Competitive Information Design (work in progress)

Frédéric Koessler, Marie Laclau & Tristan Tomala

PGMO, November 2016

Introduction

Competition between multiple information designers who want to persuade a single decision maker to choose their most preferred action.

- Game between mediators (senders) who choose information structures.
- Game between experts who choose statistical experiments about a piece of information.
- Splitting games.

Applications:

- pharmaceutical firms perform clinical trials in order to get the authorization of selling the drug
- lobbyists design studies (e.g. about climate) to influence a policy

The setup

The game unfolds as follows:

- n senders have independent types, θ_i for each i, with $\theta = (\theta_1, \ldots, \theta_n)$.
- Sender i chooses a statistical experiment $x_i : \Theta_i \to \Delta(S_i)$ with S_i a set of signals *(without knowing his type !)*. Senders play simultaneously.
- Statistical experiments are drawn.
- The Receiver gets a signal from each sender and chooses an action $a \in A$.
- Payoffs $u_i(\theta, a)$ for each Sender i, and $v(\theta, a)$ for the Receiver.

Preliminary Results

(i) General equilibrium existence result
   - It requires designers to be able to randomize over statistical experiments
   - . . . or to have a continuum of messages (even with binary actions and states)

(ii) Existence (even with finite message spaces and no randomization by the designers) and characterization of equilibria in the class of "rectangular" games

(iii) *(In progress)* Characterization of equilibria in dynamic versions of the games

Statistical experiments and splittings

Set of states Θ, prior $p \in \Delta(\Theta)$.

Definition

A statistical experiment is a mapping $x : \Theta \to \Delta(S)$ with S a set of signals.

Let,

$$p_s(\theta) = \mathbb{P}(\theta \mid s) = \frac{x(s|\theta)p(\theta)}{\sum_\eta x(s|\eta)p(\eta)} \text{ and } \lambda_s = \mathbb{P}(s) = \textstyle\sum_\eta p(\eta)x(s \mid \eta).$$

One has:

$$p(\theta) = \sum_s \mathbb{P}(s)\mathbb{P}(\theta \mid s) = \sum_s \lambda_s p_s(\theta),$$

a splitting of p.

Statistical experiments and splittings

Set of states Θ, prior $p \in \Delta(\Theta)$.

Definition

A statistical experiment is a mapping $x : \Theta \to \Delta(S)$ with S a set of signals.

Let,

$$p_s(\theta) = \mathbb{P}(\theta \mid s) = \frac{x(s|\theta)p(\theta)}{\sum_\eta x(s|\eta)p(\eta)} \text{ and } \lambda_s = \mathbb{P}(s) = \textstyle\sum_\eta p(\eta)x(s \mid \eta).$$

One has:

$$p(\theta) = \sum_s \mathbb{P}(s)\mathbb{P}(\theta \mid s) = \sum_s \lambda_s p_s(\theta),$$

a splitting of p. Conversely, if $p = \sum_s \lambda_s p_s$ is a splitting, let:

$$x(s \mid \theta) = \frac{\lambda_s p_s(\theta)}{p(\theta)},$$

then, $\mathbb{P}(\theta \mid s) = p_s(\theta)$ (Aumann-Maschler, 1967).

Information design game with one designer

The receiver. Action set A, payoff $v(\theta, a)$. Given a belief p:

$$\max_a v(p,a) = \max_a \sum_\theta p(\theta) v(\theta, a).$$

- Optimal actions $A(p)$, optimal mixed actions $Y(p) := \Delta A(p)$.
- A tie-breaking-rule (TBR) is a selection $y(p) \in Y(p)$.

Example: Two states $\{\theta_0, \theta_1\}$, two actions $\{a_0, a_1\}$, and payoffs:

$v(\theta, a) =$

	a_0	a_1
θ_0	0	-3
θ_1	0	1

The receiver takes action a_1 if his belief p on θ_0 is: $p < \frac{1}{4}$ (indifferent in $\frac{1}{4}$).

Information design game with one designer (Kamenica and Gentzkow, 2011)

The designer. He chooses $x : \Theta \to \Delta(S)$ *(without knowing the state !)*. The receiver gets a signal s, and chooses action a. Payoff $u(\theta, a)$ to the designer.

Information design game with one designer (Kamenica and Gentzkow, 2011)

The designer. He chooses $x : \Theta \to \Delta(S)$ *(without knowing the state !)*. The receiver gets a signal s, and chooses action a. Payoff $u(\theta, a)$ to the designer.

Solving the game. The receiver chooses a TBR $y(p)$. The program of the designer is then:

$$\max\{\sum_s \lambda_s \sum_\theta p_s(\theta) u(\theta, y(p_s)) : p = \sum_s \lambda_s p_s\}.$$

Information design game with one designer (Kamenica and Gentzkow, 2011)

The designer. He chooses $x : \Theta \to \Delta(S)$ *(without knowing the state !)*. The receiver gets a signal s, and chooses action a. Payoff $u(\theta, a)$ to the designer.

Solving the game. The receiver chooses a TBR $y(p)$. The program of the designer is then:

$$\max\{\sum_s \lambda_s \sum_\theta p_s(\theta) u(\theta, y(p_s)) : p = \sum_s \lambda_s p_s\}.$$

If the receiver selects the designer-preferred TBR (as in Kamenica and Gentzkow, 2011), then the induced designer's payoff as a function of the posteriors is u.s.c.: $U_y(p) = \sum_\theta p(\theta) u(\theta, y(p))$.

Information design game with one designer (Kamenica and Gentzkow, 2011)

The designer. He chooses $x : \Theta \to \Delta(S)$ *(without knowing the state !)*. The receiver gets a signal s, and chooses action a. Payoff $u(\theta, a)$ to the designer.

Solving the game. The receiver chooses a TBR $y(p)$. The program of the designer is then:

$$\max\{\sum_s \lambda_s \sum_\theta p_s(\theta) u(\theta, y(p_s)) : p = \sum_s \lambda_s p_s\}.$$

If the receiver selects the designer-preferred TBR (as in Kamenica and Gentzkow, 2011), then the induced designer's payoff as a function of the posteriors is u.s.c.: $U_y(p) = \sum_\theta p(\theta) u(\theta, y(p))$.

We then have:

$$\sup\{\sum_s \lambda_s U_y(p_s) : p = \sum_s \lambda_s p_s\} = \text{Cav}\, U_y(p)$$

where $\text{Cav}\, U_y(p)$ is the smallest concave function above $U_y(p)$ (Aumann and Maschler, 1967).

Example

Coming back to the previous exemple, and assume that the designer's payoff is 1 if a_1 and 0 otherwise.

Designer's utility

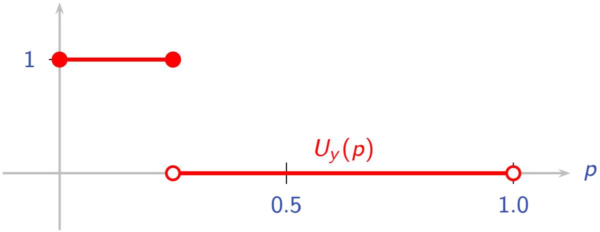

Example

Coming back to the previous exemple, and assume that the designer's payoff is 1 if a_1 and 0 otherwise.

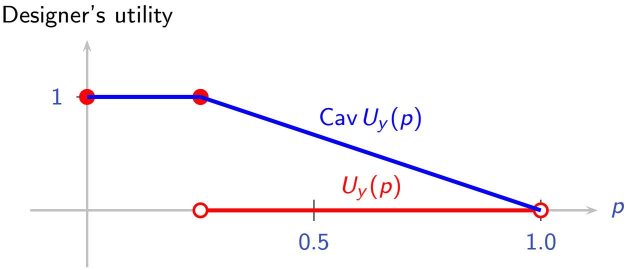

In $p = \frac{1}{2}$, the designer can induce posteriors $\frac{1}{4}$ and 1 with probabilities $\frac{2}{3}$ and $\frac{1}{3}$. His expected payoff is then $\frac{2}{3}U_y(\frac{1}{4}) + \frac{1}{3}U_y(1) = \text{Cav}\, U_y(\frac{1}{2})$.

Remarks

(i) $\sharp S = \min\{\sharp\Theta, \sharp A\}$ messages are enough (Caratheodory's theorem)

(ii) For general TBR, there might not be a maximum. No 0-equilibrium, but ε-equilibria.

(iii) There exists a TBR such that there is a maximum, thus a 0-equilibrium (favor the sender).

(iv) $\operatorname{Cav} U_y(p)$ might depend on $y(p)$.
Example: Two states $\{\theta_0, \theta_1\}$, two actions $\{a_0, a_1\}$, and payoffs:

$$v(\theta, a) = \begin{array}{|c|c|c|} \hline & a_0 & a_1 \\ \hline \theta_0 & 3 & 3 \\ \hline \theta_1 & 1 & 0 \\ \hline \end{array}$$

Remarks

(i) $\sharp S = \min\{\sharp\Theta, \sharp A\}$ messages are enough (Caratheodory's theorem)

(ii) For general TBR, there might not be a maximum. No 0-equilibrium, but ε-equilibria.

(iii) There exists a TBR such that there is a maximum, thus a 0-equilibrium (favor the sender).

(iv) $\operatorname{Cav} U_y(p)$ might depend on $y(p)$.
Example: Two states $\{\theta_0, \theta_1\}$, two actions $\{a_0, a_1\}$, and payoffs:

$$v(\theta, a) = \begin{array}{|c|c|c|} \hline & a_0 & a_1 \\ \hline \theta_0 & 3 & 3 \\ \hline \theta_1 & 1 & 0 \\ \hline \end{array}$$

Unless the problem is regular:

$$\forall a, \exists p, \ s.t. \ A(p) = \{a\}.$$

Then $\operatorname{Cav} U_y(p)$ does not depend on y and is the value of all limits of ε-equilibrium payoffs (as $\varepsilon \to 0$).

Competition

Let $\Theta = \Theta_1 \times \cdots \times \Theta_n$ and $p = p_1 \otimes \cdots \otimes p_n$. Sender i has payoff $u_i(\theta, a)$ and chooses $x_i : \Theta_i \to \Delta(S_i)$, or equivalently a splitting of p_i. This defines an $n+1$-player game Γ where:

- Senders choose experiments (splittings) simultaneously.
- Receiver gets a message from each sender (or a posterior) and chooses $a \in A$.

Let $\Gamma(M_1, \ldots, M_n, R)$ the $n+1$-player game where $\sharp S_i = M_i$.

In any (subgame perfect) equilibrium, the strategy of the receiver is a TBR $y(p) \in Y(p) = \Delta A(p)$.

Fixing a TBR y, let $\Gamma(M_1, \ldots, M_n, y)$ be the induced n-player game on senders.

Example: Being believed to be the best

Two senders, binary states with $\Theta_1 = \Theta_2 = \{0,1\}$, 0-sum. Belief about θ_1: $(1-p,p)$, belief about θ_2: $(1-q,q)$. Prior beliefs over θ_1 and θ_2 are $(\frac{1}{2},\frac{1}{2})$ each. Binary actions a_1, a_2.

Receiver's payoffs: $v(\theta,a_1) = \mathbf{1}\{\theta_1 = 1\}$, $v(\theta,a_2) = \mathbf{1}\{\theta_2 = 1\}$.

Senders' payoffs: $u_1(a_1) = -u_2(a_1) := u(a_1) = 1$, $u(a_2) = -1$.

The receiver chooses a_1 if $p > q$, a_2 if $p < q$. If $p = q$, let $y(p,q) = (\frac{1}{2},\frac{1}{2})$.

Thus $u(p,q) = \mathbf{1}\{p > q\} - \mathbf{1}\{p < q\}$.

Interpretation: the receiver can buy a product from firm 1 or firm 2.

Are 2 messages enough?

Consider a 2-splitting μ of Sender 2, $\frac{1}{2} = \lambda q + (1-\lambda) q'$ with $q \leq \frac{1}{2}$ and $q' \geq \frac{1}{2}$. The best-reply payoff of Sender 1 is $\mathrm{Cav}\, u(p,\mu)$, with:

$$u(p,\mu) = \begin{cases} -1 & \text{for } p < q, \\ -(1-\lambda) & \text{for } p = q, \\ \lambda - (1-\lambda) = 2\lambda - 1 & \text{for } q < p < q', \\ \lambda & \text{for } p = q', \\ 1 & \text{for } p > q'. \end{cases}$$

If $q \neq 0$, $q' \neq 1$ and $q \neq q'$, then $\mathrm{Cav}\, u(\frac{1}{2},\mu) \geq \lambda(2\lambda - 1) + (1-\lambda)1$ $= 1 - 2\lambda(1-\lambda) \geq \frac{1}{2}$ (and similarly for $q = 0$ and $q' < 1$).

If $q = 0$ and $q' = 1$ (full revelation), then $u(0,\mu) = -\frac{1}{2}$, $u(p,\mu) = 0$ for $0 < p < 1$ and $u(1,\mu) = \frac{1}{2}$. Then, $\mathrm{Cav}\, u(\frac{1}{2},\mu) = \frac{1}{4}$.

If $q = q'$ (no revelation), then $u(p,\mu) = -1$ for $p < \frac{1}{2}$, $u(p,\mu) = 0$ for $p = \frac{1}{2}$, $u(p,\mu) = 1$ for $p > \frac{1}{2}$. Then, $\mathrm{Cav}\, u(\frac{1}{2},\mu) = 1$.

If $q > 0$ and $q' = 1$, then $\mathrm{Cav}\, u(\frac{1}{2},\mu) \geq \lambda(2\lambda - 1) + (1-\lambda)\lambda \geq \lambda^2 \geq \frac{1}{4}$.

Are 2 messages enough?

We conclude that

$$\min_{\mu_2 \in Split_{M_2=2}(\frac{1}{2})} \sup_{\mu_1 \in Split_{M_1=2}(\frac{1}{2})} u = \frac{1}{4}.$$

By symmetry,

$$\max_{\mu_1 \in Split_{M_1=2}(\frac{1}{2})} \inf_{\mu_2 \in Split_{M_2=2}(\frac{1}{2})} u = -\frac{1}{4}.$$

There is thus no equilibrium (for these strategy sets).

An equilibrium

However, the game has a symmetric equilibrium where μ_1 and μ_2 are given by the uniform distribution over $[0,1]$. Indeed, if q is uniformly distributed then:

$$u(p,\mu_2) = \mathbb{E}u(p,q) = \int_0^p 1 + \int_p^1 -1 = 2p-1.$$

If $\sum_s \lambda_s p_s = \frac{1}{2}$, then $\sum_s \lambda_s(2p_s-1)=0$. So this is an equilibrium with payoff 0.

Interpretation 1. This is a mixed strategy equilibrium, a randomization over all splittings with two points. Draw p uniformly, split $\frac{1}{2} = \frac{1}{2}p + \frac{1}{2}(1-p)$.

Interpretation 2. This is a (pure) equilibrium of the game with infinite message sets. In state 0, draw $x \in [0,1]$ from $dF_0(x) = 2(1-x)dx$ and in state 1, draw $x \in [0,1]$ from $dF_1(x) = 2xdx$.

Existence result

Theorem

For all number of messages $M_1, \ldots, M_n$, the $n+1$-player game $\Gamma(M_1, \ldots, M_n, R)$ has an equilibrium in mixed strategies.

Proof. Fix an arbitrary TBR y_0 and consider the n-player game with payoff function $U_{y_0}(p)$. Consider $\bar{U}_{y_0}(p) = \{\lim U_{y_0}(p^n) : \lambda^n \to \lambda,\, p^n \to p\}$. We have:

$$\operatorname{co} \bar{U}_{y_0}(p) \subseteq \{\sum_s \lambda_s u(p_s, y(p_s)) \ : \ y(\cdot) \in \Delta A(\cdot)\}.$$

From Simon and Zame (1990), there exists a selection $V(p) \in \operatorname{co} \bar{U}_{y_0}(p)$ which has an equilibrium. Thus, there exists a TBR y_p such that

$$\sum_s \lambda_s(p) u(p_s, y_p(p_s))$$

has an equilibrium.

Existence result: remarks

- For existence, we cannot exclude that the TBR y^* (strategy of the receiver) depends also on the splitting profile $(\lambda, \mathbf{p})$
- If we fix an arbitrary TBR y^*, an equilibrium (and even an ε-equilibrium) may not exist
- If $|M_i| \geq |\Theta_i|$ for each i, then a (mixed) equilibrium outcome of the game with messages $M_1, \ldots, M_n$ is an equilibrium outcome of the game with arbitrary splittings (not necessarily with finite support) in which designers use pure strategies (we do not know if the converse is true)

Rectangular problems

Consider the pure persuasion case $u_i(\theta, a) = u_i(a)$, $i = 1, \dots, n$.

- The decision problem of the receiver is *rectangular* if $A = A_1 \times \cdots \times A_n$ and,

$$v(\theta_1, \dots, \theta_n, a_1, \dots, a_n) = v_1(\theta_1, a_1) + \cdots + v_n(\theta_n, a_n).$$

- The TBR is *rectangular* if $y(a_1, \dots, a_n \mid p_1, \dots p_n) = \times_i y(a_i \mid p_i)$ and $y(\cdot \mid p_i) \in \Delta A_i(p_i)$.

Rectangular problems

Consider the pure persuasion case $u_i(\theta, a) = u_i(a)$, $i = 1, \dots, n$.

- The decision problem of the receiver is *rectangular* if $A = A_1 \times \cdots \times A_n$ and,

$$v(\theta_1, \dots, \theta_n, a_1, \dots, a_n) = v_1(\theta_1, a_1) + \cdots + v_n(\theta_n, a_n).$$

- The TBR is *rectangular* if $y(a_1, \dots, a_n \mid p_1, \dots p_n) = \times_i y(a_i \mid p_i)$ and $y(\cdot \mid p_i) \in \Delta A_i(p_i)$.

Interpretation: there is one receiver for each sender i, with payoff $v_i(\theta_i, a_i)$.

Rectangular problems

Consider the pure persuasion case $u_i(\theta, a) = u_i(a)$, $i = 1, \ldots, n$.

- The decision problem of the receiver is *rectangular* if $A = A_1 \times \cdots \times A_n$ and,

$$v(\theta_1, \ldots, \theta_n, a_1, \ldots, a_n) = v_1(\theta_1, a_1) + \cdots + v_n(\theta_n, a_n).$$

- The TBR is *rectangular* if $y(a_1, \ldots, a_n \mid p_1, \ldots p_n) = \times_i y(a_i \mid p_i)$ and $y(\cdot \mid p_i) \in \Delta A_i(p_i)$.

Interpretation: there is one receiver for each sender i, with payoff $v_i(\theta_i, a_i)$.

Lemma

The set of mixed actions induced by splitting and optimal decisions

$$C_i(p_i) = \left\{ x_i \in \Delta(A_i) \ : \ \exists(\lambda_{is}, q_{is}, y_i),\ p_i = \sum_s \lambda_{is} q_{is} \text{ and } x_i(a_i) = \sum_s \lambda_{is} y_i(a_i | q_{is}) \right\}$$

is convex and closed in $\Delta(A_i)$.

Rectangular problems

Denote $G(p)$ the n-player game $(C_1(p_1), \ldots, C_n(p_n), u_1, \ldots, u_n)$. This is a finite game with convex constraints on mixed strategies. Let $E(p)$ be the non-empty and compact set of distributions of actions induced by equilibria of $G(p)$.

Theorem

(i) *$E(p)$ is the set of distributions of actions induced by all equilibria of the $(n+1)$-player rectangular game.*

(ii) *If v_i is regular for each i (no action is dominated), then we can fix any TBR and $E(p)$ is the set of limit distributions induced by all ε-equilibria (as $\varepsilon \to 0$) of the n-player information design game, given this TBR.*

Remarks.

- Number of messages has to be $\geq \sharp A_i$.
- Other proof of existence with a finite number of messages and pure strategies for the designers.

Repeated competition between designers

Consider a rectangular problem. At each stage $t = 1, 2, \ldots$, senders choose splittings simultaneously. When beliefs do not change anymore, the Receiver takes an action.

Examples:

(i) Full Revelation is an equilibrium of both the one-shot and the repeated game + No Revelation in the repeated game:

$$u(p,q) = \begin{array}{|c|c|c|} \hline 1,1 & 4,0 & 1,1 \\ \hline 0,4 & 3,3 & 0,4 \\ \hline 1,1 & 4,0 & 1,1 \\ \hline \end{array}$$

Repeated competition between designers

(ii) Same equilibrium in the one-shot and in the repeated game (No Revelation):

$u(p,q) =$

3,3	0,4	3,3
4,0	1,1	4,0
3,3	0,4	3,3

(iii) No Revelation is not an equilibrium of the repeated game although it is an equilibrium of the one-shot game + Full Revelation in the repeated game:

$u(p,q) =$

3,2	0,0	3,2
0,0	1,1	0,0
3,2	0,0	3,2

Repeated competition between designers

(iv) The Full Revelation equilibrium requires two rounds of communication:

$$u(p,q) = \begin{array}{|c|c|c|}\hline 3,2 & 0,0 & 3,2 \\ \hline 4,0 & 1,1 & 4,0 \\ \hline 3,2 & 0,0 & 3,2 \\ \hline \end{array}$$

Sender 1 has to reveal first.

Thank you !