Lecture 2: Bayesian persuasion

Fedor Sandomirskiy

May 4, 2020

Technion, Haifa & Higher School of Economics, St.Petersburg
e-mail: fedor.sandomirskiy@gmail.com
homepage: https://www.fedors.info/

Outline:

- The model of Bayesian persuasion
- Geometric approach to persuasion: Splitting lemma and $\mathrm{Cav}\,[U]$-theorem
- Action-recommendation approach: revelation-principle

The model of Bayesian persuasion

Bayesian persuasion and Information Design

Information Design

How to induce the desired behavior of a decision-maker by changing the information available to him?

- A young field. The origin:

Bayesian persuasion
E Kamenica, M Gentzkow - American Economic Review, 2011 - aeaweb.org
… work that identifies which sequences of distributions of posteriors are consistent with **Bayesian** rationality … we need only ask how Eт □ v (μ) varies over the space of **Bayes**-plausible distributions … COROLLARY 1: sender benefits from **persuasion** if and only if there exists a **Bayes** …
Cited by 1295 Related articles All 31 versions

Bayesian persuasion and Information Design

Information Design

How to induce the desired behavior of a decision-maker by changing the information available to him?

- A young field. The origin:

Bayesian persuasion
E Kamenica, M Gentzkow - American Economic Review, 2011 - aeaweb.org
… work that identifies which sequences of distributions of posteriors are consistent with **Bayesian** rationality … we need only ask how Eτ ☐ v (μ) varies over the space of **Bayes**-plausible distributions … COROLLARY 1: sender benefits from **persuasion** if and only if there exists a **Bayes** …
☆ 99 Cited by 1295 Related articles All 31 versions »

- Bayesian persuasion = the simplest model of Information Design: 2 agents:
  - Receiver: a decision maker who has no access to payoff-relevant information
  - Sender: has information, cares about the action of Receiver, can send him a signal

Bayesian persuasion and Information Design

Information Design

How to induce the desired behavior of a decision-maker by changing the information available to him?

- A young field. The origin:

  Bayesian persuasion
  E Kamenica, M Gentzkow - American Economic Review, 2011 - aeaweb.org
  … work that identifies which sequences of distributions of posteriors are consistent with **Bayesian** rationality … we need only ask how Eτ ☐ v (μ) varies over the space of **Bayes**-plausible distributions … COROLLARY 1: sender benefits from **persuasion** if and only if there exists a **Bayes** …
  Cited by 1295 Related articles All 31 versions

- Bayesian persuasion = the simplest model of Information Design: 2 agents:
  - Receiver: a decision maker who has no access to payoff-relevant information
  - Sender: has information, cares about the action of Receiver, can send him a signal
- Bayesian persuasion $\simeq$ Information Design

Bayesian persuasion and Information Design

Information Design

How to induce the desired behavior of a decision-maker by changing the information available to him?

- A young field. The origin:

> **Bayesian persuasion**
> E Kamenica, M Gentzkow - American Economic Review, 2011 - aeaweb.org
> … work that identifies which sequences of distributions of posteriors are consistent with **Bayesian** rationality … we need only ask how Eт □ v (μ) varies over the space of **Bayes**-plausible distributions … COROLLARY 1: sender benefits from **persuasion** if and only if there exists a **Bayes** …
> ☆ 99 Cited by 1295 Related articles All 31 versions »

- Bayesian persuasion $=$ the simplest model of Information Design: 2 agents:
  - Receiver: a decision maker who has no access to payoff-relevant information
  - Sender: has information, cares about the action of Receiver, can send him a signal
- Bayesian persuasion $\simeq$ Information Design
- Popularity: simplicity, explicit solutions, many applications

Bayesian persuasion and Information Design

Information Design

How to induce the desired behavior of a decision-maker by changing the information available to him?

- A young field. The origin:

Bayesian persuasion
E Kamenica, M Gentzkow - American Economic Review, 2011 - aeaweb.org
When is it possible for one person to persuade another to change her action? We consider a symmetric information model where a sender chooses a signal to reveal to a receiver, who then takes a noncontractible action that affects the welfare of both players. We derive ...
Cited by 949 Related articles All 38 versions

- Bayesian persuasion = the simplest model of Information Design: 2 agents:
  - Receiver: a decision maker who has no access to payoff-relevant information
  - Sender: has information, cares about the action of Receiver, can send him a signal
- Bayesian persuasion $\simeq$ Information Design
- Popularity: simplicity, explicit solutions, many applications

A toy example of Bayesian persuasion

A court problem

- 75% of defendants are innocent ($\theta = 0$), 25% are guilty ($\theta = 1$)
- Prosecutor (P) observes θ, Judge (J) does not
- J has two actions: to acquit ($a = 0$) or to convict ($a = 1$)
- P's utility $u_P(a, \theta) = a$ (always wants to convict)
- J's utility $u_J(a, \theta) = \mathbb{1}_{a=\theta}$ (cares about justice)

A toy example of Bayesian persuasion

A court problem

- 75% of defendants are innocent ($\theta = 0$), 25% are guilty ($\theta = 1$)
- Prosecutor (P) observes θ, Judge (J) does not
- J has two actions: to acquit ($a = 0$) or to convict ($a = 1$)
- P's utility $u_P(a, \theta) = a$ (always wants to convict)
- J's utility $u_J(a, \theta) = \mathbb{1}_{a=\theta}$ (cares about justice)

What should P do?

- Reveal no information $\Longrightarrow (a^* = 0) \Longrightarrow (u_P = 0)$
- Reveal $\theta \Longrightarrow (a^* = \theta) \Longrightarrow (u_P = \frac{1}{4})$
- Send a noisy signal $\in \{\text{innocent}, \text{maybe guilty}\}$ with some θ-dependent probabilities $\pi \Rightarrow$ the optimal payoff $u_P = \frac{1}{2}$.

A toy example of Bayesian persuasion

A court problem

- 75% of defendants are innocent ($\theta = 0$), 25% are guilty ($\theta = 1$)
- Prosecutor (P) observes θ, Judge (J) does not
- J has two actions: to acquit ($a = 0$) or to convict ($a = 1$)
- P's utility $u_P(a, \theta) = a$ (always wants to convict)
- J's utility $u_J(a, \theta) = \mathbb{1}_{a=\theta}$ (cares about justice)

What should P do?

- Reveal no information $\implies (a^* = 0) \implies (u_P = 0)$
- Reveal $\theta \implies (a^* = \theta) \implies (u_P = \frac{1}{4})$
- Send a noisy signal $\in \{\text{innocent}, \text{maybe guilty}\}$ with some θ-dependent probabilities $\pi \Rightarrow$ the optimal payoff $u_P = \frac{1}{2}$.

A toy example of Bayesian persuasion

A court problem

- 75% of defendants are innocent ($\theta = 0$), 25% are guilty ($\theta = 1$)
- Prosecutor (P) observes θ, Judge (J) does not
- J has two actions: to acquit ($a = 0$) or to convict ($a = 1$)
- P's utility $u_P(a, \theta) = a$ (always wants to convict)
- J's utility $u_J(a, \theta) = \mathbb{1}_{a=\theta}$ (cares about justice)

What should P do?

- Reveal no information $\implies (a^* = 0) \implies (u_P = 0)$
- Reveal $\theta \implies (a^* = \theta) \implies (u_P = \frac{1}{4})$
- Send a noisy signal $\in \{\text{innocent}, \text{maybe guilty}\}$ with some θ-dependent probabilities $\pi \Rightarrow$ the optimal payoff $u_P = \frac{1}{2}$.

A toy example of Bayesian persuasion

A court problem

- 75% of defendants are innocent ($\theta = 0$), 25% are guilty ($\theta = 1$)
- Prosecutor (P) observes θ, Judge (J) does not
- J has two actions: to acquit ($a = 0$) or to convict ($a = 1$)
- P's utility $u_P(a, \theta) = a$ (always wants to convict)
- J's utility $u_J(a, \theta) = \mathbb{1}_{a=\theta}$ (cares about justice)

What should P do?

- Reveal no information $\implies (a^* = 0) \implies (u_P = 0)$
- Reveal $\theta \implies (a^* = \theta) \implies (u_P = \frac{1}{4})$
- Send a noisy signal $\in \{\text{innocent}, \text{maybe guilty}\}$ with some θ-dependent probabilities $\pi \Rightarrow$ the optimal payoff $u_P = \frac{1}{2}$.

A toy example of Bayesian persuasion

A court problem

- 75% of defendants are innocent ($\theta = 0$), 25% are guilty ($\theta = 1$)
- Prosecutor (P) observes θ, Judge (J) does not
- J has two actions: to acquit ($a = 0$) or to convict ($a = 1$)
- P's utility $u_P(a, \theta) = a$ (always wants to convict)
- J's utility $u_J(a, \theta) = \mathbb{1}_{a=\theta}$ (cares about justice)

What should P do?

- Reveal no information $\Longrightarrow (a^* = 0) \Longrightarrow (u_P = 0)$
- Reveal $\theta \Longrightarrow (a^* = \theta) \Longrightarrow (u_P = \frac{1}{4})$
- Send a noisy signal $\in \{\text{innocent}, \text{maybe guilty}\}$ with some θ-dependent probabilities $\pi \Rightarrow$ the optimal payoff $u_P = \frac{1}{2}$.

Remark: Communication is possible because:

- Non-zero-sum: sometimes P and J want the same (convict guilty).
- J knows the information structure π.

A toy example of Bayesian persuasion

A court problem

- 75% of defendants are innocent ($\theta = 0$), 25% are guilty ($\theta = 1$)
- Prosecutor (P) observes θ, Judge (J) does not
- J has two actions: to acquit ($a = 0$) or to convict ($a = 1$)
- P's utility $u_P(a, \theta) = a$ (always wants to convict)
- J's utility $u_J(a, \theta) = \mathbb{1}_{a=\theta}$ (cares about justice)

What should P do?

- Reveal no information $\Longrightarrow (a^* = 0) \Longrightarrow (u_P = 0)$
- Reveal $\theta \Longrightarrow (a^* = \theta) \Longrightarrow (u_P = \frac{1}{4})$
- Send a noisy signal $\in \{\text{innocent}, \text{maybe guilty}\}$ with some θ-dependent probabilities $\pi \Rightarrow$ the optimal payoff $u_P = \frac{1}{2}$.

Remark: Communication is possible because:

- Non-zero-sum: sometimes P and J want the same (convict guilty).
- J knows the information structure π.

A toy example of Bayesian persuasion

A court problem

- 75% of defendants are innocent ($\theta = 0$), 25% are guilty ($\theta = 1$)
- Prosecutor (P) observes θ, Judge (J) does not
- J has two actions: to acquit ($a = 0$) or to convict ($a = 1$)
- P's utility $u_P(a, \theta) = a$ (always wants to convict)
- J's utility $u_J(a, \theta) = \mathbb{1}_{a=\theta}$ (cares about justice)

What should P do?

- Reveal no information $\Longrightarrow (a^* = 0) \Longrightarrow (u_P = 0)$
- Reveal $\theta \Longrightarrow (a^* = \theta) \Longrightarrow (u_P = \frac{1}{4})$
- Send a noisy signal $\in \{\text{innocent}, \text{maybe guilty}\}$ with some θ-dependent probabilities $\pi \Rightarrow$ the optimal payoff $u_P = \frac{1}{2}$.

Remark: Communication is possible because:

- Non-zero-sum: sometimes P and J want the same (convict guilty).
- J knows the <u>information structure</u> π.
  - P announces π before observing θ and cannot change it after (P has the <u>commitment power</u>).

Some other interpretations/applications:

- Employers and universities: $\theta =$quality of a student (good/bad), U wants a good placement for any student, E wants good candidates.
  - Explains coarse grading in schools, universities, and industries:[1]
    "When recruiters call me up and ask me for the three best people, I tell them, "No! I will give you the names of the six best."

    Robert J. Gordon, Econ. dept., Northwestern

[1] Ostrovsky, Schwarz (2010) Information disclosure and unraveling in matching markets. AER

Some other interpretations/applications:

- Employers and universities: $\theta =$quality of a student (good/bad), U wants a good placement for any student, E wants good candidates.
  - Explains coarse grading in schools, universities, and industries:[1]

    "When recruiters call me up and ask me for the three best people, I tell them, "No! I will give you the names of the six best."

    Robert J. Gordon, Econ. dept., Northwestern
- Buyers and Sellers : $\theta =$ quality of the product (good/bad), S wants to sell any product, B wants a good product.
  - Explains why you cannot order the apts by rating or price on AirBNB[2]

---

[1] Ostrovsky, Schwarz (2010) Information disclosure and unraveling in matching markets. AER

[2] Romanyuk, Smolin (2019) Cream skimming and information design in matching markets. AEJ

Some other interpretations/applications:

- Employers and universities: $\theta =$quality of a student (good/bad), U wants a good placement for any student, E wants good candidates.
  - Explains coarse grading in schools, universities, and industries:[1]
    "When recruiters call me up and ask me for the three best people, I tell them, "No! I will give you the names of the six best."

    Robert J. Gordon, Econ. dept., Northwestern
- Buyers and Sellers : $\theta =$ quality of the product (good/bad), S wants to sell any product, B wants a good product.
  - Explains why you cannot order the apts by rating or price on AirBNB[2]
- Police & drivers: $\theta =$ whether the region is patrolled (yes/no). P wants D to obey the speed limit, D wants to obey only if the region is patrolled.

[1] Ostrovsky, Schwarz (2010) Information disclosure and unraveling in matching markets. AER

[2] Romanyuk, Smolin (2019) Cream skimming and information design in matching markets. AEJ

Bayesian persuasion: the model

The model

- A random state $\theta \in \Theta$, $\theta \sim p \in \Delta(\Theta)$
- Sender (S)
  - selects an information structure $(M,\ \pi : \Theta \to \Delta(M))$
  - observes θ and sends a message (signal) $m \in M$ with distribution π_θ
- Receiver (R) knows (M, π) and takes an action $a \in A$ after getting m
- Payoffs $u_R(a, \theta)$ and $u_S(a, \theta)$

Bayesian persuasion: the model

The model

- A random state $\theta \in \Theta$, $\theta \sim p \in \Delta(\Theta)$
- Sender (S)
  - selects an information structure $(M,\ \pi : \Theta \to \Delta(M))$
  - observes θ and sends a message (signal) $m \in M$ with distribution π_θ
- Receiver (R) knows (M, π) and takes an action $a \in A$ after getting m
- Payoffs $u_R(a, \theta)$ and $u_S(a, \theta)$

R's problem: play the optimal reply to (M, π) and received signal m

$$a^*(m) \in \arg\max_{a \in A} \mathbb{E}[u_R(a, \theta) \mid m].$$

Standard assumption: ties are broken in favor of S.

Bayesian persuasion: the model

The model

- A random state $\theta \in \Theta$, $\theta \sim p \in \Delta(\Theta)$
- Sender (S)
  - selects an information structure $(M,\ \pi : \Theta \to \Delta(M))$
  - observes θ and sends a message (signal) $m \in M$ with distribution π_θ
- Receiver (R) knows (M, π) and takes an action $a \in A$ after getting m
- Payoffs $u_R(a, \theta)$ and $u_S(a, \theta)$

R's problem: play the optimal reply to (M, π) and received signal m

$$a^*(m) \in \arg\max_{a \in A} \mathbb{E}[u_R(a, \theta) \mid m].$$

Standard assumption: ties are broken in favor of S.

S's problem: maximize

$$\mathbb{E}\big[u_S(a^*(m), \theta)\big] \quad \text{over information structures } (M, \pi).$$

Bayesian persuasion: the model

The model

- A random state $\theta \in \Theta$, $\theta \sim p \in \Delta(\Theta)$
- Sender (S)
  - selects an information structure $(M,\ \pi : \Theta \to \Delta(M))$
  - observes θ and sends a message (signal) $m \in M$ with distribution π_θ
- Receiver (R) knows (M, π) and takes an action $a \in A$ after getting m
- Payoffs $u_R(a, \theta)$ and $u_S(a, \theta)$

R's problem: play the optimal reply to (M, π) and received signal m

$$a^*(m) \in \arg\max_{a \in A} \mathbb{E}[u_R(a, \theta) \mid m].$$

Standard assumption: ties are broken in favor of S.

S's problem: maximize

$$\mathbb{E}\big[u_S(a^*(m), \theta)\big] \quad \text{over information structures } (M, \pi).$$

Remark: $a^*(m)$ is computed $\Rightarrow$ 1-agent problem

Geometric approach to persuasion: splitting lemma and $\mathrm{Cav}\,[U]$-theorem

Persuasion as inducing posterior beliefs

Simplifying assumption: binary state $\theta \in \{0, 1\}$.
Prior $p = \mathbb{P}(\theta = 1)$, posterior $p' = p'(m) = \mathbb{P}(\theta = 1 \mid m)$.

Persuasion as inducing posterior beliefs

Simplifying assumption: binary state $\theta \in \{0, 1\}$.
Prior $p = \mathbb{P}(\theta = 1)$, posterior $p' = p'(m) = \mathbb{P}(\theta = 1 \mid m)$.

R's problem again: maximize over $a \in A$

$$\mathbb{E}\big[u_R(a,\theta) \mid m\big] = \mathbb{E}\big[\mathbb{1}_{\{\theta=0\}} u_R(a,0) + \mathbb{1}_{\{\theta=1\}} u_R(a,1) \mid m\big] =$$

$$= (1-p')u_R(a,0) + p' \cdot u_R(a,1) \implies a^* = a^*(p').$$

Persuasion as inducing posterior beliefs

Simplifying assumption: binary state $\theta \in \{0,1\}$.
Prior $p = \mathbb{P}(\theta = 1)$, posterior $p' = p'(m) = \mathbb{P}(\theta = 1 \mid m)$.

R's problem again: maximize over $a \in A$

$$\mathbb{E}\big[u_R(a,\theta) \mid m\big] = \mathbb{E}\big[\mathbb{1}_{\{\theta=0\}} u_R(a,0) + \mathbb{1}_{\{\theta=1\}} u_R(a,1) \mid m\big] =$$

$$= (1-p')u_R(a,0) + p' \cdot u_R(a,1) \quad \Longrightarrow \quad a^* = a^*(p').$$

S's payoff= $\mathbb{E}\big[u_S(a^*(p'),\theta)\big] = \mathbb{E}\Big[\mathbb{E}\big[u_S(a^*(p'),\theta) \mid m\big]\Big] =$

$$= \mathbb{E}\Big[\mathbb{E}\big[\mathbb{1}_{\{\theta=0\}} u_S(a^*(p'),0) + \mathbb{1}_{\{\theta=1\}} u_S(a^*(p'),1) \mid m\big]\Big] =$$

$$= \mathbb{E}\Big[(1-p')u_S(a^*(p'),0) + p' \cdot u_S(a^*(p'),1)\Big] = \mathbb{E}\Big[U_S(p')\Big]$$

Persuasion as inducing posterior beliefs

S's payoff=$\mathbb{E}\big[u_S(a^*(p'),\theta)\big] = \mathbb{E}\Big[\mathbb{E}\big[u_S(a^*(p'),\theta) \mid m\big]\Big] =$

$$= \mathbb{E}\Big[\mathbb{E}\big[\mathbb{1}_{\{\theta=0\}} u_S(a^*(p'),0) + \mathbb{1}_{\{\theta=1\}} u_S(a^*(p'),1) \mid m\big]\Big] =$$

$$= \mathbb{E}\Big[(1-p')u_S(a^*(p'),0) + p' \cdot u_S(a^*(p'),1)\Big] = \mathbb{E}\Big[U_S(p')\Big]$$

Notation: $\mu_{(M,\pi)} \in \Delta([0,1])$ is the distribution of $p'(m)$ induced by (M,π).

Persuasion as inducing posterior beliefs

S's payoff=$\mathbb{E}\left[u_S(a^*(p'),\theta)\right] = \mathbb{E}\Big[\mathbb{E}\left[u_S(a^*(p'),\theta) \mid m\right]\Big] =$

$$= \mathbb{E}\Big[\mathbb{E}\left[\mathbb{1}_{\{\theta=0\}} u_S(a^*(p'),0) + \mathbb{1}_{\{\theta=1\}} u_S(a^*(p'),1) \mid m\right]\Big] =$$

$$= \mathbb{E}\Big[(1-p')u_S(a^*(p'),0) + p' \cdot u_S(a^*(p'),1)\Big] = \mathbb{E}\Big[U_S(p')\Big]$$

Notation: $\mu_{(M,\pi)} \in \Delta([0,1])$ is the distribution of $p'(m)$ induced by (M,π).

Conclusion

- S cares only about $\mu_{(M,\pi)}$:

$$\mathbb{E}\Big[U_S(p')\Big] = \int_{[0,1]} U_S(x)\, d\mu_{(M,\pi)}(x).$$

If $\mu_{(M,\pi)} = \mu_{(\tilde{M},\tilde{\pi})}$, (M,π) and $(\tilde{M},\tilde{\pi})$ are payoff-equivalent.

Persuasion as inducing posterior beliefs

S's payoff=$\mathbb{E}\Big[u_S(a^*(p'),\theta)\Big] = \mathbb{E}\Big[\mathbb{E}\big[u_S(a^*(p'),\theta) \mid m\big]\Big] =$

$$= \mathbb{E}\Big[\mathbb{E}\big[\mathbb{1}_{\{\theta=0\}} u_S(a^*(p'),0) + \mathbb{1}_{\{\theta=1\}} u_S(a^*(p'),1) \mid m\big]\Big] =$$

$$= \mathbb{E}\Big[(1-p')u_S(a^*(p'),0) + p' \cdot u_S(a^*(p'),1)\Big] = \mathbb{E}\Big[U_S(p')\Big]$$

Notation: $\mu_{(M,\pi)} \in \Delta([0,1])$ is the distribution of $p'(m)$ induced by (M,π).

Conclusion

- S cares only about $\mu_{(M,\pi)}$:

$$\mathbb{E}\Big[U_S(p')\Big] = \int_{[0,1]} U_S(x)\, d\mu_{(M,\pi)}(x).$$

  If $\mu_{(M,\pi)} = \mu_{(\tilde{M},\tilde{\pi})}$, (M,π) and $(\tilde{M},\tilde{\pi})$ are payoff-equivalent.

- Instead of maximizing over (M,π), it is enough to maximize over

$$\mathcal{D}(p) = \Big\{\mu \in \Delta[0,1] \;:\; \mu = \mu_{(M,\pi)} \text{ for some } (M,\pi) \text{ with prior } p\Big\}.$$

The splitting lemma

$$\mathcal{D}(p) \subset \left\{ \mu \in \Delta([0,1]) \;:\; \int_{[0,1]} x \, d\mu(x) = p \right\}$$

Why?

The splitting lemma

$$\mathcal{D}(p) \subset \left\{ \mu \in \Delta([0,1]) \;:\; \int_{[0,1]} x \, d\mu(x) = p \right\}$$

Why? By the martingale property $\mathbb{E}[p'] = p$ (aka Bayesian plausibility).

The splitting lemma

$$\mathcal{D}(p) \subset \left\{ \mu \in \Delta([0,1]) \, : \, \int_{[0,1]} x \, d\mu(x) = p \right\}$$

Why? By the martingale property $\mathbb{E}[p'] = p$ (aka Bayesian plausibility).

The splitting lemma (Aumann, Maschler (1960ies) / folk)

These two sets are equal: for any $\mu \in \Delta([0,1])$ with $\int x \, d\mu(x) = p$ there exists (M, π) s.t. $p'(m) \sim \mu$.
One can take $M = \operatorname{supp} \mu \subset [0,1]$.

The splitting lemma (Aumann, Maschler (1960ies) / folk)

These two sets are equal: for any $\mu \in \Delta([0,1])$ with $\int x \, d\mu(x) = p$ there exists (M, π) s.t. $p'(m) \sim \mu$.
One can take $M = \operatorname{supp} \mu \subset [0,1]$.

Proof for discrete μ via belief-recommendation:

- μ : point x_k has mass μ_k, $\sum \mu_k = 1$, $\sum x_k \cdot \mu_k = p$.
- define
  $\pi_{\theta=1}$: point x_k has mass $\frac{x_k}{p} \mu_k$
  $\pi_{\theta=0}$: point x_k has mass $\frac{1-x_k}{1-p} \mu_k$
- sample $m \sim \pi_\theta$ conditional on θ
- unconditionally $m \sim \mu$: $\mathbb{P}(m = x_k) = (1-p)\frac{1-x_k}{1-p}\mu_k + p\frac{x_k}{p}\mu_k = \mu_k$.
- $p' = \mathbb{P}(\theta = 1 \mid m = x_k) = \frac{\mathbb{P}(m=x_k|\theta=1)\mathbb{P}(\theta=1)}{\mathbb{P}(m=x_k)} = \frac{x_k}{p}\mu_k \cdot p \cdot \frac{1}{\mu_k} = x_k = m$

$\square$

The splitting lemma

The splitting lemma (Aumann, Maschler (1960ies) / folk)

These two sets are equal: for any $\mu \in \Delta([0,1])$ with $\int x\, d\mu(x) = p$ there exists (M, π) s.t. $p'(m) \sim \mu$.
One can take $M = \operatorname{supp} \mu \subset [0,1]$.

Proof for discrete μ via belief-recommendation:

- μ : point x_k has mass μ_k, $\sum \mu_k = 1$, $\sum x_k \cdot \mu_k = p$.
- define
  $\pi_{\theta=1}$: point x_k has mass $\frac{x_k}{p} \mu_k$
  $\pi_{\theta=0}$: point x_k has mass $\frac{1-x_k}{1-p} \mu_k$
- sample $m \sim \pi_\theta$ conditional on θ
- unconditionally $m \sim \mu$: $\mathbb{P}(m = x_k) = (1-p)\frac{1-x_k}{1-p}\mu_k + p\frac{x_k}{p}\mu_k = \mu_k$.
- $p' = \mathbb{P}(\theta = 1 \mid m = x_k) = \frac{\mathbb{P}(m=x_k|\theta=1)\mathbb{P}(\theta=1)}{\mathbb{P}(m=x_k)} = \frac{x_k}{p}\mu_k \cdot p \cdot \frac{1}{\mu_k} = x_k = m$

$\square$

The splitting lemma

The splitting lemma (Aumann, Maschler (1960ies) / folk)

These two sets are equal: for any $\mu \in \Delta([0,1])$ with $\int x \, d\mu(x) = p$ there exists (M, π) s.t. $p'(m) \sim \mu$.
One can take $M = \operatorname{supp} \mu \subset [0,1]$.

Proof for discrete μ via belief-recommendation:

- μ : point x_k has mass μ_k, $\sum \mu_k = 1$, $\sum x_k \cdot \mu_k = p$.
- define
  $\pi_{\theta=1}$: point x_k has mass $\frac{x_k}{p}\mu_k$
  $\pi_{\theta=0}$: point x_k has mass $\frac{1-x_k}{1-p}\mu_k$
- sample $m \sim \pi_\theta$ conditional on θ
- unconditionally $m \sim \mu$: $\mathbb{P}(m = x_k) = (1-p)\frac{1-x_k}{1-p}\mu_k + p\frac{x_k}{p}\mu_k = \mu_k$.
- $p' = \mathbb{P}(\theta = 1 \mid m = x_k) = \frac{\mathbb{P}(m=x_k|\theta=1)\mathbb{P}(\theta=1)}{\mathbb{P}(m=x_k)} = \frac{x_k}{p}\mu_k \cdot p \cdot \frac{1}{\mu_k} = x_k = m$

$\square$

The splitting lemma

The splitting lemma (Aumann, Maschler (1960ies) / folk)

These two sets are equal: for any $\mu \in \Delta([0,1])$ with $\int x \, d\mu(x) = p$ there exists (M, π) s.t. $p'(m) \sim \mu$.
One can take $M = \operatorname{supp} \mu \subset [0,1]$.

Proof for discrete μ via belief-recommendation:

- μ : point x_k has mass μ_k, $\sum \mu_k = 1$, $\sum x_k \cdot \mu_k = p$.
- define
  $\pi_{\theta=1}$: point x_k has mass $\frac{x_k}{p} \mu_k$
  $\pi_{\theta=0}$: point x_k has mass $\frac{1-x_k}{1-p} \mu_k$
- sample $m \sim \pi_\theta$ conditional on θ
- unconditionally $m \sim \mu$: $\mathbb{P}(m = x_k) = (1-p)\frac{1-x_k}{1-p}\mu_k + p\frac{x_k}{p}\mu_k = \mu_k$.
- $p' = \mathbb{P}(\theta = 1 \mid m = x_k) = \frac{\mathbb{P}(m=x_k|\theta=1)\mathbb{P}(\theta=1)}{\mathbb{P}(m=x_k)} = \frac{x_k}{p}\mu_k \cdot p \cdot \frac{1}{\mu_k} = x_k = m$

$\square$

The splitting lemma

The splitting lemma (Aumann, Maschler (1960ies) / folk)

These two sets are equal: for any $\mu \in \Delta([0,1])$ with $\int x \, d\mu(x) = p$ there exists (M, π) s.t. $p'(m) \sim \mu$.
One can take $M = \operatorname{supp} \mu \subset [0,1]$.

Proof for discrete μ via *belief-recommendation:*

- μ : point x_k has mass μ_k, $\sum \mu_k = 1$, $\sum x_k \cdot \mu_k = p$.
- define
  $\pi_{\theta=1}$: point x_k has mass $\frac{x_k}{p}\mu_k$
  $\pi_{\theta=0}$: point x_k has mass $\frac{1-x_k}{1-p}\mu_k$
- sample $m \sim \pi_\theta$ conditional on θ
- unconditionally $m \sim \mu$: $\mathbb{P}(m = x_k) = (1-p)\frac{1-x_k}{1-p}\mu_k + p\frac{x_k}{p}\mu_k = \mu_k$.
- $p' = \mathbb{P}(\theta = 1 \mid m = x_k) = \frac{\mathbb{P}(m=x_k|\theta=1)\mathbb{P}(\theta=1)}{\mathbb{P}(m=x_k)} = \frac{x_k}{p}\mu_k \cdot p \cdot \frac{1}{\mu_k} = x_k = m$

$\square$

The splitting lemma

The splitting lemma (Aumann, Maschler (1960ies) / folk)

These two sets are equal: for any $\mu \in \Delta([0,1])$ with $\int x \, d\mu(x) = p$ there exists (M, π) s.t. $p'(m) \sim \mu$.
One can take $M = \operatorname{supp} \mu \subset [0,1]$.

Corollary: S's optimal payoff is

$$\max_{(M,\pi)} \mathbb{E}\Big[U_S(p'(m))\Big] = \max_{\mu \in \mathcal{D}(p)} \int_{[0,1]} U_S(x) \, d\mu(x),$$

where $\mathcal{D}(p) = \Big\{\mu \in \Delta([0,1]) \,:\, \int_{[0,1]} x \, d\mu(x) = p\Big\}$.

Concavification as martingale-optimization

Concavification of continuous f on $[0,1]$ is

$$\mathrm{Cav}\,[f](x) = \min\{\varphi(x) \,:\, \varphi \text{ is concave and } f(\cdot) \leq \varphi(\cdot)\}$$

Concavification as martingale-optimization

Concavification of continuous f on $[0,1]$ is

$$\mathrm{Cav}\,[f](x) = \min\{\varphi(x) \,:\, \varphi \text{ is concave and } f(\cdot) \leq \varphi(\cdot)\}$$

Lemma

$$\mathrm{Cav}\,[f](p) = \max_{\mu \in \mathcal{D}(p)} \int_{[0,1]} f(x) d\mu(x),$$

where $\mathcal{D}(p) = \left\{\mu \in \Delta([0,1]) \,:\, \int_{[0,1]} x \, d\mu(x) = p\right\}$

Concavification as martingale-optimization

Concavification of continuous f on $[0,1]$ is

$$\mathrm{Cav}\,[f](x) = \min\{\varphi(x) \,:\, \varphi \text{ is concave and } f(\cdot) \le \varphi(\cdot)\}$$

Lemma

$$\mathrm{Cav}\,[f](p) = \max_{\mu \in \mathcal{D}(p)} \int_{[0,1]} f(x) d\mu(x),$$

where $\mathcal{D}(p) = \left\{\mu \in \Delta([0,1]) \,:\, \int_{[0,1]} x\, d\mu(x) = p\right\}$

Proof: Denote r.h.s. by $g(p)$.

$$f(p) \le g(p) \le \mathrm{Cav}\,[f](p).$$

Concavification as martingale-optimization

Concavification of continuous f on $[0,1]$ is

$$\mathrm{Cav}\,[f](x) = \min\{\varphi(x) \,:\, \varphi \text{ is concave and } f(\cdot) \leq \varphi(\cdot)\}$$

Lemma

$$\mathrm{Cav}\,[f](p) = \max_{\mu \in \mathcal{D}(p)} \int_{[0,1]} f(x) d\mu(x),$$

where $\mathcal{D}(p) = \left\{\mu \in \Delta([0,1]) \,:\, \int_{[0,1]} x \, d\mu(x) = p\right\}$

Proof: Denote r.h.s. by $g(p)$.

$$f(p) \leq g(p) \leq \mathrm{Cav}\,[f](p) \impliedby (\delta_p \in \mathcal{D}(p) \text{ \& Jensen's inequality}).$$

Concavification as martingale-optimization

Concavification of continuous f on $[0,1]$ is

$$\mathrm{Cav}\,[f](x) = \min\{\varphi(x) \,:\, \varphi \text{ is concave and } f(\cdot) \le \varphi(\cdot)\}$$

Lemma

$$\mathrm{Cav}\,[f](p) = \max_{\mu \in \mathcal{D}(p)} \int_{[0,1]} f(x)d\mu(x),$$

where $\mathcal{D}(p) = \left\{\mu \in \Delta([0,1]) \,:\, \int_{[0,1]} x\, d\mu(x) = p\right\}$

Proof: Denote r.h.s. by $g(p)$.

$$f(p) \le g(p) \le \mathrm{Cav}\,[f](p) \impliedby (\delta_p \in \mathcal{D}(p) \text{ \& Jensen's inequality}).$$

Prove concavity of g, i.e., $\alpha g(p_1) + (1-\alpha)g(p_2) \le g(\alpha p_1 + (1-\alpha)p_2)$.
Pick optimal $\mu_1 \in \mathcal{D}(p_1)$, $\mu_2 \in \mathcal{D}(p_2)$

$$\alpha g(p_1) + (1-\alpha)g(p_2) = \int f(x)d\left(\alpha\mu_1 + (1-\alpha)\mu_2\right) \le g(\alpha p_1 + (1-\alpha)p_2).$$

Concavification as martingale-optimization

Concavification of continuous f on $[0,1]$ is

$$\mathrm{Cav}\,[f](x) = \min\{\varphi(x)\ :\ \varphi \text{ is concave and } f(\,\cdot\,) \le \varphi(\,\cdot\,)\}$$

Lemma

$$\mathrm{Cav}\,[f](p) = \max_{\mu\in\mathcal{D}(p)} \int_{[0,1]} f(x)d\mu(x),$$

where $\mathcal{D}(p) = \left\{\mu \in \Delta([0,1])\ :\ \int_{[0,1]} x\, d\mu(x) = p\right\}$

Proof: Denote r.h.s. by $g(p)$.

$$f(p) \le g(p) \le \mathrm{Cav}\,[f](p) \iff (\delta_p \in \mathcal{D}(p) \text{ \& Jensen's inequality}).$$

Prove concavity of g, i.e., $\alpha g(p_1) + (1-\alpha)g(p_2) \le g(\alpha p_1 + (1-\alpha)p_2)$.
Pick optimal $\mu_1 \in \mathcal{D}(p_1)$, $\mu_2 \in \mathcal{D}(p_2)$

$$\alpha g(p_1)+(1-\alpha)g(p_2) = \int f(x)d\left(\alpha\mu_1 + (1-\alpha)\mu_2\right) \le g(\alpha p_1+(1-\alpha)p_2).$$

So g is a concave function above $f \Rightarrow (g \ge \mathrm{Cav}\,[f]) \Rightarrow (g \equiv \mathrm{Cav}\,[f])$.

Corollary: the $\mathrm{Cav}\,[U]$-theorem for S's optimal payoff

Theorem (Kamenica, Gentzkow, 2011)

$$\max_{(M,\pi)} \mathbb{E}\Big[U_S(p')\Big] = \mathrm{Cav}\,[U_S](p)$$

Corollary: the $\mathrm{Cav}\,[U]$-theorem for S's optimal payoff

Theorem (Kamenica, Gentzkow, 2011)

$$\max_{(M,\pi)} \mathbb{E}\Big[U_S(p')\Big] = \mathrm{Cav}\,[U_S](p)$$

Proof: By the two lemmas,

$$\max_{(M,\pi)} \mathbb{E}\Big[U_S(p')\Big] = \max_{\mu\in\mathcal{D}(p)} \int_{[0,1]} U_S(x)\, d\mu(x) = \mathrm{Cav}\,[U_S](p). \quad \square$$

Corollary: the $\operatorname{Cav}[U]$-theorem for S's optimal payoff

Theorem (Kamenica, Gentzkow, 2011)

$$\max_{(M,\pi)} \mathbb{E}\Big[U_S(p')\Big] = \operatorname{Cav}[U_S](p)$$

Proof: By the two lemmas,

$$\max_{(M,\pi)} \mathbb{E}\Big[U_S(p')\Big] = \max_{\mu \in \mathcal{D}(p)} \int_{[0,1]} U_S(x)\, d\mu(x) = \operatorname{Cav}[U_S](p). \quad \square$$

Corollary: It is always enough to assume that $M \subset [0,1]$ and signals=induced beliefs.

Corollary: the $\mathrm{Cav}\,[U]$-theorem for S's optimal payoff

Theorem (Kamenica, Gentzkow, 2011)

$$\max_{(M,\pi)} \mathbb{E}\Big[U_S(p')\Big] = \mathrm{Cav}\,[U_S](p)$$

Proof: By the two lemmas,

$$\max_{(M,\pi)} \mathbb{E}\Big[U_S(p')\Big] = \max_{\mu\in\mathcal{D}(p)} \int_{[0,1]} U_S(x)\, d\mu(x) = \mathrm{Cav}\,[U_S](p). \quad \square$$

Corollary: It is always enough to assume that $M \subset [0,1]$ and signals=induced beliefs.

Remark: We will see that $|M| = 2$ is enough.

Back to the toy example

The court problem

- $p = 0.25$ are guilty ($\theta = 1$), Prosecutor (P) observes θ
- Judge (J) has two actions: to acquit ($a = 0$) or to convict ($a = 1$)
- $u_P(a, \theta) = a, \quad u_J(a, \theta) = \mathbb{1}_{a=\theta}$

Back to the toy example

The court problem

- $p = 0.25$ are guilty ($\theta = 1$), Prosecutor (P) observes θ
- Judge (J) has two actions: to acquit ($a = 0$) or to convict ($a = 1$)
- $u_P(a, \theta) = a, \quad u_J(a, \theta) = \mathbb{1}_{a=\theta}$

- Optimal action of J: $\quad a^*(p') = \begin{cases} \text{convict} & p' \geq 0.5 \\ \text{acquit} & p' < 0.5. \end{cases}$

Back to the toy example

The court problem

- $p = 0.25$ are guilty ($\theta = 1$), Prosecutor (P) observes θ
- Judge (J) has two actions: to acquit ($a = 0$) or to convict ($a = 1$)
- $u_P(a, \theta) = a, \quad u_J(a, \theta) = \mathbb{1}_{a=\theta}$

- Optimal action of J: $a^*(p') = \begin{cases} \text{convict} & p' \geq 0.5 \\ \text{acquit} & p' < 0.5. \end{cases}$
- P's payoff as a function of p' and its concavification:

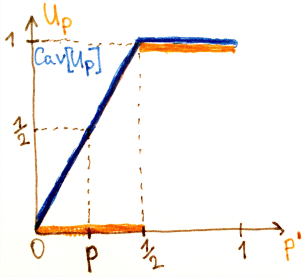

Back to the toy example

The court problem

- $p = 0.25$ are guilty ($\theta = 1$), Prosecutor (P) observes θ
- Judge (J) has two actions: to acquit ($a = 0$) or to convict ($a = 1$)
- $u_P(a, \theta) = a, \quad u_J(a, \theta) = \mathbb{1}_{a=\theta}$

- P's payoff as a function of p' and its concavification:

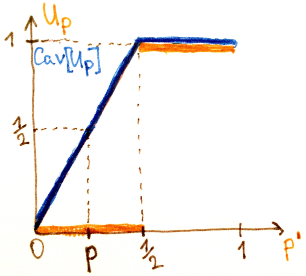

- $\mathrm{Cav}\,[U_P](0.25) = \frac{1}{2} = \frac{1}{2} U_P(0) + \frac{1}{2} U_P(0.5) = \int U_P(x)\, d\left(\frac{1}{2}\delta_0 + \frac{1}{2}\delta_{0.5}\right)$

Back to the toy example

- P's payoff as a function of p' and its concavification:

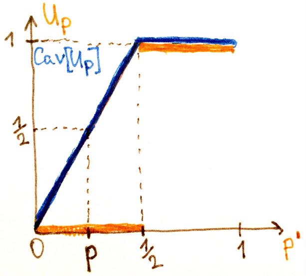

- $\mathrm{Cav}\,[U_P](0.25) = \frac{1}{2} = \frac{1}{2}U_P(0) + \frac{1}{2}U_P(0.5) = \int U_P(x)\,d\left(\frac{1}{2}\delta_0 + \frac{1}{2}\delta_{0.5}\right)$
- Two signals $m \in \{"0","0.5"\}$ with distribution π:

$$\begin{aligned}
\pi_{\theta=1} &= \left(\frac{0}{p}\cdot\frac{1}{2},\ \frac{0.5}{p}\cdot\frac{1}{2}\right) &= (0,1)\\
\pi_{\theta=0} &= \left(\frac{1-0}{1-p}\cdot\frac{1}{2},\ \frac{1-0.5}{1-p}\cdot\frac{1}{2}\right) &= \left(\frac{2}{3},\frac{1}{3}\right)
\end{aligned}$$

Why 2 signals are always enough:

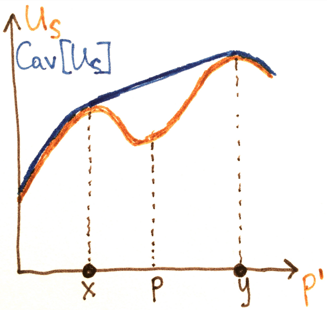

- if $\mathrm{Cav}\,[U_S](p) = U_S(p)$, send a dummy signal
- if $\mathrm{Cav}\,[U_S](p) > U_S(p)$ we can "split" p:

Why 2 signals are always enough:

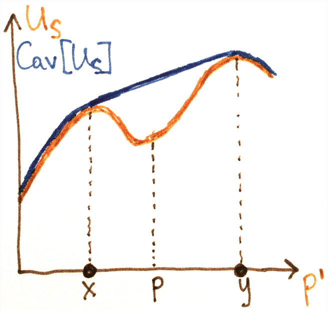

- if $\operatorname{Cav}[U_S](p) = U_S(p)$, send a dummy signal
- if $\operatorname{Cav}[U_S](p) > U_S(p)$ we can "split" p:

Why 2 signals are always enough:

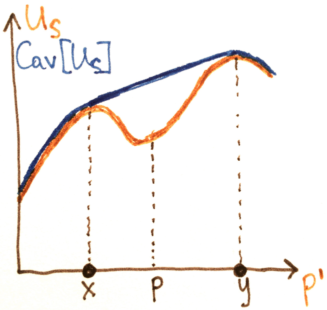

- if $\mathrm{Cav}\,[U_S](p) = U_S(p)$, send a dummy signal
- if $\mathrm{Cav}\,[U_S](p) > U_S(p)$ we can "split" p:
  find $x < p < y$ such that

$$\mathrm{Cav}\,[U_S](p) = (1-\alpha)U_S(x) + \alpha U_S(y) \quad \text{and} \quad p = (1-\alpha)x + \alpha y$$

Why 2 signals are always enough:

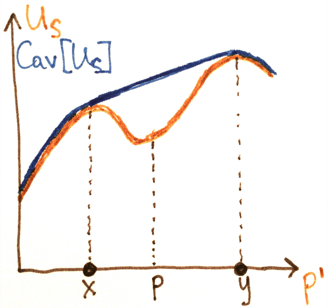

- if $\operatorname{Cav}[U_S](p) = U_S(p)$, send a dummy signal
- if $\operatorname{Cav}[U_S](p) > U_S(p)$ we can "split" p:
  find $x < p < y$ such that

$$\operatorname{Cav}[U_S](p) = (1-\alpha)U_S(x) + \alpha U_S(y) \quad \text{and} \quad p = (1-\alpha)x + \alpha y$$

$$\iff \operatorname{Cav}[U_S](p) = \int U_S \, d\mu, \text{ where } \mu = (1-\alpha)\delta_x + \alpha\delta_y \in \mathcal{D}(p)$$

Why 2 signals are always enough:

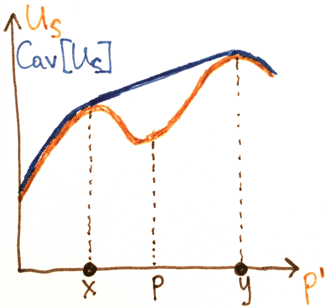

More abstract point of view:

- $\int U_S \, d\mu$ is a linear functional of μ on a convex set $\mathcal{D}(p)$
- Bauer's maximum principle: a convex functional on a convex set attains its maximum at an extreme point.

  $z \in K$ is an extreme point of a convex set K if z cannot be represented as a convex combination of two distinct points $w, w' \in K$.
- Extreme points of $\mathcal{D}(p)$ are two-point distributions $(1-\alpha)\delta_x + \alpha\delta_y$ with $(1-\alpha)x + \alpha \cdot y = p$.

Why 2 signals are always enough:

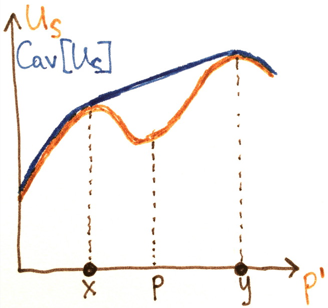

More abstract point of view:

- $\int U_S \, d\mu$ is a linear functional of μ on a convex set $\mathcal{D}(p)$
- Bauer's maximum principle: a convex functional on a convex set attains its maximum at an extreme point.

  $z \in K$ is an extreme point of a convex set K if z cannot be represented as a convex combination of two distinct points $w, w' \in K$.
- Extreme points of $\mathcal{D}(p)$ are two-point distributions $(1-\alpha)\delta_x + \alpha\delta_y$ with $(1-\alpha)x + \alpha \cdot y = p$.

Why 2 signals are always enough:

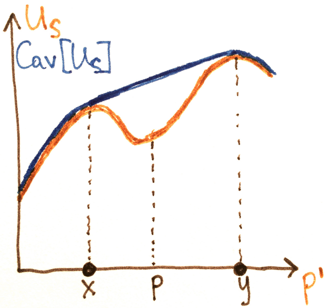

More abstract point of view:

- $\int U_S \, d\mu$ is a linear functional of μ on a convex set $\mathcal{D}(p)$
- Bauer's maximum principle: a convex functional on a convex set attains its maximum at an extreme point.

  $z \in K$ is an extreme point of a convex set K if z cannot be represented as a convex combination of two distinct points $w, w' \in K$.
- Extreme points of $\mathcal{D}(p)$ are two-point distributions $(1-\alpha)\delta_x + \alpha\delta_y$ with $(1-\alpha)x + \alpha \cdot y = p$.

Extension to $|\Theta| > 2$

All the results & proofs are the same with the following modifications:

- p' is a posterior distribution, $p' \in \Delta(\Theta)$
- $\mu_{(M,\pi)} \in \Delta(\Delta(\Theta))$

The only change: need $|\Theta|$ signals

Application to repeated games

Reminder:

- Repeated zero-sum game $G_T(p)$ with incomplete information:
  - a state $\theta \in \{0, 1\}$ with prior p. P1 observes θ, P2 does not
  - a zero-sum game A^θ is played T times, the history is observable.
  - the payoff to P1 is $\mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T} A^\theta_{i_t, j_t}\right]$
- P1 can guarantee $\mathrm{Cav}\,[u](p)$, where $u(p) = \mathrm{val}\big[(1-p)A^0 + p \cdot A^1\big]$.

Application to repeated games

Reminder:

- Repeated zero-sum game $G_T(p)$ with incomplete information:
  - a state $\theta \in \{0,1\}$ with prior p. P1 observes θ, P2 does not
  - a zero-sum game A^θ is played T times, the history is observable.
  - the payoff to P1 is $\mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T} A^\theta_{i_t,j_t}\right]$
- P1 can guarantee $\mathrm{Cav}\,[u](p)$, where $u(p) = \mathrm{val}\big[(1-p)A^0 + p\cdot A^1\big]$.

Application to repeated games

Reminder:

- Repeated zero-sum game $G_T(p)$ with incomplete information:
  - a state $\theta \in \{0, 1\}$ with prior p. P1 observes θ, P2 does not
  - a zero-sum game A^θ is played T times, the history is observable.
  - the payoff to P1 is $\mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T} A^\theta_{i_t, j_t}\right]$
- P1 can guarantee $\mathrm{Cav}\,[u](p)$, where $u(p) = \mathrm{val}\big[(1-p)A^0 + p \cdot A^1\big]$.

Application to repeated games

Reminder:

- Repeated zero-sum game $G_T(p)$ with incomplete information:
  - a state $\theta \in \{0, 1\}$ with prior p. P1 observes θ, P2 does not
  - a zero-sum game A^θ is played T times, the history is observable.
  - the payoff to P1 is $\mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T} A^\theta_{i_t,j_t}\right]$
- P1 can guarantee $\mathrm{Cav}\,[u](p)$, where $u(p) = \mathrm{val}\big[(1-p)A^0 + p \cdot A^1\big]$.

Application to repeated games

Reminder:

- Repeated zero-sum game $G_T(p)$ with incomplete information:
  - a state $\theta \in \{0, 1\}$ with prior p. P1 observes θ, P2 does not
  - a zero-sum game A^θ is played T times, the history is observable.
  - the payoff to P1 is $\mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T} A^\theta_{i_t,j_t}\right]$
- P1 can guarantee $\mathrm{Cav}\,[u](p)$, where $u(p) = \mathrm{val}\big[(1-p)A^0 + p \cdot A^1\big]$.

Question: How can P1 guarantee $\mathrm{Cav}\,[u](p)$ in a long run?

Application to repeated games

Reminder:

- Repeated zero-sum game $G_T(p)$ with incomplete information:
  - a state $\theta \in \{0,1\}$ with prior p. P1 observes θ, P2 does not
  - a zero-sum game A^θ is played T times, the history is observable.
  - the payoff to P1 is $\mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T} A^\theta_{i_t,j_t}\right]$
- P1 can guarantee $\mathrm{Cav}\,[u](p)$, where $u(p) = \mathrm{val}\big[(1-p)A^0 + p\cdot A^1\big]$.

Question: How can P1 guarantee $\mathrm{Cav}\,[u](p)$ in a long run?

- If it was a persuasion problem, P1 could send a signal m to P2 such that $\mathbb{E}[u(p'(m))] = \mathrm{Cav}\,[u](p)$
- **Idea:** P1 can use first-stage action as a signal to induce p' and then play the optimal strategy from $(1-p')A^0 + p'\cdot A^1$ at all stages.
- P1 gets at least

$$\mathbb{E}\left[\frac{1}{T}\cdot(T-1)\cdot u(p')\right] = \mathrm{Cav}\,[u]\cdot\left(1-\frac{1}{T}\right) \to \mathrm{Cav}\,[u], \quad T\to\infty.$$

Application to repeated games

Reminder:

- Repeated zero-sum game $G_T(p)$ with incomplete information:
  - a state $\theta \in \{0,1\}$ with prior p. P1 observes θ, P2 does not
  - a zero-sum game A^θ is played T times, the history is observable.
  - the payoff to P1 is $\mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T} A^\theta_{i_t,j_t}\right]$
- P1 can guarantee $\mathrm{Cav}\,[u](p)$, where $u(p) = \mathrm{val}\big[(1-p)A^0 + p\cdot A^1\big]$.

Question: How can P1 guarantee $\mathrm{Cav}\,[u](p)$ in a long run?

- If it was a persuasion problem, P1 could send a signal m to P2 such that $\mathbb{E}[u(p'(m))] = \mathrm{Cav}\,[u](p)$
- **Idea:** P1 can use first-stage action as a signal to induce p' and then play the optimal strategy from $(1-p')A^0 + p'\cdot A^1$ at all stages.
- P1 gets at least

$$\mathbb{E}\left[\frac{1}{T}\cdot(T-1)\cdot u(p')\right] = \mathrm{Cav}\,[u]\cdot\left(1-\frac{1}{T}\right) \to \mathrm{Cav}\,[u], \quad T\to\infty.$$

Application to repeated games

Reminder:

- Repeated zero-sum game $G_T(p)$ with incomplete information:
  - a state $\theta \in \{0,1\}$ with prior p. P1 observes θ, P2 does not
  - a zero-sum game A^θ is played T times, the history is observable.
  - the payoff to P1 is $\mathbb{E}\left[\frac{1}{T}\sum_{t=0}^{T} A^\theta_{i_t,j_t}\right]$
- P1 can guarantee $\mathrm{Cav}\,[u](p)$, where $u(p) = \mathrm{val}\big[(1-p)A^0 + p\cdot A^1\big]$.

Question: How can P1 guarantee $\mathrm{Cav}\,[u](p)$ in a long run?

- If it was a persuasion problem, P1 could send a signal m to P2 such that $\mathbb{E}[u(p'(m))] = \mathrm{Cav}\,[u](p)$
- **Idea:** P1 can use first-stage action as a signal to induce p' and then play the optimal strategy from $(1-p')A^0 + p'\cdot A^1$ at all stages.
- P1 gets at least

$$\mathbb{E}\left[\frac{1}{T}\cdot(T-1)\cdot u(p')\right] = \mathrm{Cav}\,[u]\cdot\left(1-\frac{1}{T}\right) \to \mathrm{Cav}\,[u], \quad T\to\infty.$$

Action-recommendation approach: revelation principle

Idea of action-recommendation approach

Definition: (M, π) is an action-recommendation (AR) information structure $\Leftrightarrow$ $M = A$ and $a^*(a) = a$ (it is in R's best interest to play the action matching the signal aka obedience constraint).

- Similar to belief-recommendation from Splitting lemma

Idea of action-recommendation approach

Definition: (M, π) is an action-recommendation (AR) information structure $\Leftrightarrow$ $M = A$ and $a^*(a) = a$ (it is in R's best interest to play the action matching the signal aka obedience constraint).

- Similar to belief-recommendation from Splitting lemma

Revelation principle

For any (M, π) there exists AR (A, ψ) with the same S's payoff.

Idea of action-recommendation approach

Definition: (M, π) is an action-recommendation (AR) information structure $\Leftrightarrow M = A$ and $a^*(a) = a$ (it is in R's best interest to play the action matching the signal aka obedience constraint).

Revelation principle

For any (M, π) there exists AR (A, ψ) with the same S's payoff.

Proof: How (M, π) works:

$$\theta \to (m \sim \pi_\theta) \to a^*(m).$$

Idea of action-recommendation approach

Definition: (M, π) is an action-recommendation (AR) information structure $\Leftrightarrow$ $M = A$ and $a^*(a) = a$ (it is in R's best interest to play the action matching the signal aka obedience constraint).

Revelation principle

For any (M, π) there exists AR (A, ψ) with the same S's payoff.

Proof: How (M, π) works:

$$\theta \to (m \sim \pi_\theta) \to a^*(m).$$

Denote ψ_θ the distribution of a^* conditional on θ.
(A, ψ) recommends the action a whenever R plays a for (M, π).

Idea of action-recommendation approach

Definition: (M, π) is an action-recommendation (AR) information structure $\Leftrightarrow$ $M = A$ and $a^*(a) = a$ (it is in R's best interest to play the action matching the signal aka obedience constraint).

Revelation principle

For any (M, π) there exists AR (A, ψ) with the same S's payoff.

Proof: How (M, π) works:

$$\theta \to (m \sim \pi_\theta) \to a^*(m).$$

Denote ψ_θ the distribution of a^* conditional on θ.
(A, ψ) recommends the action a whenever R plays a for (M, π).

- $a^*(a) = a$. Otherwise, R can improve his action for (M, π).

Idea of action-recommendation approach

Definition: (M, π) is an action-recommendation (AR) information structure $\Leftrightarrow M = A$ and $a^*(a) = a$ (it is in R's best interest to play the action matching the signal aka obedience constraint).

Revelation principle

For any (M, π) there exists AR (A, ψ) with the same S's payoff.

Proof: How (M, π) works:

$$\theta \to (m \sim \pi_\theta) \to a^*(m).$$

Denote ψ_θ the distribution of a^* conditional on θ.
(A, ψ) recommends the action a whenever R plays a for (M, π).

- $a^*(a) = a$. Otherwise, R can improve his action for (M, π).
- the payoff is the same. $\square$

Idea of action-recommendation approach

Definition: (M, π) is an action-recommendation (AR) information structure $\Leftrightarrow$ $M = A$ and $a^*(a) = a$ (it is in R's best interest to play the action matching the signal aka obedience constraint).

Revelation principle

For any (M, π) there exists AR (A, ψ) with the same S's payoff.

Proof: How (M, π) works:

$$\theta \to (m \sim \pi_\theta) \to a^*(m).$$

Denote ψ_θ the distribution of a^* conditional on θ.
(A, ψ) recommends the action a whenever R plays a for (M, π).

- $a^*(a) = a$. Otherwise, R can improve his action for (M, π).
- the payoff is the same. $\square$

Corollary:

- Restriction to AR is w.l.o.g.
- $\min\{|\Theta|, |A|\}$ signals are enough for optimal persuasion.

Idea of action-recommendation approach

Definition: (M, π) is an action-recommendation (AR) information structure $\Leftrightarrow M = A$ and $a^*(a) = a$ (it is in R's best interest to play the action matching the signal aka obedience constraint).

Revelation principle

For any (M, π) there exists AR (A, ψ) with the same S's payoff.

Proof: How (M, π) works:

$$\theta \to (m \sim \pi_\theta) \to a^*(m).$$

Denote ψ_θ the distribution of a^* conditional on θ.
(A, ψ) recommends the action a whenever R plays a for (M, π).

- $a^*(a) = a$. Otherwise, R can improve his action for (M, π).
- the payoff is the same. $\square$

Corollary:

- Restriction to AR is w.l.o.g.
- $\min\{|\Theta|, |A|\}$ signals are enough for optimal persuasion.

Persuasion as a linear program

The **obedience constraint** for (A, ψ):

$$\mathbb{E}\big[u_R(a,\theta) \mid m = a\big] \geq \mathbb{E}\big[u_R(\tilde{a},\theta) \mid m = a\big] \quad \text{for all distinct } a, \tilde{a} \in A.$$

Persuasion as a linear program

The **obedience constraint** for (A, ψ):

$$\mathbb{E}\big[u_R(a,\theta) \mid m = a\big] \geq \mathbb{E}\big[u_R(\tilde{a},\theta) \mid m = a\big] \quad \text{for all distinct } a, \tilde{a} \in A.$$

Rewrite the l.h.s.:

$$l.h.s. = \sum_{\theta\in\Theta} u_R(a,\theta) \cdot \mathbb{P}(\theta \mid m = a) = \sum_{\theta\in\Theta} u_R(a,\theta) \cdot \frac{p(\theta)\cdot\psi_\theta(a)}{\mathbb{P}(m = a)}.$$

Persuasion as a linear program

The **obedience constraint** for (A, ψ):

$$\mathbb{E}\big[u_R(a,\theta) \mid m = a\big] \geq \mathbb{E}\big[u_R(\tilde{a},\theta) \mid m = a\big] \quad \text{for all distinct } a, \tilde{a} \in A.$$

Rewrite the l.h.s. (r.h.s. is similar):

$$l.h.s. = \sum_{\theta \in \Theta} u_R(a,\theta) \cdot \mathbb{P}(\theta \mid m = a) = \sum_{\theta \in \Theta} u_R(a,\theta) \cdot \frac{p(\theta) \cdot \psi_\theta(a)}{\mathbb{P}(m = a)}.$$

Persuasion as a linear program

The **obedience constraint** for (A, ψ): Rewrite the l.h.s. (r.h.s. is similar):

$$l.h.s. = \sum_{\theta \in \Theta} u_R(a, \theta) \cdot \mathbb{P}(\theta \mid m = a) = \sum_{\theta \in \Theta} u_R(a, \theta) \cdot \frac{p(\theta) \cdot \psi_\theta(a)}{\mathbb{P}(m = a)}.$$

Obedience $\iff$

$$\sum_{\theta \in \Theta} \big(u_R(a, \theta) - u_R(\tilde{a}, \theta)\big) \cdot p(\theta) \cdot \psi_\theta(a) \geq 0 \quad \forall a \neq \tilde{a} \in A.$$

Persuasion as a linear program

The **obedience constraint** for (A, ψ):

$$\sum_{\theta \in \Theta} \big(u_R(a, \theta) - u_R(\tilde{a}, \theta)\big) \cdot p(\theta) \cdot \psi_\theta(a) \geq 0 \quad \forall a \neq \tilde{a} \in A.$$

Corollary: optimal action-recommendation = LP:

$$\textbf{maximize} \sum_{\theta, a} u_S(a, \theta) \cdot p(\theta) \cdot \psi_\theta(a)$$

$$\textbf{over } (\psi_\theta(a))_{\theta \in \Theta, a \in A} \ \textbf{ such that}$$

$$\psi_\theta(a) \geq 0, \quad \& \quad \sum_a \psi_\theta(a) = 1 \ \ \& \ \ \text{Obedience}$$

Persuasion as a linear program

The **obedience constraint** for (A, ψ):

$$\sum_{\theta \in \Theta} \big(u_R(a,\theta) - u_R(\tilde{a},\theta)\big) \cdot p(\theta) \cdot \psi_\theta(a) \geq 0 \quad \forall a \neq \tilde{a} \in A.$$

Corollary: optimal action-recommendation = LP:

$$\textbf{maximize} \sum_{\theta,a} u_S(a,\theta) \cdot p(\theta) \cdot \psi_\theta(a)$$

$$\textbf{over } (\psi_\theta(a))_{\theta \in \Theta, a \in A} \ \ \textbf{such that}$$

$$\psi_\theta(a) \geq 0, \ \ \& \ \ \sum_a \psi_\theta(a) = 1 \ \ \& \ \ \text{Obedience}$$

- Easy to solve algorithmically + structural information about solution + duality
- AR extends to n receivers, who play a game G after receiving the signals. Joint distributions of $(a_1, \ldots, a_n)$ that can be generated by AR = Bayesian Correlated Equilibria of G.

Persuasion as a linear program

The **obedience constraint** for (A, ψ):

$$\sum_{\theta \in \Theta} \big(u_R(a,\theta) - u_R(\tilde{a},\theta)\big) \cdot p(\theta) \cdot \psi_\theta(a) \geq 0 \quad \forall a \neq \tilde{a} \in A.$$

Corollary: optimal action-recommendation = LP:

$$\textbf{maximize} \sum_{\theta, a} u_S(a,\theta) \cdot p(\theta) \cdot \psi_\theta(a)$$

$$\textbf{over } (\psi_\theta(a))_{\theta \in \Theta, a \in A} \; \textbf{ such that}$$

$$\psi_\theta(a) \geq 0, \;\; \& \;\; \sum_a \psi_\theta(a) = 1 \;\; \& \;\; \text{Obedience}$$

- Easy to solve algorithmically + structural information about solution + duality
- AR extends to n receivers, who play a game G after receiving the signals. Joint distributions of $(a_1, \ldots, a_n)$ that can be generated by AR = Bayesian Correlated Equilibria of G.

References

References:

- **E. Kamenica** (2019) Bayesian persuasion and information design
  Annual Review of Economics
  A comprehensive survey of geometric approach, its extensions, and applications.
- **D.Bergemann, S. Morris** (2019) Information design: A unified perspective. Journal of Economic Literature
  Survey of action-recommendation approach, multi-receiver persuasion, and Bayesian correlated equilibrium