Bayesian Persuasion and Information Design

Dongkyu Chang

City University of Hong Kong

CEBSS Reading Group

Jan 7, 2020

Roadmap

- Bayesian Persuasion (Kamenica and Gentzkow 2011)
- Information Design: A Unified Perspective (Bergemann and Morris 2019)

Bayesian Persuasion

Motivating Example

- Prosecutor (Sender) vs Judge (Receiver)
- The judge must choose $a \in A = \{\text{acquit}, \text{convict}\}$.
- There are two possible states of the world $\Theta = \{\text{guilty}, \text{innocent}\}$.
- The judge gets utility 1 for choosing the just action and 0 otherwise.
- The prosecutor gets utility 1 if the judge convicts and 0 otherwise.
- Common prior: $\Pr(\text{guilty}) = 0.3$.

Motivating Example

- The prosecutor conducts an investigation and is required by law to report its full outcome.
- Formalize an investigation as distributions $\pi(\cdot|\text{guility})$ and $\pi(\cdot|\text{innocent})$ on some set of signal realizations.
- The prosecutor "designs" π and must honestly report the signal realization to the judge.

Motivating Example

- With a fully informative π, $a =$ convict 30% of the time.
- The prosecutor can do better by adding some noise to π.

$$\pi(i|\text{innocent}) = \frac{4}{7} \quad \pi(i|\text{guilty}) = 0$$
$$\pi(g|\text{innocent}) = \frac{3}{7} \quad \pi(g|\text{guilty}) = 1$$

$$P(\text{innocent}|i) = \frac{0.7\times\pi(i|\text{innocent})}{0.3\times\pi(i|\text{guilty})+0.7\times\pi(i|\text{innocent})} = 1 \quad \Longrightarrow \quad a = \text{acquit}$$
$$P(\text{innocent}|g) = \frac{0.7\times\pi(g|\text{innocent})}{0.3\times\pi(g|\text{guilty})+0.7\times\pi(g|\text{innocent})} = \frac{1}{2} \quad \Longrightarrow \quad a = \text{convict}$$

- The judge convicts with probability 60 percent (ex ante).

Setup

- Receiver has a continuous utility function $u(a, \theta)$ that depends on her action $a \in A$ (compact) and the state $\theta \in \Theta$ (finite).
- Sender has a continuous utility function $v(a, \theta)$.
- Common prior $\mu_0 \in \text{int}(\Delta\Theta)$.

Setup

- Sender commits to a signal $\pi = \langle T, \{\pi(\cdot|\theta)\}_{\theta \in \Theta} \rangle$ *before* θ is realized.
  - T: *finite* set of possible realizations of signal.
  - $\{\pi(\cdot|\theta)\}_{\theta \in \Theta}$: family of distributions. $\pi(t|\theta) = \Pr\{t \text{ is realized}|\theta\}$.
- Receiver observes Sender's choice of π and a signal realization $t \in T$, forms the posterior μ_t using Bayes' rule, and then takes action $\hat{a}(\mu_t)$.
- Taking Receiver's behavior as given, Sender chooses a signal π that maximizes his expected utility
- Sender preferred PBE: If Receiver has multiple optimal actions, she chooses one that maximizes Sender's expected payoff.

Setup

- Each realization $t \in T$ leads to a posterior belief $\mu_t \in \Delta(\Theta)$.

$$\mu_t(\theta) = \frac{\pi(t|\theta)\mu_0(\theta)}{\sum_{\theta' \in \Theta} \pi(t|\theta')\mu_0(\theta')} \quad \forall t, \theta$$

  and Sender's expected payoff

$$\hat{v}(\mu_t) = \mathbb{E}_{\mu_t}\left[v(\hat{a}(\mu_t), \theta)\right]$$

- Each signal leads to a distribution over posterior beliefs $\tau \in \Delta(\Delta(\Theta))$.

$$\tau(\mu) = \sum_{t:\mu_t=\mu} \sum_{\theta' \in \Theta} \pi(t|\theta')\mu_0(\theta') \quad \forall \mu \in \text{supp}(\tau).$$

- By a direct computation,

$$\sum_{\mu \in \text{supp}(\tau)} \mu(\theta)\tau(\mu) = \mu_0(\theta) \quad \forall \theta \qquad \text{(Bayes plausibility)}$$

Equilibrium

- Equilibrium:

$$\max_{\pi\in\Pi} \underbrace{\sum_{\theta\in\Theta}\left[\sum_{t\in T}\hat{v}(\mu_t)\pi(t|\theta)\right]\mu_0(\theta)}_{=\text{value of signal } \pi}$$

- Π: set of signals with finite signal realizations (i.e., $|T| < \infty$)
- A signal is called straightforward if $T = A$ and $\hat{a}(\mu_t) = t$ for any $t \in T = A$. i.e., Sender recommends action and then Receiver obeys.
- By the revelation principle, we may focus on straightforward signals.

Equilibrium

Proposition

The following are equivalent:

1. There exists a signal with value v^*;
2. There exists a straightforward signal with value v^*;
3. There exists a Bayes-plausible distribution of posteriors τ such that

$$\mathbb{E}_\tau[\hat{v}(\mu)] = v^*.$$

Equilibrium

Corollary

Sender benefits from persuasion iff there exists a Bayes-plausible distribution of posteriors τ such that

$$\mathbb{E}_\tau[\hat{v}(\mu)] > \hat{v}(\mu_0)$$

The value of an optimal signal is

$$\max_\tau \ \mathbb{E}_\tau[\hat{v}(\mu)] = \sum_\mu \hat{v}(\mu)\tau(\mu) \quad \textit{subject to} \quad \sum_{\mu \in supp(\tau)} \mu\tau(\mu) = \mu_0$$

Concavication

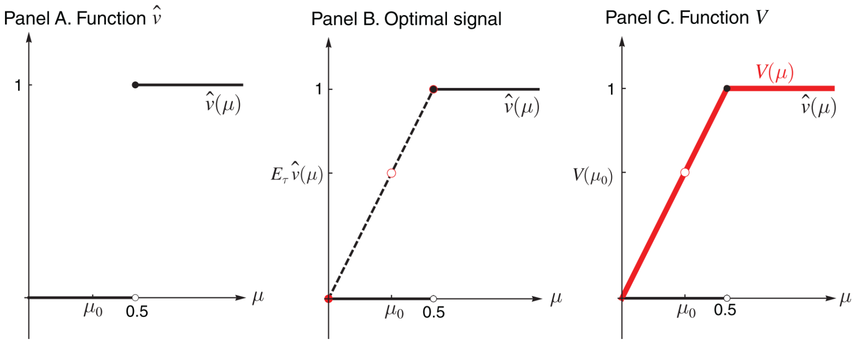

FIGURE 2. THE MOTIVATING EXAMPLE

Information Design: A Unified Perspective

Setup

- Θ: Finite set of states.
- Basic game $G = ((A_i, u_i)_{i=1}^I, \psi)$ w/ a common prior $\psi \in \Delta_{++}(\Theta)$.
- Information Structure $S = ((T_i)_{i=1}^I, \pi)$
  - T_i: a finite set of signals (or types); $T = T_1 \times \ldots \times T_I$.
  - $\pi : \Theta \to \Delta(T)$: signal distribution (or type distribution).
- Example: T_i is a singleton for each t_i (null prior information).
- Bayes Nash Equilibrium: A strategy profile $(\beta_1, \ldots, \beta_I)$ is a BNE of (G, S) if for each $i \in I$, $t_i \in T_i$, and $a_i \in A_i$ with $\beta_i(a_i|t_i) > 0$, we have

$$a_i \in \arg\max_{a_i' \in A_i} \sum_{t_{-i}, a_{-i}, \theta} \psi(\theta)\pi(t_i, t_{-i}|\theta)\beta_{-i}(a_{-i}|t_{-i})u_i(a_i', a_{-i}, \theta)$$

Correlated Equilibrium (with Complete Information)

- For now suppose $\Theta = \{\theta_\varnothing\}$. Hence, there is no incomplete information.
- Signals have no information content. However, players may use signals as a correlation device.
- Now suppose there is a mediator.

1. $\theta = \theta_\varnothing$ is (trivially) realized.
2. t is realized and each player i observes t_i.
3. A mediator recommends an action for each player according to a decision rule $\sigma : T \to \Delta(A)$.
4. Each player i chooses an action.

Correlated Equilibrium (with Complete Information)

- $\sigma : T \to \Delta(A)$ is *obedient* if for each i, $t_i \in T_i$, and $a_i \in A_i$

$$a_i \in \arg\max_{a_i' \in A_i} \sum_{a_{-i}} \sum_{t_{-i}} \pi(t_i, t_{-i}|\theta_\varnothing)\sigma(a_i, a_{-i}|t_i, t_{-i})u_i(a_i', a_{-i}, \theta).$$

- Any correlated eqm can be implementable w/ an σ.

Correlated Equilibrium (with Incomplete Information)

- From now on, suppose $|\Theta| \geq 2$.
- Now signals may have information content (about the true state).
- Several alternatives:
  - The mediator observes t and θ directly (omniscient mediator);
  - The mediator does not observe t but she can elicit t_i from player i;
  - The mediator does not observe t and cannot elicit t from players, etc.

Bayes Correlated Equilibrium (Bergemann and Morris, 2016)

- Decision rule:

$$\sigma : T \times \Theta \to \Delta(A).$$

- A decision rule is *obedient* for (G, S) if for each i, $t_i \in T_i$ and $a_i \in A_i$,

$$a_i \in \arg\max_{a_i' \in A_i} \sum_{a_{-i}, t_{-i}, \theta} \frac{\psi(\theta)\pi(t_i, t_{-i}|\theta)\sigma(a_i, a_{-i}|t_i, t_{-i}, \theta)}{P(a_i, t_i)} u_i(a_i', a_{-i}, \theta).$$

- $P(a_i, t_i) = \sum_{\tilde{a}_{-i}, \tilde{t}_{-i}, \tilde{\theta}} \psi(\tilde{\theta})\pi(t_i, \tilde{t}_{-i}|\tilde{\theta})\sigma(a_i, \tilde{a}_{-i}|t_i, \tilde{t}_{-i}, \tilde{\theta}) = \text{(constant)}.$
- A decision rule σ is a Bayes correlated equilibrium (BCE) of (G, S) if it is obedient for (G, S).
- The notion of BCE suits to situations with an omniscient mediator.

Information Design Problem

- Now consider the problem of information design.
- Timeline:

1. The information designer picks and commits to a rule for providing the players with extra messages (π and σ).
2. The true state θ is realized, and each player's type t_i is privately realized.
3. The players receive extra messages according to the info designer's rule.
4. The players pick their actions based on their prior information and the messages from the information designer, and then payoffs are realized.

Information Design with Omniscient Information Designer

- Suppose the information designer is omniscient:

1. the information designer can observe the realization of θ and t, thus
2. the information designer's messages can be contingent on the realization of θ and t.

- Revelation Principle: We may assume (i) the information designer's messages directly recommend each player which action to take, and (ii) all players obey these recommendations.

- Given this restriction, and also the assumption of the designer's omniscience, the information designer is choosing a decision rule

$$\sigma : T \times \Theta \to \Delta(A).$$

Information Design with Omniscient Information Designer

Proposition (Revelation Principle)

An omniscient information designer can attain decision rule σ if and only if it is a BCE, i.e., if it satisfies obedience.

Investment Example

- There is a bad state (B) and a good state (G). The two states are equally likely: $\psi(G) = \psi(B) = \frac{1}{2}$.
- There is one player (the "firm") who can decide to *invest* or *not invest*.

$u(a, \theta)$	B	G
invest	-1	x
not invest	0	0

where $x \in (0, 1)$.

- The information designer (the "government") is interested in maximizing the probability of investment independent of the state.

$$1 = v(\text{invest}, \theta) > v(\text{not invest}, \theta) = 0 \quad \theta = B, G.$$

Investment Example (with Single Player and No Prior Information)

- First, suppose that the firm has no prior information (T is a singleton).
- $\sigma(\theta)$ specifies the prob of investment, denoted by p_θ, conditional on θ.
- Obedience constraints:

$$p_B - p_G x \leq 0 \leq (1 - p_B) - (1 - p_G)x \iff \tfrac{1-p_B}{1-p_G} \geq x \geq \tfrac{p_B}{p_G}$$

- Information designer solves the linear programming:

$$\max\ p_G + p_B \quad \text{s.t.} \quad \text{obedience constraints and } 0 \leq p_G, p_B \leq 1.$$

Investment Example (with Single Player and No Prior Information)

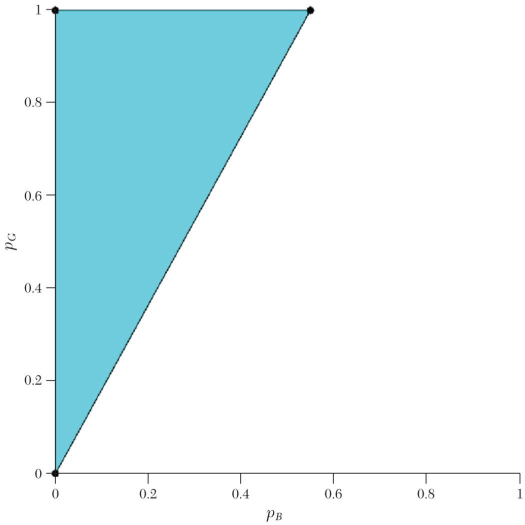

Figure 1. Investment Probability with Uninformed Player: $x = 55/100$

Investment Example (with Single Player and Prior Information)

- Next, suppose the firm receives a "correct" signal with prob $q > 1/2$.
- $\sigma(\theta, t)$ specifies the prob of investment $p_{\theta t}$ conditional on (θ, t).
- Obedience constraints:

$$
\begin{aligned}
qp_{Gg}x + (1-q)p_{Bg}(-1) &\geq 0 \quad \text{for } (t,a) = (g, \text{invest}) \\
(1-q)p_{Gb}x + qp_{Bb}(-1) &\geq 0 \quad \text{for } (t,a) = (b, \text{invest}) \\
q(1-p_{Gg})x + (1-q)(1-p_{Bg})(-1) &\leq 0 \quad \text{for } (t,a) = (g, \text{not invest}) \\
(1-q)(1-p_{Gb})x + q(1-p_{Bb})(-1) &\leq 0 \quad \text{for } (t,a) = (b, \text{not invest})
\end{aligned}
$$

- Information designer solves the linear programming:

$$
\max_{p_{Gg}, p_{Gb}, p_{Bg}, p_{Bb}} \quad qp_{Gg} + (1-q)p_{Gb} + (1-q)p_{Bg} + qp_{Bb}
$$

$$
\text{s.t. obedience constraints and } p_{Gg}, p_{Gb}, p_{Bg}, p_{Bb} \in [0,1].
$$

Investment Example (with Single Player and Prior Information)

- We are interested in *ex ante* investment probabilities

$$p_G = qp_{Gg} + (1-q)p_{Gb} \quad \text{and} \quad p_B = (1-q)p_{Bg} + qp_{Bb}.$$

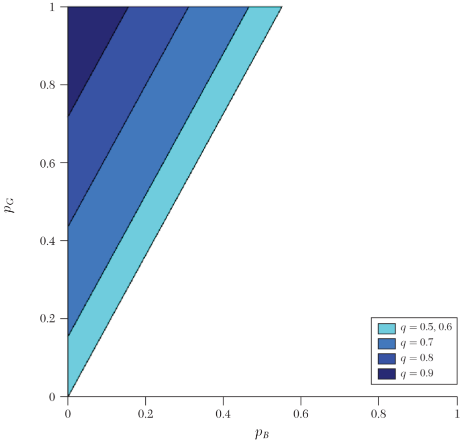

Figure 2. Investment Probability with Informed Player: $x = 55/100$

Investment Example (with Two Players and No Prior Information)

- There are two firms who simultaneously decide whether to invest. Firm 1's payoffs are given as follows (and symmetrically for firm 2):

$\theta = B$	invest	not
invest	$-1 + \epsilon$	-1
not	0	0

$\theta = G$	invest	not
invest	$x + \epsilon$	x
not	0	0

- The government wants to maximize the sum over each individual firm's probability of investment.
- We can focus on symmetric decision rules, given the symmetry of the basic game, for any symmetric objective of the information designer.
- $p_\theta = P\{\text{each firm invests}\}$ and $r_\theta = P\{\text{both invest}\}$ in each state θ.

Investment Example (with Two Players and No Prior Information)

- Obedience constraints

$$0 \le \tfrac{1}{2}(p_B - r_B)(-1) + \tfrac{1}{2}(p_G - r_G)x + \tfrac{1}{2}r_B(-1+\epsilon) + \tfrac{1}{2}r_G(x+\epsilon)$$

$$\begin{aligned} 0 \ge\ & \tfrac{1}{2}(p_B - r_B)(-1+\epsilon) + \tfrac{1}{2}(p_G - r_G)(x+\epsilon) \\ & + \tfrac{1}{2}(1 + r_B - 2p_B)(-1) + \tfrac{1}{2}(1 + r_G - 2p_G)x \end{aligned}$$

- Feasibility constraint: $\max\{0, 2p_\theta - 1\} \le r_\theta \le p_\theta$ for both $\theta = G, B$.
- Again, the optimal decision rule can be identified by solving a linear programming.

Investment Example (with Two Players and No Prior Information)

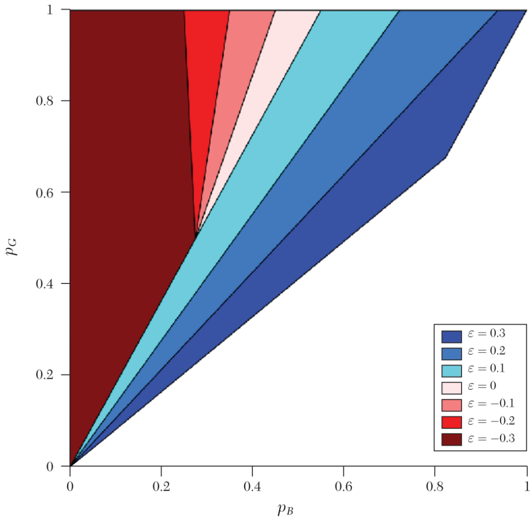

Figure 3. Investment Probability with Negative or Positive Strategic Term ε

Investment Example (with Two Players and No Prior Information)

- With strategic complementarity ($\epsilon > 0$) the optimal decision rule is

$\theta = B$	invest	not
invest	$\frac{x+\epsilon}{1-\epsilon}$	0
not	0	$\frac{1-x-2\epsilon}{1-\epsilon}$

$\theta = G$	invest	not
invest	1	0
not	0	0

and

$$p_G = r_G = 1 \text{ and } p_B = r_B = \frac{x+\epsilon}{1-\epsilon}.$$

- The optimal decision rule entails public messages.

Investment Example (with Two Players and No Prior Information)

- With strategic substitutes ($\epsilon < 0$) and the extra assumptions $x > 1/2$ and $|\epsilon| \leq x - 1/2$, the optimal decision rule is

$\theta = B$	invest	not
invest	0	$x + \epsilon$
not	$x + \epsilon$	$1 - 2x - 2\epsilon$

$\theta = G$	invest	not
invest	1	0
not	0	0

- The optimal decision rule entails private messages.

Information Design with Private Information

- We have thus far assumed the designer knows not only the true state θ but also the players' prior information about the state.
- What if the designer cannot condition recommendations on players' prior information?
  - Information design with elicitation
  - Information design without elicitation

Information Design with Private Information and Elicitation

- Suppose the decision maker can elicit the players' private information.
- We will now require an incentive compatibility condition that entails both *truth telling* as well as *obedience.*
- A decision rule $\sigma : T \times \Theta \to \Delta(A)$ is *incentive compatible* for (G, S) if for each i and $t_i \in T_i$

$$\sum_{a_{-i}, t_{-i}, \theta} \psi(\theta)\pi(t_i, t_{-i}|\theta)\sigma(a_i, a_{-i}|t_i, t_{-i}, \theta)u_i(a_i, a_{-i}, \theta)$$
$$\geq \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta)\pi(t_i, t_{-i}|\theta)\sigma(\delta(a_i), a_{-i}|t_i', t_{-i}, \theta)u_i(a_i, a_{-i}, \theta)$$

for all $t_i' \in T$ and $\delta_i : A_i \to A_i$.

- Prevent "double deviations."

Information Design with Private Information but without Elicitation

Proposition

An information designer with elicitation can attain a decision rule if and only if it is incentive compatible.

Information Design with Private Information but without Elicitation

- Suppose the decision maker cannot elicit private information.
- The designer has to offer each player a *contingent* recommendation, a *vector* of action recommendations, where each entry is an action recommendation for a specific type.
- The set of feasible recommendations to player i is given by $B_i = A_i^{T_i}$. Also, define $B = \prod_{i=1}^{I} B_i$.
- $\sigma : T \times \Theta \to \Delta(A)$ is *publicly feasible* if there exists a contingent recommendation $\phi : \Theta \to \Delta(B)$ s.t. for each $a \in A$, $t \in T$, and $\theta \in \Theta$ with $\pi(t|\theta) > 0$,

$$\sigma(a|t,\theta) = \sum_{b \in B : b(t) = a} \phi(b|\theta).$$

In this case, we say that σ is induced by ϕ.

- $\sigma : T \times \Theta \to \Delta(A)$ is *publicly feasible obedient* if there exists a contingent recommendation $\phi : \Theta \to \Delta(B)$ such that (i) ϕ induces σ, and (ii) ϕ satisfies obedience in the sense that for each i, $t_i \in T_i$ and $b_i \in B_i$,

$$\sum_{a_{-i}, t_{-i}, \theta} \psi(\theta)\phi(b_i, b_{-i}|\theta)\pi(t_i, t_{-i}|\theta)u_i(b_i(t_i), b_{-i}(t_{-i}), \theta)$$
$$\geq \sum_{a_{-i}, t_{-i}, \theta} \psi(\theta)\phi(b_i, b_{-i}|\theta)\pi(t_i, t_{-i}|\theta)u_i(a', b_{-i}(t_{-i}), \theta) \quad \forall a_i' \in A_i.$$

Proposition

An information designer without elicitation can attain a decision rule if and only if it is publicly feasible obedient.

Information Design with Private Information

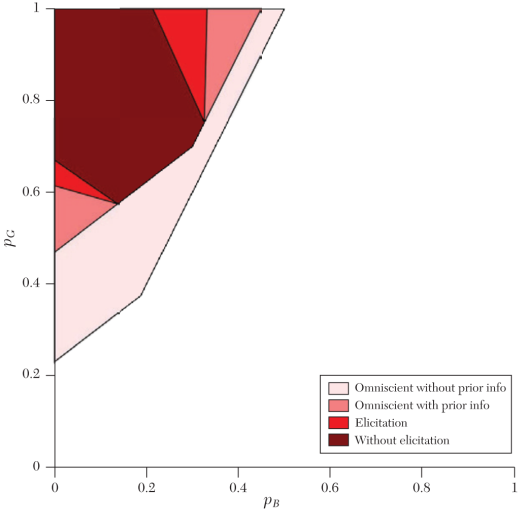

Figure 7. Investment Probability under Different Information Design Scenarios