Mechanism Design

10: Information Acquisition (and Mechanism Design)

Egor Starkov

Københavns Unversitet
Fall 2024

Introduction

- Throughout the course we assumed the players have some private information that the designer wants to extract
- But where does this information come from?
    - I.e.: what incentives do agents have to acquire info?
    - Especially if they know the designer is going to exploit this knowledge?
- In the previous lecture (info design) we asked how the designer would prefer to inform the players. But an obvious question is: what do *the players* want to learn?
- This lecture is a compilation from a variety of published and working papers, including Asher Wolinsky's lecture notes. If you are interested in models of information acquisition, see a recent survey by Maćkowiak et al. [2023].
    - Note: the notation in this slide deck is not very polished and may be confusing in places.

Grossman-Stiglitz paradox

- The efficient market hypothesis says that in financial markets, all available information is already incorporated into asset prices
- But then no single trader wants to investigate the fundamental value of the assets
- But then how does information feed into the prices?

Grossman-Stiglitz paradox

- The efficient market hypothesis says that in financial markets, all available information is already incorporated into asset prices
- But then no single trader wants to investigate the fundamental value of the assets
- But then how does information feed into the prices?

- This is known as the Grossman and Stiglitz [1980] paradox.
- We've had examples where all private info is extracted – so why would players bother acquiring it?

Grossman-Stiglitz vs Cremer-McLean

Example: Oil companies compete for mining rights, choose how much info to acquire about the oil field richness.

- Players' info/values would be correlated with the true state – the amount of oil in the field – and so with one another's, from the point of view of an unaware seller

Grossman-Stiglitz vs Cremer-McLean

Example: Oil companies compete for mining rights, choose how much info to acquire about the oil field richness.

- Players' info/values would be correlated with the true state – the amount of oil in the field – and so with one another's, from the point of view of an unaware seller
- Seller's interim best option is to extract all info & surplus via Cremer-McLean mechanism.
- Firms know they'll get no surplus either way – so why bother learning?

Grossman-Stiglitz vs Cremer-McLean

Example: Oil companies compete for mining rights, choose how much info to acquire about the oil field richness.

- Players' info/values would be correlated with the true state – the amount of oil in the field – and so with one another's, from the point of view of an unaware seller
- Seller's interim best option is to extract all info & surplus via Cremer-McLean mechanism.
- Firms know they'll get no surplus either way – so why bother learning?
- What would the seller prefer? Learning or no learning?
  - If there are guaranteed gains from trade – probably don't care, expected price is the same either way. E.g.: you own a home, and some oil has been found under it. You don't care how much oil is there, you want to sell it either way.
  - Otherwise might want the buyers to learn, since info increases total surplus – how to induce learning?

Grossman-Stiglitz vs Cremer-McLean

Example: Oil companies compete for mining rights, choose how much info to acquire about the oil field richness.

- Players' info/values would be correlated with the true state – the amount of oil in the field – and so with one another's, from the point of view of an unaware seller
- Seller's interim best option is to extract all info & surplus via Cremer-McLean mechanism.
- Firms know they'll get no surplus either way – so why bother learning?
- What would the seller prefer? Learning or no learning?
  - If there are guaranteed gains from trade – probably don't care, expected price is the same either way. E.g.: you own a home, and some oil has been found under it. You don't care how much oil is there, you want to sell it either way.
  - Otherwise might want the buyers to learn, since info increases total surplus – how to induce learning?
  - I.e., designer only cares about info to the extent that it would affect the allocation – does not want the players to acquire any info beyond that, since it would entail info rents (costly to designer)

This slide deck:

1. How to model information?
2. How to rank information?
3. Costs of information
4. Entropy cost in decision problems
5. Mechanism design with costly learning
6. Mechanism design with costly learning: Examples

How to model information? – Probability spaces

- Let's start from the basics – how to model information?

How to model information? – Probability spaces

- Let's start from the basics – how to model information?
  - Info must be about "something", so we need to define the "something" first.

How to model information? – Probability spaces

- Let's start from the basics – how to model information?
  - Info must be about "something", so we need to define the "something" first.
- A probability space $(\Omega, \mathcal{B}, \mathbb{P})$ consists of:
  - set of states $\omega \in \Omega$,
  - set of events $\mathcal{F}$, where any event is a set of states:
    any $f \in \mathcal{F}$ is a set $f \subseteq \Omega$
  - a probability measure $\mathbb{P}$ over events in $\mathcal{F}$.

How to model information? – Probability spaces

- Let's start from the basics – how to model information?
  - Info must be about "something", so we need to define the "something" first.
- A probability space $(\Omega, \mathcal{B}, \mathbb{P})$ consists of:
  - set of states $\omega \in \Omega$,
  - set of events $\mathcal{F}$, where any event is a set of states: any $f \in \mathcal{F}$ is a set $f \subseteq \Omega$
  - a probability measure $\mathbb{P}$ over events in $\mathcal{F}$.
- If that seems too abstract, it's good enough to think of the following special case:
  - set of states is $\Omega = [0, 1]$
  - set of events is a Borel σ-algebra $\mathcal{B}([0, 1])$ (think "set of all possible subsets of $[0, 1]$ that we care about")
  - probability measure $\mathbb{P}$ is Lebesgue measure over $\mathcal{B}([0, 1])$ (implying a uniform distribution: $\omega \sim U[0, 1]$).

Random variables

- Then any random variable θ can be represented as a function $\theta : \Omega \to \Theta$. So a probability of any given realization $\theta_k \in \Theta$ would be given by $\mathbb{P}\left\{\omega \mid \theta(\omega) = \theta_k\right\} = \mathbb{P}\left\{\theta^{-1}(\theta_k)\right\}$.

Random variables

- Then any random variable θ can be represented as a function $\theta : \Omega \to \Theta$. So a probability of any given realization $\theta_k \in \Theta$ would be given by $\mathbb{P}\{\omega \mid \theta(\omega) = \theta_k\} = \mathbb{P}\left\{\theta^{-1}(\theta_k)\right\}$.
  - E.g., a coin flip $\theta \in \{h, t\}$ can be described as $\theta = \begin{cases} h & \text{if } \omega \in [0, 0.5], \\ t & \text{if } \omega \in (0.5, 1]. \end{cases}$
  - A dice roll $\theta \in \{1, ..., 6\}$ can be described as $\theta = \begin{cases} 1 & \text{if } \omega \in [0, 1/6], \\ ... \end{cases}$

Random variables

- Then any random variable θ can be represented as a function $\theta : \Omega \to \Theta$. So a probability of any given realization $\theta_k \in \Theta$ would be given by $\mathbb{P}\left\{\omega \mid \theta(\omega) = \theta_k\right\} = \mathbb{P}\left\{\theta^{-1}(\theta_k)\right\}$.
  - E.g., a coin flip $\theta \in \{h, t\}$ can be described as $\theta = \begin{cases} h & \text{if } \omega \in [0, 0.5], \\ t & \text{if } \omega \in (0.5, 1]. \end{cases}$
  - A dice roll $\theta \in \{1, ..., 6\}$ can be described as $\theta = \begin{cases} 1 & \text{if } \omega \in [0, 1/6], \\ ... \end{cases}$
- Why do we need this structure? To describe random variables that can be arbitrarily correlated!

Partitions

- **Example:** Suppose the payoff-relevant state is $\theta_0 \in \{h, t\}$, equiprobable.
- Suppose an agent can receive a binary signal $\theta_1 \in \{u, d\}$ s.t. $\mathbb{P}(u) = 0.5$. How to optimally design such a signal?

Partitions

- **Example:** Suppose the payoff-relevant state is $\theta_0 \in \{h, t\}$, equiprobable.
- Suppose an agent can receive a binary signal $\theta_1 \in \{u, d\}$ s.t. $\mathbb{P}(u) = 0.5$. How to optimally design such a signal?

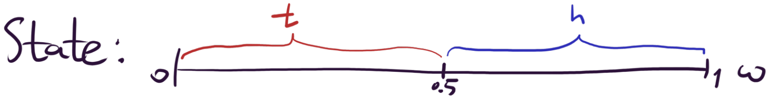

Partitions

- **Example:** Suppose the payoff-relevant state is $\theta_0 \in \{h, t\}$, equiprobable.
- Suppose an agent can receive a binary signal $\theta_1 \in \{u, d\}$ s.t. $\mathbb{P}(u) = 0.5$. How to optimally design such a signal?

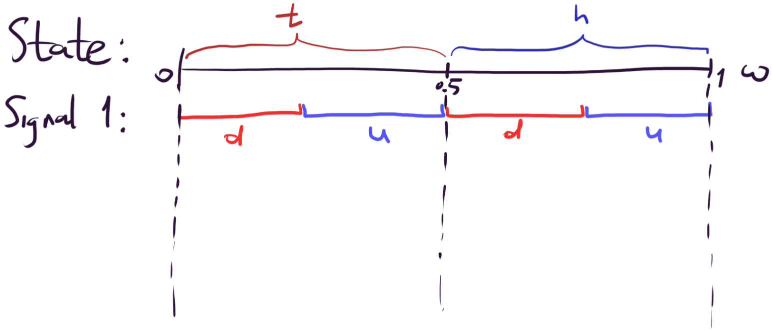

Partitions

- **Example:** Suppose the payoff-relevant state is $\theta_0 \in \{h, t\}$, equiprobable.
- Suppose an agent can receive a binary signal $\theta_1 \in \{u, d\}$ s.t. $\mathbb{P}(u) = 0.5$. How to optimally design such a signal?

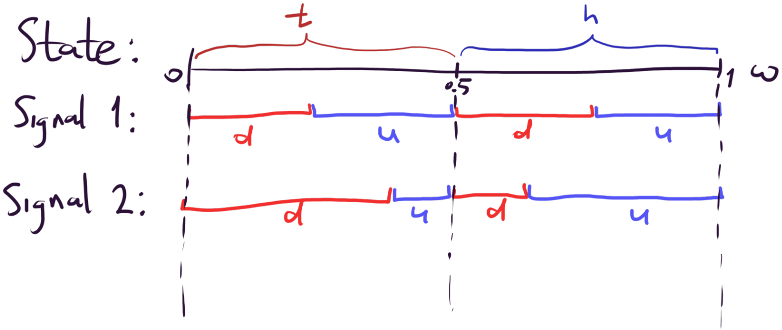

Partitions

- **Example:** Suppose the payoff-relevant state is $\theta_0 \in \{h, t\}$, equiprobable.
- Suppose an agent can receive a binary signal $\theta_1 \in \{u, d\}$ s.t. $\mathbb{P}(u) = 0.5$. How to optimally design such a signal?

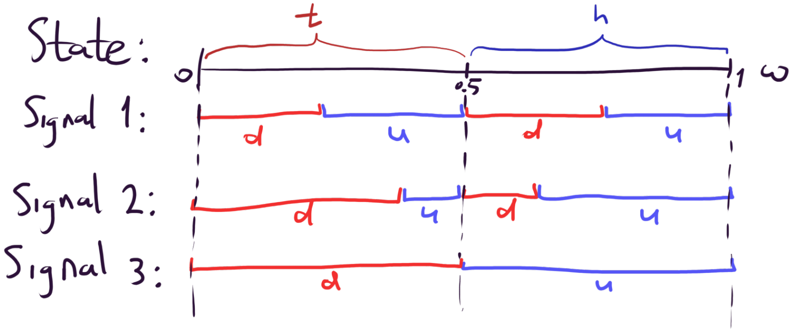

Which of the signals in the picture above is the most informative about θ_0?

Conditional probabilities

- ...but this is not how we modelled information throughout this course?
- Well, it's true that you can capture all of this partitional structure via conditional probabilities:
  - Describe θ_0 in terms of distribution of its realizations, $\Phi_0 \in \Delta\Theta_0$ (there's implicit conditioning on ω),
  - describe θ_1 in terms of distributions of its realizations conditional on realizations of θ_0: $\Phi_1|\theta_0 \in \Delta\Theta_1$, etc
- It is not, however, clear, in general, whether any collection of such conditional probabilities yields a sane joint distribution. Further, conditional probabilities do not uniquely pin down a partitional representation.

This slide deck:

How to rank information? – Garblings

- Next question: how to rank signal structures? What does it mean to be more informative?

How to rank information? – Garblings

- Next question: how to rank signal structures? What does it mean to be more informative?
- One option: a more informative signal contains all info from the less informative signal (and maybe more). Formally:
- **Definition:** signal θ_0 is a garbling of θ_1 if $\mathbb{P}\{\omega|\theta_0, \theta_1\} = \mathbb{P}\{\omega|\theta_1\}$ for all $\omega \in \Omega$ and all realizations $\theta_0 \in \Theta_0, \theta_1 \in \Theta_1$.
    - In words, θ_0 conveys no additional information on top of θ_1.
    - So θ_1 is more informative than its garbling θ_0.

How to rank information? – Garblings

- Next question: how to rank signal structures? What does it mean to be more informative?
- One option: a more informative signal contains all info from the less informative signal (and maybe more). Formally:
- **Definition:** signal θ_0 is a garbling of θ_1 if $\mathbb{P}\{\omega|\theta_0, \theta_1\} = \mathbb{P}\{\omega|\theta_1\}$ for all $\omega \in \Omega$ and all realizations $\theta_0 \in \Theta_0, \theta_1 \in \Theta_1$.
  - In words, θ_0 conveys no additional information on top of θ_1.
  - So θ_1 is more informative than its garbling θ_0.
- Note that this is not a complete ordering – we may be unable to compare two arbitrary signal structures and say that one is necessarily more informative than another.
  - E.g., in the figure above, we can't compare signals' informativeness (about ω)

How to rank information? – Garblings

- Next question: how to rank signal structures? What does it mean to be more informative?
- One option: a more informative signal contains all info from the less informative signal (and maybe more). Formally:
- **Definition:** signal θ_0 is a garbling of θ_1 if $\mathbb{P}\{\omega|\theta_0, \theta_1\} = \mathbb{P}\{\omega|\theta_1\}$ for all $\omega \in \Omega$ and all realizations $\theta_0 \in \Theta_0, \theta_1 \in \Theta_1$.
  - In words, θ_0 conveys no additional information on top of θ_1.
  - So θ_1 is more informative than its garbling θ_0.
- Note that this is not a complete ordering – we may be unable to compare two arbitrary signal structures and say that one is necessarily more informative than another.
  - E.g., in the figure above, we can't compare signals' informativeness (about ω)
  - But can always say that a perfect signal $\theta_p(\omega) = \omega \ \forall \omega \in \Omega$ is more informative than any other signal
  - and that an uninformative signal $\theta_u(\omega) = \theta_u(\omega') \ \forall \omega, \omega' \in \Omega$ is a garbling of any signal.

What does it mean to be a garbling? 1 – Partitions

- Going back to partitions, note that every signal structure (or random variable) θ defines a *partition* of Ω.
- **Definition:** signal structure θ_0 is coarser than θ_1 if $\theta_1^{-1}(\theta_1(\omega)) \subseteq \theta_0^{-1}(\theta_0(\omega))$ for all $\omega \in \Omega$.

What does it mean to be a garbling? 1 – Partitions

- Going back to partitions, note that every signal structure (or random variable) θ defines a *partition* of Ω.
- **Definition:** signal structure θ_0 is coarser than θ_1 if $\theta_1^{-1}(\theta_1(\omega)) \subseteq \theta_0^{-1}(\theta_0(\omega))$ for all $\omega \in \Omega$.

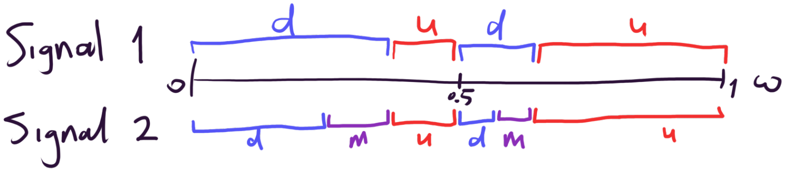

Here, signal [structure] 1 is coarser than signal 2.

What does it mean to be a garbling? 1 – Partitions

- Going back to partitions, note that every signal structure (or random variable) θ defines a *partition* of Ω.
- **Definition:** signal structure θ_0 is coarser than θ_1 if $\theta_1^{-1}(\theta_1(\omega)) \subseteq \theta_0^{-1}(\theta_0(\omega))$ for all $\omega \in \Omega$.

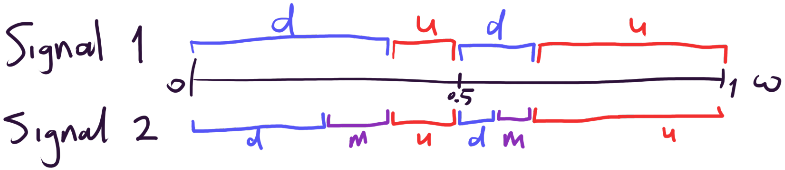

Here, signal [structure] 1 is coarser than signal 2.

- **Proposition:** θ_0 is coarser than θ_1 if and only if θ_0 is a garbling of θ_1

A quick test: is signal 1 coarser than signal 2?

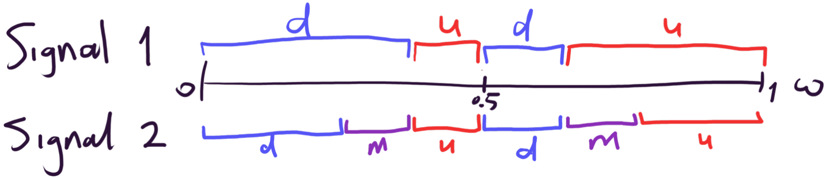

A quick test: is signal 1 coarser than signal 2?

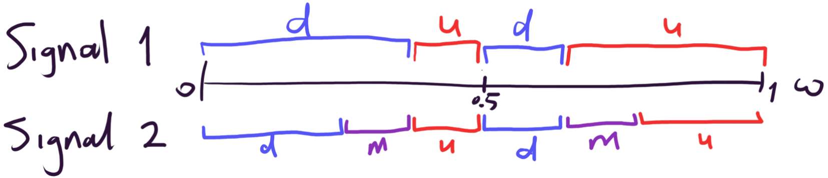

No! Set *m* below is not included in either *d*, or *u* above!

What does it mean to be a garbling? 2 – Conditional probabilities

If we turn to conditional probability representations, then we can phrase the condition as follows:

Proposition (Blackwell's conditions)

θ_0 is a garbling of θ_1 if and only if $\exists z : \Theta_1 \to \Delta\Theta_0$ s.t.:

$$\theta_0(s_0|\omega) = \sum_{s_1 \in \Theta_1} z(s_0|s_1)\theta_1(s_1|\omega) \quad \forall \omega \in \Omega, s_0 \in \Theta_0.$$

Blackwell's Theorem

Consider an agent with an arbitrary utility function $u(x, \omega)$, who chooses some action/outcome $x \in X$ in order to maximize their expected utility. Suppose the agent does not know ω, but receives signal θ.

Blackwell's Theorem

Consider an agent with an arbitrary utility function $u(x, \omega)$, who chooses some action/outcome $x \in X$ in order to maximize their expected utility. Suppose the agent does not know ω, but receives signal θ.

Blackwell's Theorem

$\mathbb{E}\left[u(x^*(\theta), \omega) \mid \theta_1\right] \geq \mathbb{E}\left[u(x^*(\theta), \omega) \mid \theta_0\right]$ for all utility functions u
if and only if θ_0 is a garbling of θ_1.

Blackwell's Theorem

Consider an agent with an arbitrary utility function $u(x, \omega)$, who chooses some action/outcome $x \in X$ in order to maximize their expected utility. Suppose the agent does not know ω, but receives signal θ.

Blackwell's Theorem

$\mathbb{E}\left[u(x^*(\theta), \omega) \mid \theta_1\right] \geq \mathbb{E}\left[u(x^*(\theta), \omega) \mid \theta_0\right]$ for all utility functions u
if and only if θ_0 is a garbling of θ_1.

- Theorem says that in a decision problem (i.e., absent any strategic concerns), an agent would always prefer a *more informative* (less garbled) signal.
- Hence if learning is free, agent always gets a perfect signal.

Blackwell's Theorem

Consider an agent with an arbitrary utility function $u(x, \omega)$, who chooses some action/outcome $x \in X$ in order to maximize their expected utility. Suppose the agent does not know ω, but receives signal θ.

Blackwell's Theorem

$\mathbb{E}\left[u(x^*(\theta), \omega) \mid \theta_1\right] \geq \mathbb{E}\left[u(x^*(\theta), \omega) \mid \theta_0\right]$ for all utility functions u
if and only if θ_0 is a garbling of θ_1.

- Theorem says that in a decision problem (i.e., absent any strategic concerns), an agent would always prefer a *more informative* (less garbled) signal.
- Hence if learning is free, agent always gets a perfect signal.
- In mechdesign environment, if learning is unobservable (to designer/other players), the same holds true.

Blackwell's Theorem

Consider an agent with an arbitrary utility function $u(x, \omega)$, who chooses some action/outcome $x \in X$ in order to maximize their expected utility. Suppose the agent does not know ω, but receives signal θ.

Blackwell's Theorem

$\mathbb{E}\left[u(x^*(\theta), \omega) \mid \theta_1\right] \geq \mathbb{E}\left[u(x^*(\theta), \omega) \mid \theta_0\right]$ for all utility functions u
if and only if θ_0 is a garbling of θ_1.

- Theorem says that in a decision problem (i.e., absent any strategic concerns), an agent would always prefer a *more informative* (less garbled) signal.
- Hence if learning is free, agent always gets a perfect signal.
- In mechdesign environment, if learning is unobservable (to designer/other players), the same holds true.
- But what if learning is costly? Actually, how to even impose a cost on information?

This slide deck:

Putting a price on information

- What assumptions on information cost function are reasonable?
  - *Note:* this is a very active research area, and there's heated debate about answers to this question. I present one answer here, but it is not universally accepted.

Putting a price on information

- What assumptions on information cost function are reasonable?
  - *Note:* this is a very active research area, and there's heated debate about answers to this question. I present one answer here, but it is not universally accepted.
- One approach: measure somehow the information $H(\phi)$ contained in belief $\phi \in \Delta(\Omega)$, and let the cost of a signal be proportional to the expected amount of information it adds:

$$C(\theta) = \lambda \cdot \left[\mathbb{E}[H(\phi(\theta))] - H(\phi_0)\right].$$

How to measure information?

- Then how should our measure $H(\phi)$ of information in belief ϕ look like?
- Say that we derive information $\psi(\mathbb{P}(E))$ from the fact that event $E \subset \Omega$ realized.
  (Then we could define $H(\phi) \equiv \mathbb{E}\left[\psi(\phi(\omega))\right]$ when Ω is finite.)

How to measure information?

- Then how should our measure $H(\phi)$ of information in belief ϕ look like?
- Say that we derive information $\psi(\mathbb{P}(E))$ from the fact that event $E \subset \Omega$ realized. (Then we could define $H(\phi) \equiv \mathbb{E}\left[\psi(\phi(\omega))\right]$ when Ω is finite.)
- Consider the three following (natural?) **axioms**:

1. Positivity: $\psi(p) > 0$ for any $p \in (0, 1)$.
2. Continuity: $\psi(p)$ is continuous in p.
3. Additivity – information from two independent events is the sum of informations: if $E_1, E_2 \subset \Omega$ are independent, then $\psi\left(\mathbb{P}\{E_1 \cap E_2\}\right) = \psi(\mathbb{P}\{E_1\}) + \psi(\mathbb{P}\{E_2\})$.

How to measure information?

- Then how should our measure $H(\phi)$ of information in belief ϕ look like?
- Say that we derive information $\psi(\mathbb{P}(E))$ from the fact that event $E \subset \Omega$ realized. (Then we could define $H(\phi) \equiv \mathbb{E}\left[\psi(\phi(\omega))\right]$ when Ω is finite.)
- Consider the three following (natural?) axioms:
  1. Positivity: $\psi(p) > 0$ for any $p \in (0, 1)$.
  2. Continuity: $\psi(p)$ is continuous in p.
  3. Additivity – information from two independent events is the sum of informations: if $E_1, E_2 \subset \Omega$ are independent, then $\psi\left(\mathbb{P}\{E_1 \cap E_2\}\right) = \psi(\mathbb{P}\{E_1\}) + \psi(\mathbb{P}\{E_2\})$.
- **Claim:** if the axioms above hold, then $\psi(p) = C \ln(p)$ with $C < 0$.

How to measure information?

- Then how should our measure $H(\phi)$ of information in belief ϕ look like?
- Say that we derive information $\psi(\mathbb{P}(E))$ from the fact that event $E \subset \Omega$ realized. (Then we could define $H(\phi) \equiv \mathbb{E}\left[\psi(\phi(\omega))\right]$ when Ω is finite.)
- Consider the three following (natural?) axioms:
  1. Positivity: $\psi(p) > 0$ for any $p \in (0,1)$.
  2. Continuity: $\psi(p)$ is continuous in p.
  3. Additivity – information from two independent events is the sum of informations: if $E_1, E_2 \subset \Omega$ are independent, then $\psi\left(\mathbb{P}\{E_1 \cap E_2\}\right) = \psi(\mathbb{P}\{E_1\}) + \psi(\mathbb{P}\{E_2\})$.
- **Claim:** if the axioms above hold, then $\psi(p) = C \ln(p)$ with $C < 0$. *Proof:* from additivity, $\forall m, n \in \mathbb{N}$: $\psi(p^n) = n\psi(p)$ and $\psi\left(p^{\frac{1}{m}}\right) = \frac{1}{m}\psi(p)$, so $\psi\left(p^{\frac{n}{m}}\right) = \frac{n}{m}\psi(p)$. By continuity then, $\psi(p^a) = a\psi(p)\ \forall a \in \mathbb{R}$, implying $\psi(p) = \psi(e^{\ln p}) = \psi(e)\ln p = C \ln p$. Positivity implies $C < 0$.

How to measure information? – Entropy

- So with axioms above, $\psi(p) \propto -\ln(p)$ is the amount of information contained in probability p.
- And then $H(\phi) \propto \mathbb{E}\left[\psi(\phi(\omega))\right] = -\sum_\omega \phi(\omega)\ln(\phi(\omega))$ is the amount of information in belief ϕ (if Ω finite; let $0\ln 0 \equiv 0$).
- This $H(\phi)$ is called Shannon entropy, it's a measure of uncertainty (high entropy $=$ high uncertainty).

How to measure information? – Entropy

- So with axioms above, $\psi(p) \propto -\ln(p)$ is the amount of information contained in probability p.
- And then $H(\phi) \propto \mathbb{E}[\psi(\phi(\omega))] = -\sum_{\omega} \phi(\omega) \ln(\phi(\omega))$ is the amount of information in belief ϕ (if Ω finite; let $0 \ln 0 \equiv 0$).
- This $H(\phi)$ is called Shannon entropy, it's a measure of uncertainty (high entropy = high uncertainty).
- We can model cost of signal θ as the expected decrease in entropy:

$$C(\theta) = -\lambda \Big[\mathbb{E}[H(\phi|\omega)] - H(\phi_0) \Big]$$

How to measure information? – Entropy

- So with axioms above, $\psi(p) \propto -\ln(p)$ is the amount of information contained in probability p.
- And then $H(\phi) \propto \mathbb{E}\left[\psi(\phi(\omega))\right] = -\sum_{\omega} \phi(\omega) \ln(\phi(\omega))$ is the amount of information in belief ϕ (if Ω finite; let $0 \ln 0 \equiv 0$).
- This $H(\phi)$ is called Shannon entropy, it's a measure of uncertainty (high entropy = high uncertainty).
- We can model cost of signal θ as the expected decrease in entropy:

$$C(\theta) = -\lambda \Big[\mathbb{E}\left[H(\phi|\omega)\right] - H(\phi_0) \Big]$$

- A more general class of posterior-separable cost functions allows a wide(r) range of possible $\psi(p)$.

This slide deck:

Decision problems with learning

- Consider an agent with some utility $u(x, \theta_0)$, where $\theta_0 \in \Theta_0$ is the payoff-relevant state, Θ_0 finite.
- Before choosing x, agent can choose signal structure θ_1 subject to entropy cost $C(\theta_1)$, so the final payoff is

$$\mathbb{E}\left[u(x(\theta_1), \theta_0)\right] - C(\theta_1)$$

- What are the properties of the optimal signal θ_1?
  - see Maćkowiak et al. [2023] for more details

Optimal learning

- Observation 1: learning any state perfectly (or excluding any state) is infinitely costly, so never optimal $\Rightarrow$ *supp*$(\theta_1(\omega)) = \Theta_0$ for all ω.
- Observation 2: it is never optimal to acquire more than one signal per action. Thus, the optimal signal is an action recommendation.

This slide deck:

Mechanism design with costly learning

- Consider a mechanism design problem with learning:
  - Nature draws $\omega \in \Omega$, unobservable to everyone
  - Designer chooses some mechanism $\Gamma = (A, g)$
  - Each player $i \in \{1, ..., N\}$ privately chooses a signal structure $\theta_i : \Omega \to \Theta_i$
  - Each player i privately observes signal realization $\theta_i = \theta_i(\omega)$ and chooses $a_i \in A_i$
  - Outcome $g(a) \in X$ is realized.
- There is no unified theory yet, so here are some bits and pieces and tips and tricks, following Ravid [2020], Mensch [2022], Larionov and Yamashita [2024]

Revelation principle?

- The first thing we learned in this course is the revelation principle. When types are endogenous, we can't have that. Or can we?
- As usual, assume that designer can select eqm
- Say that designer can *recommend* a signal structure to each player
- Each player i would have an option of reporting one of these "recommended" signal realizations $\theta_i \in \Theta_i$
  - Think that designer offers a menu of possible posterior beliefs $\phi|\theta \in \Delta(\Omega)$ a player can report to the mechanism. These are i's available actions.
  - At the learnng stage, i's optimal signal structure produces at most one signal per action – a "recommendation" of which action to take.
  - In equilibrium, the suggested posteriors are exactly the same as those induced by the buyer's optimal signal.

Incentive compatibility?

- Player i has two kinds of deviations:
  1. Acquire a different signal structure than the suggested one
  2. Misreport their signal realization
- Papers show that IC2 is typically not binding in mechanisms with learning; one mainly needs to worry about IC1.

Individual rationality?

- If player i has an outside option $\underline{U}_i$ from not participating in the mechanism, it is relevant at two points in time:
  1. i may refuse to acquire info and take $\underline{U}_i$ instead
  2. after observing θ_i, i may refuse to report anything to the mechanism and take $\underline{U}_i$ instead.
- Due to the latter, i's signal structure may include a recommendation to "run".

This slide deck:

Screening with costly learning

- Thereze [2023] considers a screening problem, where a buyer can acquire info about their valuation.
- The fundamental buyer's valuation θ_0 is binary: $\theta_0 \in \{\underline{\theta}, \bar{\theta}\}$; producing a good is costly for the seller-designer.

Screening with costly learning

- Thereze [2023] considers a screening problem, where a buyer can acquire info about their valuation.
- The fundamental buyer's valuation θ_0 is binary: $\theta_0 \in \{\underline{\theta}, \bar{\theta}\}$; producing a good is costly for the seller-designer.
- In the optimal mechanism, the buyer acquires a binary signal with posterior beliefs $\theta_1 \in [\underline{\theta}, \bar{\theta}]$.
- Without learning, we had "no distortion at the top": the high-valuation buyer was served for sure (efficient); the low-valuation buyer was served with lower probability (distortion)
- With learning, allocations for both learned-types θ_1 are distorted downwards
- This is because seller must leave more rents to the buyer, to incentivize the buyer to not learn more.

Envelope representation?

- Mensch [2022] considers a Myerson screening problem, where buyer learns their valuation.
- The fundamental buyer's valuation is $\theta_0 \in \Theta_0$, where Θ_0 is finite. Production is free for the seller.

Envelope representation?

- Mensch [2022] considers a Myerson screening problem, where buyer learns their valuation.
- The fundamental buyer's valuation is $\theta_0 \in \Theta_0$, where Θ_0 is finite. Production is free for the seller.

- He derives a sort of an envelope condition for the optimal menu $\{k, t(k), \phi(\cdot|k)\}$, where k is the trading probability, $t(k)$ is the transfer/price, and $\phi(\cdot|k) \in \Delta\Theta_0$ is the buyer's interim belief about their valuation conditional on acquiring info that results in k.

Correlating information

- Larionov and Yamashita [2024] explore a bilateral trade setting, where designer offers a mechanism to a buyer and a seller
- The fundamental product quality is $\theta_0 \in \Theta_0$, where Θ_0 is finite. Players' valuations are arbitrary functions $v_B(\theta_0), v_S(\theta_0)$.

Correlating information

- Larionov and Yamashita [2024] explore a bilateral trade setting, where designer offers a mechanism to a buyer and a seller
- The fundamental product quality is $\theta_0 \in \Theta_0$, where Θ_0 is finite. Players' valuations are arbitrary functions $v_B(\theta_0), v_S(\theta_0)$.
- They show that efficiency requires that buyer & seller acquire perfectly correlated signals about v. The designer can then use sort of a cross-verification mechanism, but must leave rents to the players to incentivize them to learn.
- Larionov et al. [2023] and Jiang and Whitmeyer [2024] make similar points in auction settings.

Conclusion

- Learning is fun
- Accounting for (costly) learning substantially changes the economic predictions of our models
- There's still a lot to explore here!

References I

S. J. Grossman and J. E. Stiglitz. On the impossibility of informationally efficient markets. *The American economic review*, 70(3):393–408, 1980. Publisher: JSTOR.

S. Jiang and M. Whitmeyer. (Almost) full-surplus extraction with endogenous information. mimeo, 2024.

D. Larionov and T. Yamashita. Bilateral Trade with Costly Information Acquisition. mimeo, 2024.

D. Larionov, H. Pham, T. Yamashita, and S. Zhu. First Best Implementation With Costly Information Acquisition. *SSRN Electronic Journal*, 2023. ISSN 1556-5068. doi: 10.2139/ssrn.4321432. URL `https://www.ssrn.com/abstract=4321432`.

B. Maćkowiak, F. Matějka, and M. Wiederholt. Rational Inattention: A Review. *Journal of Economic Literature*, 61(1):226–273, Mar. 2023. ISSN 0022-0515. doi: 10.1257/jel.20211524. URL `https://pubs.aeaweb.org/doi/10.1257/jel.20211524`.

References II

J. Mensch. Screening inattentive buyers. *American Economic Review*, 112(6):1949–1984, 2022. URL https://www.aeaweb.org/articles?id=10.1257/aer.20201098. Publisher: American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203.

D. Ravid. Ultimatum Bargaining with Rational Inattention. *American Economic Review*, 110 (9):2948–2963, Sept. 2020. ISSN 0002-8282. doi: 10.1257/aer.20170620. URL https://www.aeaweb.org/articles?id=10.1257/aer.20170620.

J. Thereze. Screening Costly Information. 2023.