Information, contracts and competition

Chapter IV
Principal Agent Models with many agents and many principal

This part of the course give an elementary model of competition between principals, principal being competing on the mechanism they propose.

Fall 2018

Università di Roma "Tor Vergata" - A. Chassagnon

Competing Principal in a non cooperative game

As we already said a Principal-agent model can be understood as an extension of a noncooperative game. However, here, at the ground level, we consider principal engaged in a non cooperative competing between themself.

The key element will be to understand the role of the agents in this competition.

Roadmap

0) Introduction
1) On competing mechanism under exclusive competition
2) Extensions

1. On competing mechanism under exclusive competition,

From - A. Attar, E. Campioni, G. Piaser, P. (2018). On competing mechanism under exclusive competitions. *Games and Economic Behavior*, Volume 111, September 2018, Pages 1-15.

Agents and Principals

J principals, I agents which type is unknown to the principals. The joint distribution of types if common knowledge

We refer to a scenario in which several principals (indexed by $j \in \mathcal{J} = \{1, \ldots, J\}$) contract with several agents (indexed by $i \in \mathcal{I} = \{1, \ldots, I\}$). Each agent i has private information about her type $\omega^i \in \Omega^i$. Each Ω^i is taken to be finite, and we let $p(\omega)$ be the probability of the array of types $\omega = (\omega^1, \ldots, \omega^I) \in \Omega = \times_{i \in \mathcal{I}} \Omega^i$.

Decisions (actions) of the Principal and of the agents

Participation is a key element of this competing mechanism. Indeed, participation was not really discussed in the Myerson game with one principal, as the participation with that principal was a natural issue. Here, with many principals, the first thing that the principal wants to contract is the participation with agents.

The decisions of the agents are restricted to participation only. And the participation is made before the principal receive any information.

The decisions of the principal are finite (technical assumption only).

We let $x_j \in X_j$ be a decision available to principal j, with X_j being a finite set for every $j \in \mathcal{J}$. Similarly, $a^i_j \in A^i_j = \{Y, N\}$ represents the decision of agent i to participate with principal j, in which N stands for not participating. We take $v_j : X \times A \times \Omega \to \mathbb{R}$ and $u^i : X \times A \times \Omega \to \mathbb{R}$ to be the payoff functions of principal j and of agent i, respectively, with $X = \times_{j \in \mathcal{J}} X_j$ and $A = \times_{i \in \mathcal{I}} \times_{j \in \mathcal{J}} A^i_j$.

- We will explore carrefully the not-so-simple structure of the set A.
- There are $I + J$ payoff functions to consider, but at the end, there will be J mechanism and a partition of size J of the I agents.

Exclusive communication

- Principal decisions will depend upon Ω and A, but also on the communication scheme that we consider below

Here, "messages" are the agents reporting

Communication occurs via the public mechanisms posted by principals, and via the messages sent by agents: each agent i sends a private message $m_j^i \in M_j^i$ to principal j. We let each M_j^i be finite, include the element $\{\emptyset\}$ corresponding to the information "agent i does not communicate with principal j", and be sufficiently rich that $\Omega^i \subset M_j^i$ for every i and every j.[4]

One agent is supposed to send a message to a particular principal if he participates with him. However, To make the notation easier, we will say that every agent send a message to every principal, the $\emptyset$ message being systematically the one sent by an agent that does not participate with one of the principal, to that principal.

Assumption E. *The set of participation and communication decisions for each agent i is $S^i = \{(m^i, a^i) \in M^i \times A^i \; : \; a_j^i = Y$ for at most one j and $m_j^i = \emptyset$ iff $a_j^i = N\}$, with $M^i = \underset{j \in \mathcal{J}}{\times} M_j^i$ and $A^i = \underset{j \in \mathcal{J}}{\times} A_j^i$.*

J Mechanisms

Messages are by nature private messages. We suppose at the first glance that the principal can receive message from any agent, that is the set of message is $M_j = \times_{i \in \mathcal{I}} M_j^i$.

The principal receive also the array of participation decisions $a_j = (a_j^1, a_j^2, ..., a_j^I) \in A_j$

Each principal perfectly observes the set of agents who participate with him. Hence, principal j can make his decisions contingent on the array of messages m_j he receives, with $m_j = (m_j^1, m_j^2, \ldots, m_j^I) \in M_j = \underset{i \in \mathcal{I}}{\times} M_j^i$, and on the array of participation decisions $a_j = (a_j^1, a_j^2, \ldots, a_j^I) \in A_j = \underset{i \in \mathcal{I}}{\times} A_j^i = \{Y, N\}^{|\mathcal{I}|}$. Formally, we say that a mechanism posted by principal j is a mapping $\gamma_j : M_j \times A_j \to \Delta(X_j)$. We refer to $\gamma_j(m_j, a_j)$ as to the probability distribution over X_j induced by the array of messages and actions (m_j, a_j), and to $\gamma_j(x_j | m_j, a_j)$ as to the probability that $\gamma_j(m_j, a_j)$ assigns to the decision $x_j \in X_j$. We let Γ_j be the set of mechanisms available to principal j, and denote $\Gamma = \underset{j \in \mathcal{J}}{\times} \Gamma_j$. Observe that each $\Gamma_j = (\Delta(X_j))^{|M_j \times A_j|}$ is compact in the product topology.

Remark : Notice that in response to the array of messages and actions addressed to him the $\gamma_j(m_j, a_j)$ decision of principal j is a distribution of his actions (belonging to $\Delta(X_j)$)

Timing

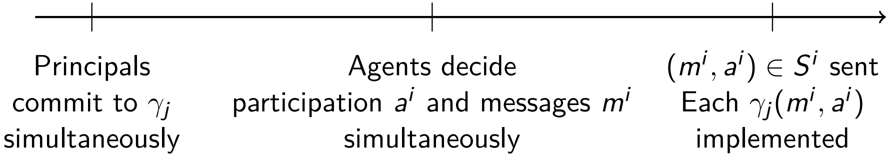

- Consider the possibility for the principal to choose a mixed strategy, that is a probability distribution over Γ_j, Γ_j being the set of all the mechanisms available to principal j

[what is the size of Γ_j, and of $\Gamma = \times \Gamma_j$?]

- We also allow agents to play mixed strategies

Agents strategies and payoffs

Here, when agents play mixed strategy, $\lambda^i(\gamma, \omega^i)$ is the joint probability distribution over message and participation,

Under Assumption E, a strategy for agent i is a measurable mapping, $\lambda^i : \Gamma \times \Omega^i \to \Delta(S^i)$, that associates to each type and to each array of posted mechanisms a probability distribution over such restricted decisions. We denote $\lambda^i(\gamma, \omega^i)$ the joint probability distribution over $(m^i, a^i) \in S^i$ for agent i of type ω^i given γ, and $\lambda^i(m^i, a^i | \gamma, \omega^i)$ the (joint) probability assigned to the pair (m^i, a^i) by $\lambda^i(m^i, a^i)$. The corresponding expected payoff to type ω^i of agent i is:

$$U^i\left(\gamma, \lambda; \omega^i\right) = \sum_{\Omega^{-i}} \sum_{M \times A} \sum_{X} u^i\left(x, a, \omega^i, \omega^{-i}\right) \prod_{j \in \mathcal{J}} \gamma_j\left(x_j | m_j, a_j\right) \prod_{h \in \mathcal{I}} \lambda^h\left(m^h, a^h | \gamma, \omega^h\right) p\left(\omega^{-i} | \omega^i\right) \tag{1}$$

We consider here the interim payoff and the ex ante payoff of the principal :

$$V_j\left(\gamma_j, \gamma_{-j}, \lambda\right) = \sum_{\Omega} \sum_{A} \sum_{X} v_j\left(x, a, \omega\right) \sum_{M} \prod_{k \in \mathcal{J}} \gamma_k\left(x_k | m_k, a_k\right) \prod_{i \in \mathcal{I}} \lambda^i\left(m^i, a^i | \gamma, \omega^i\right) p\left(\omega\right). \tag{2}$$

bayesian equilibrium

We let G^{Γ} be the competing-mechanism game induced by a given Γ.[5] As in Epstein and Peters (1999) and Han (2007), we focus on subgame perfect Nash equilibria (SPNE) of the game G^{Γ} in which principals play pure strategies. The strategies (γ, λ) constitute an SPNE of G^{Γ} if:

1. λ is a continuation equilibrium. That is, for every $\gamma \in \Gamma$, the strategies $\left(\lambda^{i}, \lambda^{-i}\right)$ constitute a Bayes–Nash equilibrium of the agents' game induced by γ;
2. given γ_{-j} and λ, $\gamma_j \in \underset{\gamma_j' \in \Gamma_j}{argmax}\, V_j\left(\gamma_j', \gamma_{-j}, \lambda\right)$ for every $j \in \mathcal{J}$.

Exemple : competitive insurance

In Rothchild and Stliglitz, menus are not contingent to participation. Here, they are.

In their canonical analysis, Rothschild and Stiglitz (1976) study strategic competition between insurance companies for the exclusive right to serve several ex-ante identical agents. Each agent's type space is $\Omega = \left\{\omega^L, \omega^H\right\}$. Uncertainty is idiosyncratic: each customer faces a binary risk on her endowment $e \in \{e_b, e_g\}$. Let $f_b(\omega)$ be the probability of the individual state b for an agent of type $\omega \in \Omega$, and $f_g(\omega) = 1 - f_b(\omega)$. These random variables are independently distributed across all agents and identically distributed across agents of the same type. Each agent privately observes her type. The payoff to an agent of type ω is $f_b(\omega)u(e_b+d_b)+(1-f_b(\omega))\,u(e_g+d_g)$, where $(d_b, d_g) \in \mathbb{R}^2$ is the state-contingent coverage purchased from the company it participates with. Attention is restricted to the game in which insurers post (deterministic) direct mechanisms which assign to each agent a pair of state-contingent insurance contracts, as a function of her participation decision and declared type. In addition, such mechanisms induce agents to be truthful to the company they participate with.[10] Communication is only allowed between each insurer and its customers, implying that Assumption E is satisfied.

Failure of the revelation principle

In that economy, we obtain more equilibrium that in the original Rothchild and Stiglitz economy, which is the restriction of this economy to direct mechanism. This is what can be called a failure of the revelation principle

1. The example shows that a monopolistic allocation can arise at equilibrium in a standard insurance setting with exclusive competition. This obtains by letting principals use indirect mechanisms, but the same allocation cannot be supported in an equilibrium of G^D. The (equilibrium) indirect mechanism of $P1$ makes available a system of threats that allow agents to punish his opponent when he attempts at profitably deviating. Indeed, following any such deviation, there is an equilibrium of the agents' game that keeps $P2$ excluded from trade and yields every type ω the coverage $(\hat{d}_b(\omega), \hat{d}_g(\omega))$. Importantly, the threats are effective despite the exclusive nature of competition, captured by Assumption E. Direct mechanisms turn out not to be flexible enough to reproduce all these threats, which shrinks the set of equilibrium outcomes they can support. In this respect, the

RSI

tumers). An insurer can provide a state-contingent coverage $(d_b, d_g) \in \mathcal{D}$ to each of the customers who participate with him. The set of feasible insurance contracts $\mathcal{D} \subset \mathbb{R}^2_+$ is compact. The payoff to type $\omega \in \left\{\omega^L, \omega^H\right\}$ of each agent $i = 1, 2$ is

$$U(d_b, d_g; \omega) = f_b(\omega)\, u(e_b + d_b) + (1 - f_b(\omega))\, u(e_g + d_g), \tag{4}$$

with (e_b, e_g) such that $e_b < e_g$ being its state-contingent endowment.[16] For each ω, the function $U(.,.;\omega)$ is twice continuously differentiable, strictly quasi-concave and the standard single-crossing property is satisfied. We let ω^H be the "high risk" type: one hence has $f_b(\omega^H) > f_b(\omega^L)$ for each agent.

RSII

A particular strategy of P1 leads to the monopole allocation, with the addition of two arbitrary messages m and m'.

i) if both agents participate with him sending a pair of messages $(\omega, \tilde{\omega}) \in \left\{\omega^L, \omega^H\right\}^2$, he provides the coverages $\left(d_b^M(\omega), d_g^M(\omega)\right)$ and $\left(d_b^M(\tilde{\omega}), d_g^M(\tilde{\omega})\right)$, respectively;

ii) if both agents participate with him sending a pair of messages in the set $\{m, m'\}^2$, he provides the coverage $\left(\hat{d}_b(\omega^L), \hat{d}_g(\omega^L)\right) = \underset{d \in D}{argmax}\, U(d_b, d_g; \omega^L)$ to each agent sending m, and the coverage $\left(\hat{d}_b(\omega^H), \hat{d}_g(\omega^H)\right) = \underset{d \in D}{argmax}\, U(d_b, d_g; \omega^H)$ to each agent sending m';

iii) for every other combination of participation choices and messages, he provides the coverage $\left(d_b^M(\omega), d_g^M(\omega)\right)$ to the agent sending the message ω and the coverage $d_b(\emptyset) = d_g(\emptyset) = 0$ to the agent sending the messages m, m' or $\emptyset$.

In addition, let the other principal $P2$ offer the null coverage $d_b(\emptyset) = d_g(\emptyset) = 0$ for every participation and communication choices.

This outcome can be supported in a pure strategy equilibrium of a competing-mechanism game $G\Gamma$ in which each agent's message set contains, beyond her individual types and the degenerate message $\{\emptyset\}$, the two additional messages m and m'.