

Bayesian games, Repeated Games, Information Design and Persuasion

Tristan Tomala (HEC Paris)

“ReLaX” Workshop on Games, CMI, February 1 - February 4, 2021

Outline

- 1 Bayesian games
- 2 Zero-sum games
- 3 Sender Receiver games
- 4 Information design
- 5 Games of information design
- 6 Splitting games

Bayesian games

A Bayesian game, or game with incomplete information is given by:

- A finite set of players $i = 1, \dots, N$;
- A set of states $\Omega = \Theta \times \prod_i M_i$ with a prior probability distribution $p \in \Delta(\Omega)$;
- A set of actions A_i for each player i ;
- A payoff function $u_i : \Omega \times A \rightarrow \mathbb{R}$ for each player i .

- 1 A state $\omega = (\theta, m_1, \dots, m_N)$ is drawn from p , Player i is informed of m_i .
- 2 Players choose actions simultaneously.

The game at the second stage is possibly dynamic: actions represent pure strategies in the continuation (extensive form) game.

Bayes-Nash equilibrium

A Bayes-Nash equilibrium is a strategy profile $\sigma = (\sigma_1, \dots, \sigma_N)$ with $\sigma_i : M_i \rightarrow \Delta(A_i)$ such that:

For each player i , each message m_i with $p(m_i) > 0$, each action a_i ,

$$\mathbb{E}[u_i(\theta, m_i, m_{-i}, \sigma_i(m_i), \sigma_{-i}(m_{-i})) \mid m_i] \geq \mathbb{E}[u_i(\theta, m_i, m_{-i}, a_i, \sigma_{-i}(m_{-i})) \mid m_i]$$

This is an Nash equilibrium of the normal form game with payoff $\mathbb{E}[u_i(\theta, m, \sigma(m))]$.

This exists when:

- States and actions are finite;
- States are finite or countable, actions sets are compact metric, payoffs are continuous.

Outline

- ① Bayesian games
- ② Zero-sum games
- ③ Sender Receiver games
- ④ Information design
- ⑤ Games of information design
- ⑥ Splitting games

Zero-sum games

We consider two-player zero-sum games with independent types (for simplicity).

- There are two players, the finite state space is $K \times L$ with prior distribution $p(k)q(l)$.
- Action sets are A, B . Payoffs are $u_1 = -u_2 = g(k, l, a, b)$.

The game has a value if

$$\sup_{\sigma} \inf_{\tau} \mathbb{E}_{p,q}[g(k, l, \sigma, \tau)] = \inf_{\tau} \sup_{\sigma} \mathbb{E}_{p,q}[g(k, l, \sigma, \tau)] := V(p, q)$$

- If $|L| = 1$ or $|K| = 1$, this is a game with Lack of Information on one-side. (Otherwise, on both sides).

Lack on one-side $|L| = 1$

Examples: $K = \{k_0, k_1\}$, $p = (\frac{1}{2}, \frac{1}{2})$. Let G_k be the payoff matrix in state k .

$$G_{k_0} = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix} \quad G_{k_1} = \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix}$$

$$G_{k_0} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad G_{k_1} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$$

$$G_{k_0} = \begin{pmatrix} 4 & 0 & 2 \\ 4 & 0 & -2 \end{pmatrix} \quad G_{k_1} = \begin{pmatrix} 0 & 4 & -2 \\ 0 & 4 & 2 \end{pmatrix}$$

Concavity

Let $G(p)$ be the game with prior distribution p and let $V(p)$ be its value.

Theorem (Aumann Maschler)

$V(\cdot)$ is a concave function of p .

This follows from the following statement:

Suppose that for each q , player 1 guarantees $f(q)$ in $G(q)$ and let $p = \sum_m \lambda_m p_m$. Then, player 1 guarantees $\sum_m \lambda_m f(p_m)$ in $G(p)$.

The theorem follows with $f(q) = V(q)$:

$$V(p) \geq \sum_m \lambda_m V(p_m)$$

We prove now this statement.

Experiments and splittings

Let M be a finite set of messages. A statistical (Blackwell) experiment is $x : K \rightarrow \Delta(M)$. We have

$$\mathbb{P}(k, m) = p(k)x(m|k) = \lambda_m p_m(k)$$

where $\lambda_m = \mathbb{P}(m) = \sum_k p(k)x(m|k)$ and

$$p_m(k) = \mathbb{P}(k | m) = \frac{\mathbb{P}(k, m)}{\mathbb{P}(m)} = p(k)x(m|k)/\lambda_m$$

Also,

$$p(k) = \sum_m \mathbb{P}(k, m) = \sum_m \lambda_m p_m(k)$$

Splitting

A splitting of $p \in \Delta(K)$ is a distribution (random posterior) \tilde{p} such that $\mathbb{E}(\tilde{p}) = p$. With finite support, this is convex combination

$$p = \sum_m \lambda_m p_m$$

Any experiment induces a splitting.

Splitting lemma

Conversely, since

$$\mathbb{P}(k, m) = p(k)x(m|k) = \lambda_m p_m(k)$$

Splitting lemma (Aumann Maschler, 65)

Any splitting of p

$$p = \sum_m \lambda_m p_m$$

is induced by the experiment $x(m|k) = \lambda_m p_m(k)/p(k)$

Auxiliary game

- Consider the game $\tilde{G}(p)$ where player 1 can send a message m to player 2 before choosing actions. Let $\tilde{V}(p)$ be its value.
- If $p = \sum_m \lambda_m p_m$, then player 1 can send the message m and guarantee $f(p_m)$ afterwards. So $\tilde{V}(p) \geq \sum_m \lambda_m f(p_m)$.
- Consider now the game $\tilde{G}^*(p)$ where player 2 does not observe the message. Player 2 has less strategies in $\tilde{G}^*(p)$ than in $\tilde{G}(p)$. So, $\tilde{V}^*(p) \geq \tilde{V}(p)$.
- But then $\tilde{G}^*(p) \equiv G(p)$. So

$$V(p) \geq \tilde{V}^*(p) \geq \tilde{V}(p) \geq \sum_m \lambda_m f(p_m)$$

QED.

Concavification

Let $u(p)$ be the value of the game without information:

$u(p)$ is the value of the matrix game $D(p) = \sum_k p(k)G_k$.

Let $\text{Cav } u(p)$ be its concave closure:

$$\begin{aligned}\text{Cav } u(p) &= \inf\{f(p) : f \geq u, f \text{ concave}\} \\ &= \sup\left\{\sum_m \lambda_m u(p_m) : \sum_m \lambda_m p_m = p\right\}\end{aligned}$$

Lemma

$$V(p) \geq \text{Cav } u(p)$$

Repeated game

$G_n(p)$ is the n -stage repeated game with average payoff where:

- State k is drawn from p once and for all and known by player 1;
- Actions are chosen at each stage and are observed; Payoffs are not observed.

$V_n(p)$ is the value of $G_n(p)$. $V_\infty(p)$ is the (uniform) value of $G_\infty(p)$ (if it exists).

Theorem (Aumann Maschler)

- 1 $\text{Cav } u(p) \leq V_n(p) \leq \text{Cav } u(p) + \frac{cte}{\sqrt{n}}$
- 2 $V_\infty(p) = \text{Cav } u(p)$.

We know that $V_n(p), V_\infty(p) \geq \text{Cav } u(p)$.

Examples

$$G_{k_0} = \begin{pmatrix} 1 & 1 \\ 0 & 0 \end{pmatrix} \quad G_{k_1} = \begin{pmatrix} 0 & 0 \\ 1 & 1 \end{pmatrix}$$


$u(p) = \max\{p, 1 - p\}$ convex. Full revelation.

$$G_{k_0} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad G_{k_1} = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$$

$u(p) = p(1 - p)$ concave. No revelation.

Examples

$$G_{k_0} = \begin{pmatrix} 4 & 0 & 2 \\ 4 & 0 & -2 \end{pmatrix} \quad G_{k_1} = \begin{pmatrix} 0 & 4 & -2 \\ 0 & 4 & 2 \end{pmatrix}$$



Partial revelation

$$V_n(p) \leq \text{Cav } u(p) + \frac{cte}{\sqrt{n}}$$

- Fix a strategy σ of player 1 and a history of actions $h_n = (a_1, b_1, \dots, a_{n-1}, b_{n-1})$. Let $p_n(k) = \mathbb{P}(k|h_n)$.
- Interpret $\sigma(a_n = m|k, h_n)$ as an experiment $x(m|k)$.
- $\lambda_m = \sum_k p_n(k)x(m|k)$ is the probability of playing $a_n = m$.
- Let $\tau(h_n)$ be a best-response to λ in $D(p_n)$:

$$\sum_k p_n(k) G_k(\lambda, \tau(h_n)) \leq u(p_n)$$

We have

$$\begin{aligned} \mathbb{E}[G_k(a_n, b_n)|h_n] &\leq \sum_k p_n(k) G_k(\lambda, \tau(h_n)) + C \sum_k p_n(k) \sum_m |x(m|k) - \lambda_m| \\ &\leq u(p_n) + C \sum_m \lambda_m \sum_k |p_{n+1}(k) - p_n(k)| \end{aligned}$$

Then (with the Martingale property),

$$\mathbb{E}[\bar{G}_n] \leq \mathbb{E} \frac{1}{n} \sum_t u(p_t) + \frac{cte(p)}{\sqrt{n}} \leq \text{Cav } u(p) + \frac{cte(p)}{\sqrt{n}}$$

$$V_\infty(p) \leq \text{Cav } u(p)$$

We need to prove that player 2 guarantees $\text{Cav } u(p)$ in the infinite (limit) game.

- There exists a vector $\xi \in \mathbb{R}^K$ such that

$$\text{Cav } u(p) = \langle p, \xi \rangle \text{ and } \forall q, u(q) \leq \langle q, \xi \rangle$$

Given history h_n , define a vector $v \in \mathbb{R}^k$ with $v_k = (1/n) \sum_t G_k(a_t, b_t)$. Define $q \in \Delta(K)$ by

$$q_k = \frac{(v_k - \xi_k)_+}{\sum_j (v_j - \xi_j)_+}$$

- If q is well defined, player 2 plays an optimal strategy in $D(q)$.
- Player 2 plays arbitrarily otherwise.

$$V_\infty(p) \leq \text{Cav } u(p)$$

Let $v_n = (1/n) \sum_t G(a_t, b_t)$ and $\pi(v_n)$ be the projection on $C = \{w \leq \xi\}$.

By construction, $v_n - \pi(v_n) = \lambda q$ with $\lambda > 0$ and

$$\langle \mathbb{E}[G(a_{n+1}, b_{n+1}) \mid h_n], q \rangle \leq u(q) \leq \langle q, \xi \rangle = \langle q, \pi(v_n) \rangle$$

For any strategy σ of player 1, if $v_n \notin C$,

$$\mathbb{E}[\langle G(a_{n+1}, b_{n+1}) - \pi(v_n), v_n - \pi(v_n) \rangle \mid h_n] \leq 0$$

If a sequence of vectors is such that

$$\langle x_{n+1} - \pi(\bar{x}_n), \bar{x}_n - \pi(\bar{x}_n) \rangle \leq 0$$

then $d(\bar{x}_n, C) \rightarrow 0$. QED

Outline

- ① Bayesian games
- ② Zero-sum games
- ③ Sender Receiver games**
- ④ Information design
- ⑤ Games of information design
- ⑥ Splitting games

Sender Receiver Games

A Sender Receiver game is a two-player Bayesian game where:

- 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 transmitted imperfectly.

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.
- The only equilibrium is babbling (non-informative).

Equilibrium condition for player 2

- A behavior strategy for player 1 is $\sigma : \Theta \rightarrow \Delta(M)$ denoted $\sigma(m | \theta)$.
- A strategy σ of player 1 is formally an experiment and is equivalent to a splitting $p = \sum_m \lambda_m p_m$.
- A behavior strategy for player 2 is $\tau : M \rightarrow \Delta(A)$ denoted $\tau(a | m)$.
- Player 2 who receives the message m has belief p_m .

Player 2 with belief p over the states, chooses a mixed action from the set:

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

- 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 ,

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

- Notice that $\sigma(m | \theta) > 0 \Leftrightarrow p_m(\theta) > 0$. Player 1 is indifferent between all messages (and thus posteriors) induced with positive probability.
- So 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 (Aumann-Hart, Forges)

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

Compare with $\text{Cav } u(p)$. Convexify the non-revealing payoffs. Indifference condition for player 1.

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.

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 $\sigma(m_1 | \theta_1) = \frac{3}{4} = \sigma(m_2 | \theta_2)$.
- This equilibrium payoff dominates NR and CR.
- Player 1 obtains the maximal payoff 3.

Outline

- ① Bayesian games
- ② Zero-sum games
- ③ Sender Receiver games
- ④ **Information design**
- ⑤ Games of information design
- ⑥ Splitting games

Bayesian persuasion

Consider a sender-receiver environment $p \in \Delta(\Theta)$, $u(\theta, a)$, $v(\theta, a)$.

- Player 1 is an *information designer* if he can choose the information structure $x : \Theta \rightarrow \Delta(M)$ which delivers a message to player 2.
- Player 1 is no longer informed of the state, but chooses how information is disclosed to player 2.
- The sender receiver game is played à la Stackelberg:
 - Player 1 chooses the experiment x once and for all. It is publicly observed.
 - Nature draws θ from p and m from x .
 - Player observes x and chooses an action.
- Player 1 is able to commit on its randomizations: distributions of messages are observable and verifiable.

This model is called Bayesian Persuasion (Kamenica and Gentzkow, 2011).

Receiver


Player 2 with belief p over the states, chooses a mixed action from the set:

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

An equilibrium strategy of player 2 is a tie-breaking-rule (TBR): a selection $\gamma(p) \in Y(p)$.

Example

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



Solving the game

A strategy σ of player 1 is an experiment and is equivalent to a splitting $p = \sum_m \lambda_m p_m$.

Given that the receiver chooses TBR $\gamma(p)$, the program of the sender is

$$\max \left\{ \sum_m \lambda_m \sum_{\theta} p_m(\theta) u(\theta, \gamma(p_m)) : p = \sum_m \lambda_m p_m \right\}$$

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


Lemma (Kamenica-Gentzkow, 11)

The equilibrium payoff of player 1 is

$$\sup \left\{ \sum_m \lambda_m U_{\gamma}(p_m) : p = \sum_m \lambda_m p_m \right\} = \text{Cav } U_{\gamma}(p)$$

Example

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



Outline

- ① Bayesian games
- ② Zero-sum games
- ③ Sender Receiver games
- ④ Information design
- ⑤ Games of information design**
- ⑥ Splitting games

Games between information designers

Environment

- Finite multidimensional state space $\Theta = \Theta_1 \times \cdots \times \Theta_n$.
- Prior $p^0 \in \Delta(\Theta)$ (possibly correlated).
- Designer $i = 1, \dots, n$ chooses a statistical experiment $x_i : \Theta_i \rightarrow \Delta(M_i)$ (*information structure*), where $M_i = M_i^1 \times \cdots \times M_i^k$ is a finite set of message profiles.
- Each agent $j = 1, \dots, k$ chooses an action $a_j \in A_j$.
- Designer i has payoff $u_i(a, \theta_1, \dots, \theta_n)$, with $a \in A = A_1 \times \cdots \times A_k$.
- Agent j has payoff $v_j(a, \theta_1, \dots, \theta_n)$.

Timing of the $(n + k)$ -player Information Design Game G_M

Finite sets of messages are given.

- 1 Designers choose experiments (x_1, \dots, x_n) simultaneously.
- 2 Agents publicly observe (x_1, \dots, x_n) .

A k -player Bayesian subgame $G_M(x)$ is then played:

- 3 Nature draws the state $\theta = (\theta_1, \dots, \theta_n)$ according to $p^0 \in \Delta(\Theta)$, and a uniformly distributed sunspot $\omega \in [0, 1]$.
- 4 A profile of messages $m_i = (m_i^1, \dots, m_i^k)$ from each designer i is drawn with probability $x_i(m_i | \theta_i)$.
- 5 Agents publicly observe ω , each agent j privately observes the profile of messages $m^j \in M^j = \prod_{i \in N} M_i^j$ and chooses an action a_j .

Subgame Perfect Equilibrium

For each profile of experiments of the designers $x \in \prod_i \Delta(M_i)^{\Theta_i}$, the Bayesian game $G_M(x)$ is a finite game, so its set of Nash equilibrium outcomes

$$\mathcal{E}_M(x) \subseteq \{y : M \rightarrow \Delta(A)\} = \Delta(A)^M$$

is non-empty.

Since there is public correlation, $\mathcal{E}_M(x)$ is convex (set of public correlated equilibria of $G_M(x)$).

Theorem (Existence, Koessler et al. 2021)

For every profile of finite message sets, the $(n + k)$ -player information design G_M game admits a subgame perfect equilibrium (x, y) , $y \in \mathcal{E}_M(x)$.

Relies on Simon and Zame (1990).

Outline

- ① Bayesian games
- ② Zero-sum games
- ③ Sender Receiver games
- ④ Information design
- ⑤ Games of information design
- ⑥ Splitting games

Splitting games

- There are two finite sets K, L and initial “states” $p^0 \in P = \Delta(K)$, $q^0 \in Q = \Delta(L)$. At every stage n , in states (p^n, q^n) ,
- Player 1 chooses a splitting $s \in S(p^n) = \left\{ s \in \Delta(P) : E_s(p^n) = p^n \right\}$,
- Player 2 chooses a splitting $t \in S(q^n) = \left\{ t \in \Delta(Q) : E_t(q^n) = q^n \right\}$,
- the next states (p^{n+1}, q^{n+1}) are drawn from $s \otimes t$.
- The stage payoff is $u(p^n, q^n)$ with $u : P \times Q \rightarrow \mathbb{R}$ continuous.

Discounted game: $\sum \lambda(1 - \lambda)^{n-1} u(p^n, q^n)$.

Long game: $u(p^\infty, q^\infty)$.

Splitting games

- Splitting games are games of information design: Player 1 designs information about k , Player 2 about l .

At the end of the day, a decision maker takes an optimal decision $z(p, q)$ as a function of his posterior beliefs (p, q) .

$u(p, q)$ is the payoff of Player 1 induced by this decision.

- These are special stochastic games with “acyclic” transitions.

Back to repeated games with incomplete information

- Splitting games were introduced as proxy for repeated games with incomplete information on both sides where:
 - Player 1 knows k , Player 2 knows l ,
 - they repeatedly choose actions i, j which are observed,
 - payoffs G_{ij}^{kl} are unobserved.
 - Let,

$$u(p, q) = \text{Val} \left(\sum_{k,l} p^k q^l G_{ij}^{kl} \right)$$

Concave–Convex functions

What is the correct generalization of $\text{Cav } u$?

- $\text{Cav}_p \text{Vex}_q u(p, q)$. This is the MaxMin of the undiscounted repeated game with incomplete information on both sides.
- $\text{Vex}_q \text{Cav}_p u(p, q)$. This is the MinMax of the undiscounted repeated game with incomplete information on both sides.
- In general

$$\text{Cav}_p \text{Vex}_q u(p, q) < \text{Vex}_q \text{Cav}_p u(p, q)$$

Therefore the undiscounted repeated game has no value.

Mertens-Zamir

Yet,

Theorem (Mertens-Zamir, 1971)

In the repeated game with incomplete information on both sides, $v(p, q) := \lim v_\lambda(p, q)$ exists and is the unique continuous function such that

$$v(p, q) = \text{Cav} \min(u, v)(p, q) = \text{Vex} \max(u, v)(p, q)$$

- v is concave–convex and $\text{Cav}_p \text{Vex}_q u(p, q) < v(p, q) < \text{Vex}_q \text{Cav}_p u(p, q)$.
- Sketch of proof: As long as $u(p, q) \geq v(p, q)$, Player 1 stays silent.
As soon as $u(p, q) < v(p, q)$, Player 1 Plays optimally in the discounted game.
Player 1 thus gets a convex combination of $\max(u, v)(p_t, q_t)$ which is $\geq \text{Vex} \max(u, v)(p, q)$.
- MZ prove existence and uniqueness of v .

Splitting game

In the splitting game with stage payoff $u(p, q)$, then

Theorem

$v = \lim v_\lambda$ (Laraki 2001). It is also the value of the undiscounted game (Oliu Barton 2017) and of the long game (Koessler et al. 2021).

One can show that v is the unique function such that

- If $u(p, q) < v(p, q)$, there exists $s \in S(p)$ such that $v(p, q) = v(s, q) \geq u(s, q)$
- If $u(p, q) > v(p, q)$, there exists $t \in S(q)$ such that $v(p, q) = v(p, t) \leq u(p, t)$

This defines an optimal strategy for each player.

Approximation of the MZ function

Assume $|K| = |L| = 1$.

- Consider grids $P^n = \{0, \frac{1}{n}, \dots, \frac{i}{n}, 1\}$, $Q^m = \{0, \frac{1}{m}, \dots, \frac{j}{m}, 1\}$ on $[0, 1] = \Delta(K) = \Delta(L)$.
- Let $U = (u_{ij})$ be the matrix $u_{i,j} = u(\frac{i}{n}, \frac{j}{m})$.

Lemma (Koessler et al. 2021)

There exists a unique matrix V which is column-concave, row-convex such that $v_{00} = u_{00}$, $v_{01} = u_{01}$, $v_{10} = u_{10}$, $v_{11} = u_{11}$ and,

$$\text{if } v_{i,j} > u_{i,j}, \text{ then } v_{i,j} = \frac{1}{2}(v_{i-1,j} + v_{i+1,j}),$$

$$\text{if } v_{i,j} < u_{i,j}, \text{ then } v_{i,j} = \frac{1}{2}(v_{i,j-1} + v_{i,j+1}).$$

- This matrix can be found by solving linear systems.
- This induces a piecewise linear function which uniformly approximates the function u .

Thank you !

REFERENCES

- R. J. Aumann and S. Hart. Long cheap talk. *Econometrica*, 71(6):1619–1660, 2003.
- R. J. Aumann, M. B. Maschler, and R. Stearns. *Repeated Games of Incomplete Information*. MIT Press, Cambridge, Massachusetts, 1995.
- D. Bergemann and S. Morris. Information design, bayesian persuasion, and bayes correlated equilibrium. *The American Economic Review*, 106(5): 586–591, 2016.
- Françoise Forges. Correlated equilibrium in games with incomplete information revisited. *Theory and decision*, 61(4):329–344, 2006.
- Françoise Forges. An approach to communication equilibria. *Econometrica*, 54(6):1375–1385, 1986.
- Françoise Forges. Five legitimate definitions of correlated equilibrium in games with incomplete information. *Theory and Decision*, 35:277–310, 1993.
- Françoise Forges. Games with incomplete information: from repetition to cheap talk and persuasion. *mimeo*, 2019.
- E. Kamenica and M. Gentzkow. Bayesian Persuasion. *American Economic Review*, 101(6):2590–2615, 2011.
- Frédéric Koessler, Marie Laclau, Jérôme Renault, and Tristan Tomala. Long information design. *mimeo*, 2021a.

Frédéric Koessler, Marie Laclau, Jérôme Renault, and Tristan Tomala.

Splitting games over finite sets. *mimeo*, 2021b.

Frédéric Koessler, Marie Laclau, and Tristan Tomala. Interactive information design. *forthcoming in Mathematics of Operations Research*, 2021c.

J-F. Mertens and S. Zamir. The value of two-person zero-sum repeated games with lack of information on both sides. *International Journal of Game Theory*, 1(1):39–64, 1971.

J-F. Mertens, S. Sorin, and S. Zamir. *Repeated games*. Cambridge University Press, 2015.

Roger B Myerson. Bayesian equilibrium and incentive-compatibility: An introduction. *Social goals and social organization: Essays in memory of Elisha Pazner*, pages 229–260, 1985.

M. Orlin-Barton. The splitting game: Uniform value and optimal strategies. *Dynamic Games and Applications*, pages 1–23, 2017.

L. K. Simon and W. R. Zame. Discontinuous games and endogenous sharing rules. *Econometrica*, pages 861–872, 1990.

S. Sorin. *A First Course on Zero-Sum Repeated Games*. Springer, 2002.

I. Taneva. Information design. *mimeo*, 2016.