Lecture on Sender Receiver Games

Tristan Tomala, HEC Paris

Equilibria in Games: existence, selection and dynamics. Santiago, Chile, January 2017.

Outline

Sender Receiver Games

A Sender Receiver Game is a special kind of game with incomplete information.

- The set of states is Θ with a prior probability p.
- Player 1 knows the state and chooses a message $m \in M$.
- Player 2 does not know the state and observes m. Then he chooses an action $a \in A$.
- Player 1's payoff is $u(\theta, a)$, player 2's payoff is $v(\theta, a)$.

Player 1 is an expert, Player 2 is a decision maker.
Information is transmitted strategically.
Since utilities are different, information is imperfectly transmitted.

Example

Consider the Sender-Receiver game defined by:

	a_1	a_2	
θ_1	$(1,5)$	$(0,1)$	$\frac{1}{2}$
θ_2	$(1,-10)$	$(0,1)$	$\frac{1}{2}$

- Player 1 only cares about player 2 taking action a_1.
- For player 2, a_2 is a safe bet, a_1 is a risky bet.
- A behavior strategy for player 1 is $\sigma : \Theta \to \Delta(M)$ denoted $\sigma(m \mid \theta)$.
- A behavior strategy for player 2 is $\tau : M \to \Delta(A)$ denoted $\tau(a \mid m)$.

Example

- There exists a non-informative, babbling equilibrium.
- There is no other equilibrium.
- If $\tau(a_1 \mid m) > \tau(a_1 \mid m')$, then player 1 will choose m for sure.

Outline

Crawford and Sobel

This model is due to Crawford and Sobel (1982).

- The set of states is $\Theta = [0, 1]$ with the Lebesgue measure.
- $M = [0, 1]$, $A = \mathbb{R}$.
- Payoff of player 2: $v(\theta, a) = -(a - \theta)^2$.
- Payoff of player 1: $u(\theta, a) = -(a - \theta - b)^2$, $b > 0$.

Ideally, the decision maker wants to choose $a = \theta$. The expert has a bias and wants the DM to choose $\theta + b$.

There is a non-informative equilibrium where $\sigma(m \mid \theta)$ does not depend on θ and $\tau(a \mid m)$ does not depend on m. Then $a^* = \frac{1}{2}$.

Partitional equilibrium

Let n and integer and a partition

$$0 = t_0 < t_1 < \cdots < t_{n-1} < t_n = 1$$

Consider the strategy of player 1,

$$\sigma(\theta) = \begin{cases} m_k \text{ if } \theta \in [t_{k-1}, t_k) \\ m_n \text{ if } \theta \in [t_{n-1}, 1] \end{cases}$$

with $m_1 < \cdots < m_n$ (all different).

The best-reply of player 2 is $a_k = \tau(m_k) = \mathbb{E}(\theta \mid m_k) = \frac{t_{k-1}+t_k}{2}$ (choose action 0 for other messages).

Given this strategy of player 2, is σ a best-reply for player 1?

Partitional equilibrium

- Player 1 only sends messages from $\{m_1, \ldots, m_n\}$.
- If $\theta \in [t_{k-1}, t_k)$ player 1 should send m_k but may want to deviate to m_{k+1}.
- The equilibrium condition is $u(t_k, a_k) = u(t_k, a_{k+1})$, if $\theta = t_k$, player 1 is indifferent. That is,

$$(t_k - \frac{t_{k-1} + t_k}{2} - b)^2 = (t_k - \frac{t_k + t_{k+1}}{2} - b)^2$$

Equivalently,

$$t_k + b = \frac{\frac{t_{k-1}+t_k}{2} + \frac{t_k+t_{k+1}}{2}}{2} = \frac{t_{k-1} + 2t_k + t_{k+1}}{4}$$

By induction, $t_k = kt_1 + \frac{k(k-1)}{2}4b$.

Partitional equilibrium

- It results $1 = t_n = nt_1 + n(n-1)2b$, $t_1 = 1/n - 2(n-1)b$ and $t_k = k/n - 2k(n-k)b$.

Crawford and Sobel

There exists an n-partitional equilibrium for $n \leq N(b)$ with $N(b)$ the largest n such that $b < \frac{1}{2n(n-1)}$.

- $N(b) \to \infty$ as $b \to 0$.
- Equilibria with larger n are more efficient.
- Full revelation is not possible for $b > 0$.
- Extension beyond quadratic utilities: strictly concave with a unique maximum which increases with θ.

Outline

Cheap Talk

From now on we assume all sets to be finite.

- A state $\theta \in \Theta$ is drawn according to p and told to player 1.
- Player 1 sends a message $m \in M$ to player 2, the strategy is denoted $\sigma(m \mid \theta)$.
- Player 2 chooses an action a, the strategy is denoted $\tau(a \mid m)$.
- Payoffs are $u(\theta, a)$, $v(\theta, a)$.

There exists a babbling equilibrium where σ does not depend on θ and τ does not depend on m.

Let's look at equilibrium conditions on both sides.

Splittings

Take set of states Θ, prior $p \in \Delta(\Theta)$. The strategy of player 1 can be thought of a *statistical experiment* $\sigma : \Theta \to \Delta(M)$.

- $p_m(\theta) = \mathbb{P}(\theta \mid m) = \frac{\sigma(m|\theta)p(\theta)}{\sum_\zeta \sigma(m|\zeta)p(\zeta)}$
- $\lambda_m = \mathbb{P}(m) = \sum_\zeta p(\zeta)\sigma(m \mid \zeta)$.

One has,

$$p(\theta) = \sum_m \mathbb{P}(m)\mathbb{P}(\theta \mid m) = \sum_m \lambda_m p_m(\theta)$$

a splitting of p into a convex combination of posteriors.

Conversely, if $p = \sum_m \lambda_m p_m$ is a splitting let,

$$\sigma(m \mid \theta) = \frac{\lambda_m p_m(\theta)}{p(\theta)}$$

then, $\mathbb{P}(\theta \mid m) = p_m(\theta)$ (Aumann-Maschler 65).

Equilibrium condition for player 2

If player 2 has a belief p over the states, he chooses a mixed action from the set:

$$Y(p) = \{\alpha \in \Delta(A) : \sum_a \alpha(a) \sum_\theta p(\theta)v(\theta, a) = \max_a \sum_\theta p(\theta)v(\theta, a)\}$$

Given a strategy σ of player 1, player 2 who receives the message m has belief p_m.

At a best-reply, player 2 chooses $y_m \in Y(p_m)$.

Equilibrium condition for player 1

Suppose that player 2 chooses $y_m \in \Delta(A)$ after receiving message m.

Consider the best-reply of player 1.

- For each state θ and each message m,

$$\sigma(m \mid \theta) > 0 \Rightarrow u(\theta, y_m) = \max_{m'} u(\theta, y_{m'}) := U_\theta$$

Notice that $\sigma(m \mid \theta) > 0 \Leftrightarrow p_m(\theta) > 0$. Player 1 is indifferent between all messages (and thus posteriors) induced with positive probability.

For each m,

- $u(\theta, y_m) \leq \max_{m'} u(\theta, y_{m'}) = U_\theta$ with equality if $p_m(\theta) > 0$.

Equilibrium characterization

Consider the set $\mathcal{E}$ of tuples $(p, U, V) \in \Delta(\Theta) \times \mathbb{R}^{\Theta} \times \mathbb{R}$, such that $\exists y \in \Delta(A)$:

(i) $U_\theta \geq u(\theta, y)$ with equality if $p(\theta) > 0$,

(ii) $y \in Y(p)$,

(iii) $V = \sum_\theta p(\theta) v(\theta, y)$.

Denote

$$\text{conv}_U(\mathcal{E}) = \left\{ \sum \lambda_m (p_m, U, V_m) : \forall m, (p_m, U, V_m) \in \mathcal{E} \right\}$$

Theorem (Forges 94)

There is an equilibrium of $\Gamma(p)$ with payoff (U, V) if and only if $(p, U, V) \in \text{conv}_U(\mathcal{E})$

Full revelation

A game with full revelation in equilibrium.

	a_1	a_2
θ_1	$(1,1)$	$(0,0)$
θ_2	$(0,0)$	$(3,3)$

No revelation

A game with no revelation in equilibrium.

	a_1	a_2	a_3	
θ_1	$(3,-4)$	$(2,-1)$	$(1,0)$	p
θ_2	$(3,0)$	$(2,-1)$	$(1,-4)$	$1-p$

- We have,

$$Y(p) = \begin{cases} \{a_1\} \text{ if } p < 1/4 \\ \{a_2\} \text{ if } 1/4 < p < 3/4 \\ \{a_3\} \text{ if } p > 3/4 \end{cases}$$

- It is not possible to split p and keep player 1 indifferent.

No revelation

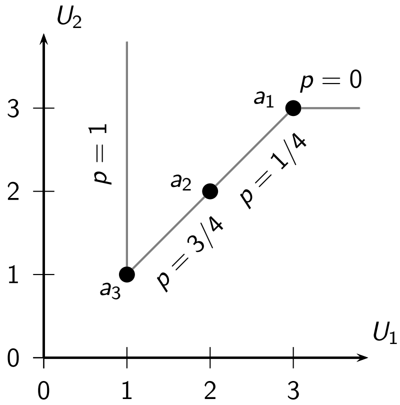

Partial revelation

A game with partial revelation in equilibrium.

	a_1	a_2	a_3	a_4	a_5	
θ_1	$(1,10)$	$(3,8)$	$(0,5)$	$(3,0)$	$(1,-8)$	p
θ_2	$(1,-8)$	$(3,0)$	$(0,5)$	$(3,8)$	$(1,10)$	$1-p$

We have,

$$Y(p) = \begin{cases} \{a_1\} \text{ if } p > 4/5 \\ \{a_2\} \text{ if } 4/5 > p > 5/8 \\ \{a_3\} \text{ if } 5/8 > p > 3/8 \\ \{a_4\} \text{ if } 3/8 > p > 1/5 \\ \{a_5\} \text{ if } 1/5 > p \end{cases}$$

Partial revelation

- For $p = 1/4$, player 2 plays a_4 and player 1's payoff are $(3, 3)$.
- For $p = 3/4$, player 2 plays a_2 and player 1's payoff are $(3, 3)$.
- Player 1 can induce those beliefs by the splitting

$$\frac{1}{2} = (1/2)\frac{1}{4} + (1/2)\frac{3}{4}$$

- Or by the strategy $\sigma(m_1 \mid \theta_1) = \frac{3}{4} = \sigma(m_2 \mid \theta_2)$.
- This equilibrium payoff dominates NR and CR.

Partial revelation

Another game with partial revelation in equilibrium.

	a_1	a_2	a_3	a_4	
θ_1	$4,0$	$2,7$	$5,9$	$1,10$	p
θ_2	$1,10$	$4,7$	$4,4$	$2,0$	$1-p$

$$Y(p) = \begin{cases} \{a_1\} & \text{if } p < 3/10 \\ \{a_2\} & \text{if } 3/10 < p < 3/5 \\ \{a_3\} & \text{if } 3/5 < p < 4/5 \\ \{a_4\} & \text{if } p > 4/5 \end{cases}$$

To make player 1 indifferent, player 2 has to randomize.

Partial revelation

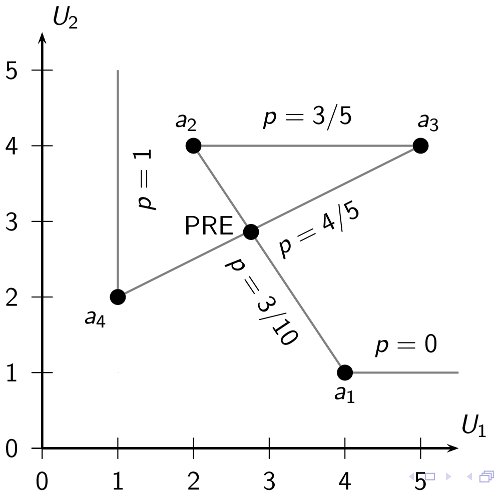

Outline

Long Cheap Talk

- Do multiple rounds of messages change equilibria?
- If only player 1 talks, then this does not change anything, up to a larger message space.
- In the long cheap talk model, at every stage $t = 1, \ldots, T, \ldots$, both players send a cheap talk message to the other (simultaneously).
- Let $\Gamma_T(p)$ be the cheap talk game with T rounds.

Jointly controlled lotteries

Consider a two-player game $G = (A, B, u(a,b), v(a,b))$ and let $M = \mathbb{Z}_M$ be a finite message space. Denote $E(G)$ the set of Nash equilibrium payoffs of G.

Consider the cheap talk game G^*_M:

- First, players select messages simultaneously. Messages are publicly observed.
- Then, G is played.

Claim

$$\text{conv}_M E(G) \subseteq E(G^*_M) \subseteq \text{conv} E(G)$$

- After every profile of public messages, a Nash equilibrium of G is played.
- If X is uniformly distributed in $\mathbb{Z}_M$, then for each y, $X + y$ is also uniformly distributed.

Jointly controlled lotteries

In the game,

	a_1	a_2	a_3	a_4	a_5	
θ_1	$(1,10)$	$(3,8)$	$(0,5)$	$(3,0)$	$(1,-8)$	p
θ_2	$(1,-8)$	$(3,0)$	$(0,5)$	$(3,8)$	$(1,10)$	$1-p$

We can get $\frac{1}{2}((3,3),6)+\frac{1}{2}((1,1),10)$ by a two-step procedure:

- With a JCL, randomize $\frac{1}{2}$PRE$+\frac{1}{2}$FRE. No information revealed.
- Implement the selected equilibrium. Information revealed, player 1 indifferent.

Di-convexity

- A set $E \subseteq \Delta(\Theta) \times \mathbb{R}^{\Theta} \times \mathbb{R}$ is di-convex if for all p, U, the sections $\{(U, V) : (p, U, V) \in E\}$ and $\{(p, V) : (p, U, V) \in E\}$ are convex.
- $\mathrm{di-co}\,(E)$ is the smallest di-convex superset of E.

- A di-martingale is a martingale $(p_n, U_n, V_n)_n$ such that at each step, $p_{n+1} = p_n$ or $U_{n+1} = U_n$.
- $\mathrm{di-span}\,(E)$ is the set of starting points of di-martingales with limit in E.

Characterization

Theorem (Hart 85, Forges 94, Aumann Hart 03)

- *The payoff vector* (U, V) *is an equilibrium payoff of the game* $\Gamma_T(p)$ *for some* T *if and only if* $(p, U, V) \in \mathrm{di-co}\,(\mathcal{E})$.
- *The payoff vector* (U, V) *is an equilibrium payoff of the game* $\Gamma_\infty(p)$ *if and only if* $(p, U, V) \in \mathrm{di-span}\,(\mathcal{E})$.

- An equilibrium is constructed by alternating JCL and splittings such that player 1 is indifferent.
- This generates a di-martingale: at a JCL step, p is constant, at a splitting step U is constant.

Example (Forges 90)

An employer (the DM) chooses to offer a job j_1, j_2, j_3 or j_4, or no job (action j_0) to a candidate (the expert).

The candidate has two possible types θ_1 et θ_2.

	j_1	j_2	j_0	j_3	j_4	
θ_1	6, 10	10, 9	0, 7	4, 4	3, 0	p
θ_2	3, 0	4, 4	0, 7	10, 9	6, 10	$1-p$

The results in the following picture.

Example

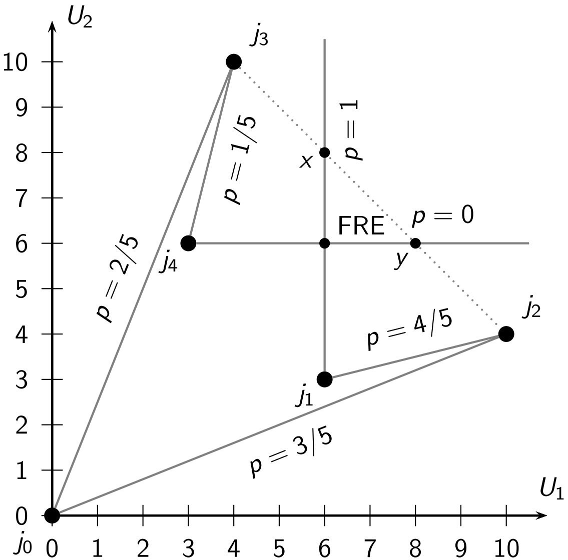

Example

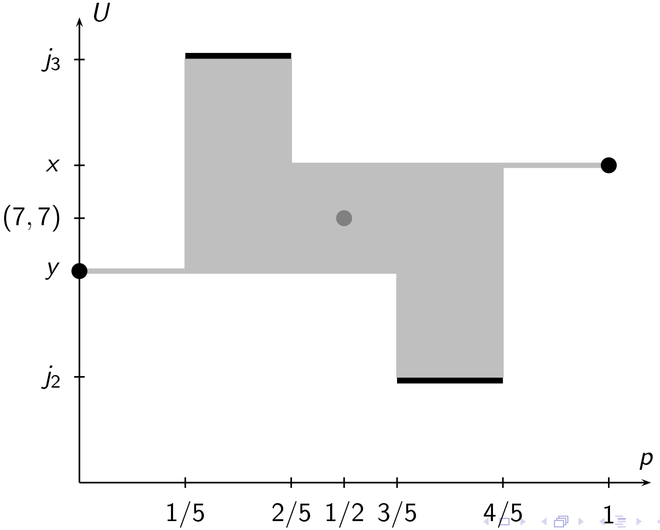

Outline

Communication equilibria

We start from the game given by $p \in \Delta(\Theta)$, $u(\theta, a)$, $v(\theta, a)$.

A communication equilibrium is given by:

- A set of messages M, a set of recommendations R and a mapping $\mu : M \to \Delta(R)$.
- An equilibrium (σ, τ) of the game where:
  - Player 1 knowing θ chooses a message m.
  - A recommendation is drawn with probability $\mu(r \mid m)$ and sent to player 2.
  - Player 2 chooses an action.

Note that the equilibrium is the tuple (μ, σ, τ).

If μ simply transmits the message ($\mu(m \mid m) = 1$), we get a Nash equilibrium.

Communication equilibria

Take μ and (σ, τ) an equilibrium and let,

$$\mu^*(a \mid \theta) = \sum_{m,r} \tau(a \mid r)\mu(r \mid m)\sigma(m \mid \theta)$$

Then, (μ^*, Id, Id) is a canonical communication equilibrium.

Revelation principle

Any communication equilibrium outcome can be obtained by a canonical communication equilibrium where $\mu : \Theta \to \Delta(A)$, $\sigma(\theta) = \theta$, $\tau(a) = a$.

Communication equilibria

$\mu : \Theta \to \Delta(A)$ is a communication equilibrium if and only if:

- $\forall \theta, \theta', \sum_a \mu(a \mid \theta) u(\theta, a) \geq \sum_a \mu(a \mid \theta') u(\theta, a)$
- $\forall a, a', \sum_\theta p(\theta) \mu(a \mid \theta) v(\theta, a) \geq \sum_\theta p(\theta) \mu(a \mid \theta) v(\theta, a')$

The set of communication equilibria is convex and contains the Nash equilibrium outcomes of the cheap talk games.

It also contains the outcomes of long cheap talk (the mediator can replicate the communication protocol).

Mediated talk dominates cheap talk

Consider the following example.

	a_1	a_2	a_3	
θ_1	$(3,3)$	$(1,2)$	$(0,0)$	$1/2$
θ_2	$(2,0)$	$(3,2)$	$(1,3)$	$1/2$

- The babbling equilibrium payoff is $((1,3),2)$.
- This is the only equilibrium payoff with (long) cheap talk.
- $\mu(\theta_1) = \frac{1}{2}a_1 + \frac{1}{2}a_2$, $\mu(\theta_2) = a_2$ is a communication equilibrium with payoff $((2,3),2.25)$.

Outline

Repeated sender receiver game

We still have $p \in \Delta(\Theta)$, $u(\theta, a)$, $v(\theta, a)$.

- A sequence of states $(\theta_1, \ldots, \theta_t, \ldots)$ is drawn i.i.d. from p.
- At each stage t, player 1 is informed of θ_t, sends a message $m_t = \hat{\theta}_t$ to player 2 who chooses an action a_t. Messages and actions are perfectly observed.
- The average payoff is $\sum_t (1-\delta)\delta^{t-1} u_i(\theta_t, a_t)$.

Repeated sender receiver game

In any equilibrium of the repeated game, payoffs (x, y) should be individually rational:

$$x \geq \underline{u}(p) = \min_a u(p, a) = \min_a \sum_\theta p(\theta) u(\theta, a)$$

Player 2 can punish player 1 by choosing the worst action.

$$y \geq \underline{v}(p) = \max_a v(p, a)$$

Player 1 can punish player 2 by revealing no information.

Repeated sender receiver game

Given $\mu : \Theta \to \Delta(A)$ denote $U(\mu) = \sum_\theta p(\theta) \sum_a \mu(a \mid \theta) u(\theta, a)$

Given $Q : \Theta \to \Delta(\Theta)$ denote,
$pQ(\theta') = \sum_\theta p(\theta) Q(\theta' \mid \theta)$ and $Q\mu(a \mid \theta) = \sum_{\theta'} Q(\theta' \mid \theta)\mu(a \mid \theta')$.

μ is Incentive Compatible if:

$$pQ = Q \text{ and } U(Q\mu) \geq U(\mu) \Rightarrow Q\mu = Q$$

- Q represents a deviation of player 1 from truth telling. If $pQ = Q$ it is statistically undetectable by player 2.
- IC means that an undetectable deviation is not profitable.
- This is satisfied by any communication equilibrium.

Repeated sender receiver game

Theorem (Renault Solan Vieille 13, Renou Tomala 15)

If μ is incentive compatible and $V(\mu) \geq \underline{v}(p)$, then for all $\varepsilon > 0$, there exists δ_ε such that for all $\delta > \delta_\varepsilon$, in all equilibria of the repeated game,

$$\max_{\theta,a} \mathbb{E}\left|\sum_t (1-\delta)\delta^{t-1}1\{\theta_t = \theta, a_t = a\} - p(\theta)\mu(a \mid \theta)\right| \leq \varepsilon$$

- All equilibria of games long enough implement the mechanism μ.
- Player 2 tests the statistical distribution of reports of player 1 and punishes if they don't match the theoretical distributions. See Lehrer 90 and Gossner 95.
- The construction extends to ergodic (recurrent aperiodic) Markov chains. See Escobar and Toikka 2013.

Example

	a_1	a_2	a_3	
θ_1	$(3,3)$	$(1,2)$	$(0,0)$	$1/2$
θ_2	$(2,0)$	$(3,2)$	$(1,3)$	$1/2$

We want to implement $\mu(\theta_1) = a_1$, $\mu(\theta_2) = a_3$. This is individually rational.

If $(1/2, 1/2)Q = (1/2, 1/2)$, then Q is bi-stochastic with diagonal element $1 - \alpha$. Then,

$$U(Q\mu) = \frac{1}{2}[(1-\alpha)3 + \alpha 0] + \frac{1}{2}[(1-\alpha)1 + \alpha 2] = 2 - \alpha < 2 = U(\mu)$$

Outline

Persuasion

Some variations of the model are called Persuasion instead of cheap talk.

- Models of Certifiable communications: Message sets depend on the state. (Prove that you cannot play the piano). Forges and Koessler 2005, 2008.
- Bayesian Persuasion: The sender announces and fixes his strategy before knowing the state. Kamenica and Gentzkow 2011.

  Now called Information Design. The sender is not informed of the state but chooses an information structure or a statistical experiment.

Statistical experiments and splittings

Set of states Θ, prior $p \in \Delta(\Theta)$. A *statistical experiment* is $x : \Theta \to \Delta(S)$ with S a set of signals. Let,

- $p_s(\theta) = \mathbb{P}(\theta \mid s) = \frac{x(s|\theta)p(\theta)}{\sum_\zeta x(s|\zeta)p(\zeta)}$
- $\lambda_s = \mathbb{P}(s) = \sum_\zeta p(\zeta)x(s \mid \theta)$.

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).

Decision maker or receiver

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

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

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

Example. $\{\theta_0, \theta_1\}$, $\{a_0, a_1\}$, $v(\cdot, a_0) = 0$, $v(\theta_0, a_1) = -3$, $v(\theta_1, a_1) = 1$.

Information designer, expert, sender

- The sender chooses $x : \Theta \to \Delta(S)$ (prior to knowing the state!).
- The receiver gets signal s, chooses action a.
- Payoff $u(\theta, a)$ to the sender, $v(\theta, a)$ to the receiver.

Solving the game

The receiver chooses a TBR $\gamma(p)$. The program of the sender is

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

Pure persuasion: $u(a)$.

$$\max\{\sum_s \lambda_s u(\gamma(p_s)) : p = \sum_s \lambda_s p_s\}$$

Denote $U_\gamma(p) = \sum_\theta p(\theta) u(\theta, \gamma(p))$ then,

Kamenica-Gentzkow

$$\sup\{\sum_s \lambda_s U_\gamma(p_s) : p = \sum_s \lambda_s p_s\} = \text{Cav } U_\gamma(p)$$

Remarks

1. $\sharp S = \sharp\Theta$ messages are enough.
2. There might not be a maximum. No 0-equilibrium, but ε-equilibria.
3. There exists a TBR such that there is a maximum, thus a 0-equilibrium (favor the sender).
4. Cav $U_\gamma(p)$ might depend on γ.

Unless the problem is regular:

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

Then Cav $U_\gamma(p)$ does not depend on γ and is the value of all limits of ε-EQ payoffs (as $\varepsilon \to 0$).

References

- Aumann and Hart, Econometrica 2003. Long Cheap Talk.
- Aumann and Maschler, 1995. Repeated Games with Incomplete Information.
- Crawford and Sobel, Econometrica 1982. Strategic Information Transmission.
- Escobar and Toikka, Econometrica 2013. Efficiency in Games with Markovian Private Information.
- Forges, 1994, Essays in Game Theory: In Honor of Michael Maschler. Non-zero sum repeated games and information transmission.
- Forges and Koessler, 2005, Journal of Mathematical Economics. Communication equilibria with partially verfiable types.
- Forges and Koessler, 2008, Journal of Economic Theory. Long Persuasion Games.
- Gossner, 1995, International Journal of Game Theory. The Folk theorem for finitely repeated games with mixed strategies,
- Hart, Mathematics of Operations Research, 1985. Non zero sum two person repeated games with incomplete information.
- Kamenica and Gentzkow, American Economic Review, 2011. Bayesian Persuasion.
- Lehrer, 1990, International Journal of Game Theory. Nash equilibria of n-player repeated games with semi-standard information.
- Renault Solan Vieille, Journal of Economic Theory 2013. Dynamic Sender Receiver Games.

Thank You!