that

- Each principal offers a T^i.
- The agent picks a $t \in T := \times_{i=1}^{N} T^i$.

Hence in the delegation game

- Each principal chooses a $\tilde{\sigma}_i(T^i) \in \Delta(\mathcal{T}^i(M^i))$.
- The agent chooses a $\tilde{\sigma}_0(\tau \mid \theta, T)$.

Delegation principle (2)

Theorem. For each $\sigma^* \in \text{PBE}(\Gamma(M))$, $\exists \tilde{\sigma}^* \in \text{PBE}(\mathcal{T}(M))$ such that $\mu_{\sigma^*}(x \mid \theta) = \mu_{\tilde{\sigma}^*}(x \mid \theta)$ for all x and θ.

It means that the specification of (γ^i, M^i) matters only in the distributions over actions it generates.

- Directly offering a choice of distributions to the agent is without loss of generality.

Improves tractability

The delegation principle is useful because $T^i \subseteq \Delta(X^i)$.

Hence, $\mathcal{T}^i(M^i)$ cannot be greater than the cardinality of the power set of $\Delta(X^i)$.

It means that there is a limit to the complexity of the mechanisms offered by principals.

Universal mechanisms: Epstein and Peters (1999)

To restore the revelation principle, Epstein and Peters (1999) establish a **universal type space** that includes

- The *usual* agent's type, $\theta \in \Theta$.
- The **market information** about other principal's mechanisms.

The main challenge consists in establishing the existence of such a language.

- Each principal must include messages to describe the other principal's mechanisms.
- Potential **infinite regress** problem.
- Epstein and Peters (1999) show that there exists a well-defined universal type space.

The environment (1)

Two principals, two agents.

Let X_0 denote the set of "simple" actions.

- Similar to X before.

Let Θ denote the **valuation space.**

- With F the CDF for valuations.

This change of terminology is intentional.

The environment (2)

Each agent can **participate with at most one principal.**

- Let $P := \{0, 1\}$ denote participation decisions.

Principals can condition their choice of simple action $x_0 \in X_0$ to **participation decisions.**

- That is, each principal chooses in the set $(X_0)^{P^2}$.
- Let x_c denote a typical element of this set.

Define π_1 and π_2 the probability that agent 1 and 2 participates in a given principal mechanism.

The environment (3)

Now we define the **action space** by:

$$X := (X_0)^{P^2} \times [0,1]^2.$$

An action is a triple $x = (x_c, \pi_1, \pi_2)$.

Payoffs are defined as follows:

$$v : X \times \Theta^2 \to [0,1], \qquad \text{(Principals)}$$
$$u : X \times \Theta \to [0,1], \qquad \text{(Agents)}$$

where u represents the payoff of an agent participating to a given principal's mechanism.

The environment (4)

Participation is exclusive in this setting.

Hence, it does not correspond to our example of the failure of the revelation principle.

But the same kind of contractual externality among principals exists because of the participation decisions of agents.

Indirect mechanisms

Let $\gamma \in \Gamma$ denote a mechanism and M be the message space.

Mechanisms are defined as

$$\gamma : M^2 \to X.$$

Let $\gamma = (\gamma_c, \gamma_{\pi_1}, \gamma_{\pi_2})$ where

$$\begin{aligned} \gamma_c &: M^2 \to (X_0)^{P^2}, \\ (\gamma_{\pi_1}, \gamma_{\pi_2}) &: M^2 \to [0,1]^2, \end{aligned}$$

For each pair of messages, a mechanism announces a contingent simple action γ_c and a participation recommendation γ_{π_i} to each agent.

Participation and communication

It is therefore assumed that agents communicate with **both principals** before choosing with whom to participate.

- Possible to restrict mechanisms such that communication is tied to participation.
- Like in McAfee (1993) and Peters (1997).

Communication and participation strategies

Communication strategy: Let $\tilde{c}$ be a measurable mapping

$$\tilde{c} : \Theta \times \Gamma^2 \to M,$$

where $\tilde{c}(\theta, \gamma, \gamma')$ represents the message sent by type θ to the principal offering γ given that the other principal offers γ'.

Participation strategy: Let $\tilde{\pi}$ be a measurable mapping

$$\tilde{\pi} : \Theta \times \Gamma^2 \to [0, 1],$$

where $\tilde{\pi}(\theta, \gamma, \gamma')$ is defined similarly.

Continuation equilibrium

We say that $(\tilde{c}, \tilde{\pi})$ is a **continuation equilibrium** if no agent has any incentive to deviate from the communication and participation strategies $\tilde{c}$ and $\tilde{\pi}$ for any $\theta \in \Theta$ and any pair $(\gamma, \gamma') \in \Gamma^2$.

It is key to understand that it **is not for a given pair of mechanisms**.

Implemented actions

It is possible to describe the action forthcoming at the principal offering mechanism γ by

$$m_\gamma(\theta,\theta',\gamma',\gamma'') := (\gamma_c(\tilde{c}(\theta,\gamma,\gamma'),\tilde{c}(\theta',\gamma,\gamma'')),\tilde{\pi}(\theta,\gamma,\gamma'),\tilde{\pi}(\theta',\gamma,\gamma'')),$$

for agents of type θ and θ'.

Where θ acts as if the other principal offers γ' and θ' as if they offer γ''.

Expected payoffs and equilibrium

Assume principals play strategies δ, $\delta' \in \Delta(\Gamma)$.

Then the expected payoff of the principal choosing δ is given by

$$V(\delta; \delta', \tilde{c}, \tilde{\pi}) := \int v(m_\gamma(\theta, \theta', \gamma', \gamma'), \theta, \theta') dF(\theta) dF(\theta') d\delta'(\gamma') d\delta(\gamma).$$

Hence $(\tilde{c}, \tilde{\pi}, \delta^*)$ is an equilibrium relative to γ if $(\tilde{c}, \tilde{\pi})$ is a continuation equilibrium and

$$\delta^* \in \arg\max_{\delta \in \Delta(\Gamma)} V(\delta; \delta^*, \tilde{c}, \tilde{\pi}).$$

Goal

Show that the set of indirect mechanisms Γ can be "reproduced" by a universal set of mechanisms such that

(i) Agents report there private information truthfully.
(ii) Agents obey participation recommendations.

Epstein and Peters (1999) develop "language", T, rich enough for agents to report mechanisms.

Language and type space

Let T be the language to describe mechanisms.

- And $\Theta \times T$ is seen as an agent's generalized type space.

Define the mapping $m : \Theta^2 \times T^2 \to X$.

- And let $X^{\Theta^2 \times T^2}$ denote the set of such mappings.

Then m can be seen as a **"direct mechanism"** from the agents' type space to actions.

- $m(\theta, \theta', t', t'')$ assigns an action directly to reports (θ, t') and (θ', t'').

Language

By definition, if T is a language to describe mechanisms then it should the case that there exists a one-to-one mapping

$$\psi : T \rightarrow X^{\Theta^2 \times T^2}.$$

That is, an element in $t \in T$ describes a direct mechanism m in $X^{\Theta^2 \times T^2}$.

$\psi(t)$ can be interpreted as the direct mechanism described by t.

- And $\psi(t)(\theta, \theta', t', t'')$ as the action induces by reports $(\theta, \theta', t', t'')$.

Hence T refers to direct mechanisms that have the same language as an input.

Main result

The main result of Epstein and Peters (1999) is that for every indirect mechanism γ there exists a T and a ψ that can replicate the allocation of Γ.

What is the nature of T? How to described mechanisms?

The idea is that, a mechanism induces a **distribution on payoffs.**

- Hence a mechanism can be described by the payoffs it induces.
- Mechanisms inducing the same payoffs can be seen as "equivalent".