Preparing for the Worst but Hoping for the Best: Robust (Bayesian) Persuasion

Piotr Dworczak Alessandro Pavan

May 28, 2020

Quick introduction to Bayesian persuasion

Kamenica and Gentzkow (**AER, 2011**, > 1300 citations)

Quick introduction to Bayesian persuasion

Kamenica and Gentzkow (**AER, 2011**, > 1300 citations)

- A game between a Sender and a Receiver;
- State $\omega \in \Omega$ (finite), distributed according to a common prior $\mu_0 \in \Delta\Omega$;
- The Sender **commits** to a signal $q : \Omega \to \Delta(S)$;
- The Receiver observes $s \in S$, updates beliefs to μ_0^s according to Bayes' rule, and takes an optimal action

$$a^\star(\mu_0^s) \in \text{argmax}_{a \in A} \mathbb{E}_{\omega \sim \mu_0^s}[u(a, \omega)].$$

- The Sender selects q to maximize

$$\mathbb{E}_{\omega \sim \mu_0} \mathbb{E}_{s \sim q(\omega)}[v(a^\star(\mu_0^s), \omega)].$$

Quick introduction to Bayesian persuasion

The judge example:

- Sender=Prosecutor, Receiver=Judge, (Suspect)

Quick introduction to Bayesian persuasion

The judge example:

- Sender=Prosecutor, Receiver=Judge, (Suspect)
- $\Omega = \{G, I\}$, $\mu_0 = 1/2$.

Quick introduction to Bayesian persuasion

The judge example:

- Sender=Prosecutor, Receiver=Judge, (Suspect)
- $\Omega = \{G, I\}$, $\mu_0 = 1/2$.
- $A = \{convict, acquit\}$.

Quick introduction to Bayesian persuasion

The judge example:

- Sender=Prosecutor, Receiver=Judge, (Suspect)
- $\Omega = \{G, I\}$, $\mu_0 = 1/2$.
- $A = \{convict, acquit\}$.
- Judge's payoff: $u(acquit, \omega) = 0$, and

$$u(convict, \omega) = \begin{cases} -2, & \omega = I, \\ 1, & \omega = G. \end{cases}$$

Quick introduction to Bayesian persuasion

The judge example:

- Sender=Prosecutor, Receiver=Judge, (Suspect)
- $\Omega = \{G, I\}$, $\mu_0 = 1/2$.
- $A = \{convict, acquit\}$.
- Judge's payoff: $u(acquit, \omega) = 0$, and

$$u(convict, \omega) = \begin{cases} -2, & \omega = I, \\ 1, & \omega = G. \end{cases}$$

- Thus, Judge convicts if she believes the Suspect to be guilty with probability $2/3$ or more.

Quick introduction to Bayesian persuasion

The judge example:

- Sender=Prosecutor, Receiver=Judge, (Suspect)
- $\Omega = \{G, I\}$, $\mu_0 = 1/2$.
- $A = \{convict, acquit\}$.
- Judge's payoff: $u(acquit, \omega) = 0$, and

$$u(convict, \omega) = \begin{cases} -2, & \omega = I, \\ 1, & \omega = G. \end{cases}$$

- Thus, Judge convicts if she believes the Suspect to be guilty with probability $2/3$ or more.
- Prosecutor's payoff:

$$v(a, \omega) = \begin{cases} 1, & a = convict, \\ 0, & a = acquit. \end{cases}$$

Motivation

- Bayesian persuasion/ information design

Motivation

- Bayesian persuasion/ information design
  - designer knows agents' sources of information

Motivation

- Bayesian persuasion/ information design
  - designer knows agents' sources of information
  - trusts her ability to coordinate Receivers on actions most favorable to her

Motivation

- Bayesian persuasion/ information design
  - designer knows agents' sources of information
  - trusts her ability to coordinate Receivers on actions most favorable to her
  - optimal information structure **sensitive to fine details of agents' beliefs**

Motivation

- Bayesian persuasion/ information design
  - designer knows agents' sources of information
  - trusts her ability to coordinate Receivers on actions most favorable to her
  - optimal information structure **sensitive to fine details of agents' beliefs**
- In many problems of interest,

Motivation

- Bayesian persuasion/ information design
  - designer knows agents' sources of information
  - trusts her ability to coordinate Receivers on actions most favorable to her
  - optimal information structure **sensitive to fine details of agents' beliefs**
- In many problems of interest,
  - agents' sources of information (both before and after receiving Sender's information) unknown

Motivation

- Bayesian persuasion/ information design
  - designer knows agents' sources of information
  - trusts her ability to coordinate Receivers on actions most favorable to her
  - optimal information structure **sensitive to fine details of agents' beliefs**
- In many problems of interest,
  - agents' sources of information (both before and after receiving Sender's information) unknown
  - Sender may not trust her ability to coordinate Receivers

Motivation

- Bayesian persuasion/ information design
  - designer knows agents' sources of information
  - trusts her ability to coordinate Receivers on actions most favorable to her
  - optimal information structure **sensitive to fine details of agents' beliefs**
- In many problems of interest,
  - agents' sources of information (both before and after receiving Sender's information) unknown
  - Sender may not trust her ability to coordinate Receivers
- Quest for **robustness**

This Paper

- New solution concept that accounts for such uncertainty/ambiguity

This Paper

- New solution concept that accounts for such uncertainty/ambiguity
- Lexicographic approach to the problem

This Paper

- New solution concept that accounts for such uncertainty/ambiguity
- Lexicographic approach to the problem
  - Step 1: "Preparing for the worst"

This Paper

- New solution concept that accounts for such uncertainty/ambiguity
- Lexicographic approach to the problem
  - Step 1: "Preparing for the worst"
    - designer seeks to protect herself against possibility that Nature provides information and coordinates agents on actions to minimize the designer's payoff

This Paper

- New solution concept that accounts for such uncertainty/ambiguity
- Lexicographic approach to the problem
  - Step 1: "Preparing for the worst"
    - designer seeks to protect herself against possibility that Nature provides information and coordinates agents on actions to minimize the designer's payoff
  - Step 2: "Hoping for the best"

This Paper

- New solution concept that accounts for such uncertainty/ambiguity
- Lexicographic approach to the problem
  - Step 1: "Preparing for the worst"
    - designer seeks to protect herself against possibility that Nature provides information and coordinates agents on actions to minimize the designer's payoff
  - Step 2: "Hoping for the best"
    - designer maximizes over all worst-case optimal policies assuming Nature and Receivers play favorably to her

This Paper

- New solution concept that accounts for such uncertainty/ambiguity
- Lexicographic approach to the problem
  - Step 1: "Preparing for the worst"
    - designer seeks to protect herself against possibility that Nature provides information and coordinates agents on actions to minimize the designer's payoff
  - Step 2: "Hoping for the best"
    - designer maximizes over all worst-case optimal policies assuming Nature and Receivers play favorably to her
- **Robust solutions**

This Paper

- New solution concept that accounts for such uncertainty/ambiguity
- Lexicographic approach to the problem
  - Step 1: "Preparing for the worst"
    - designer seeks to protect herself against possibility that Nature provides information and coordinates agents on actions to minimize the designer's payoff
  - Step 2: "Hoping for the best"
    - designer maximizes over all worst-case optimal policies assuming Nature and Receivers play favorably to her
- **Robust solutions**
  - best-case optimal among worst-case optimal ones

Results

- Separation theorem – general characterization
- Properties of robust solutions
- Implications for various persuasion models – applications
- Equivalence to the weighted objective model

Literature

- **Bayesian persuasion**
  - ...Calzolari and Pavan (2006), Brocas and Carillo (2007), Rayo-Segal (2010), Kamenica and Gentzkow (2011), Ely (2017),...
  - Surveys
    - Bergemann and Morris (2019)
    - Kamenica (2019)
- **Information design with adversarial coordination**
  - Inostroza and Pavan (2018)
  - Mathevet, Perego, Taneva (2019)
  - Morris et al. (2019)
  - Ziegler (2019)
- **Persuasion with unknown beliefs**
  - Kolotilin et al. (2017)
  - Laclau and Renou (2017)
  - Guo and Schmaya (2018)
  - Hu and Weng (2019)
  - Kosterina (2019)
- **Max-max over max-min design**
  - Borgers (2017)

Plan

1. Model
2. Robust Solutions
3. Separation Theorem
4. Properties of Robust Solutions
5. Weighted Objective
6. Applications
7. Conditionally-independent Robust Solutions (extension)

Model

Model: Environment

- Payoff-relevant state: $\omega \in \Omega$ (finite)

Model: Environment

- Payoff-relevant state: $\omega \in \Omega$ (finite)
- (Common) Prior: $\mu_0 \in \Delta\Omega$

Model: Environment

- Payoff-relevant state: $\omega \in \Omega$ (finite)
- (Common) Prior: $\mu_0 \in \Delta\Omega$
- Sender's "signal"

Model: Environment

- Payoff-relevant state: $\omega \in \Omega$ (finite)
- (Common) Prior: $\mu_0 \in \Delta\Omega$
- Sender's "signal"
  - $q : \Omega \to \Delta\mathcal{S}$

Model: Environment

- Payoff-relevant state: $\omega \in \Omega$ (finite)
- (Common) Prior: $\mu_0 \in \Delta\Omega$
- Sender's "signal"
  - $q : \Omega \to \Delta\mathcal{S}$
  - $\mathcal{S}$: signal realizations

Model: Environment

- Payoff-relevant state: $\omega \in \Omega$ (finite)
- (Common) Prior: $\mu_0 \in \Delta\Omega$
- Sender's "signal"
  - $q : \Omega \to \Delta\mathcal{S}$
  - $\mathcal{S}$: signal realizations
- (Reduced-form description of) Sender's payoff, given induced posterior $\mu \in \Delta\Omega$

Model: Environment

- Payoff-relevant state: $\omega \in \Omega$ (finite)
- (Common) Prior: $\mu_0 \in \Delta\Omega$
- Sender's "signal"
  - $q : \Omega \to \Delta\mathcal{S}$
  - $\mathcal{S}$: signal realizations
- (Reduced-form description of) Sender's payoff, given induced posterior $\mu \in \Delta\Omega$
  - $\overline{V}(\mu)$: "highest" payoff; an u.s.c. function

Model: Environment

- Payoff-relevant state: $\omega \in \Omega$ (finite)
- (Common) Prior: $\mu_0 \in \Delta\Omega$
- Sender's "signal"
  - $q : \Omega \to \Delta\mathcal{S}$
  - $\mathcal{S}$: signal realizations
- (Reduced-form description of) Sender's payoff, given induced posterior $\mu \in \Delta\Omega$
  - $\overline{V}(\mu)$: "highest" payoff; an u.s.c. function
  - $\underline{V}(\mu)$: "lowest" payoff; a l.s.c. function

Model: Environment

- Payoff-relevant state: $\omega \in \Omega$ (finite)
- (Common) Prior: $\mu_0 \in \Delta\Omega$
- Sender's "signal"
  - $q : \Omega \to \Delta\mathcal{S}$
  - $\mathcal{S}$: signal realizations
- (Reduced-form description of) Sender's payoff, given induced posterior $\mu \in \Delta\Omega$
  - $\overline{V}(\mu)$: "highest" payoff; an u.s.c. function
  - $\underline{V}(\mu)$: "lowest" payoff; a l.s.c. function
- Difference between $\overline{V}$ and $\underline{V}$:

Model: Environment

- Payoff-relevant state: $\omega \in \Omega$ (finite)
- (Common) Prior: $\mu_0 \in \Delta\Omega$
- Sender's "signal"
  - $q : \Omega \to \Delta\mathcal{S}$
  - $\mathcal{S}$: signal realizations
- (Reduced-form description of) Sender's payoff, given induced posterior $\mu \in \Delta\Omega$
  - $\overline{V}(\mu)$: "highest" payoff; an u.s.c. function
  - $\underline{V}(\mu)$: "lowest" payoff; a l.s.c. function
- Difference between $\overline{V}$ and $\underline{V}$:
  - equilibrium selection (multiple Receivers)

Model: Environment

- Payoff-relevant state: $\omega \in \Omega$ (finite)
- (Common) Prior: $\mu_0 \in \Delta\Omega$
- Sender's "signal"
  - $q : \Omega \to \Delta\mathcal{S}$
  - $\mathcal{S}$: signal realizations
- (Reduced-form description of) Sender's payoff, given induced posterior $\mu \in \Delta\Omega$
  - $\overline{V}(\mu)$: "highest" payoff; an u.s.c. function
  - $\underline{V}(\mu)$: "lowest" payoff; a l.s.c. function
- Difference between $\overline{V}$ and $\underline{V}$:
  - equilibrium selection (multiple Receivers)
  - tie-breaking (single Receiver)

Model: Environment

- Payoff-relevant state: $\omega \in \Omega$ (finite)
- (Common) Prior: $\mu_0 \in \Delta\Omega$
- Sender's "signal"
  - $q : \Omega \to \Delta\mathcal{S}$
  - $\mathcal{S}$: signal realizations
- (Reduced-form description of) Sender's payoff, given induced posterior $\mu \in \Delta\Omega$
  - $\overline{V}(\mu)$: "highest" payoff; an u.s.c. function
  - $\underline{V}(\mu)$: "lowest" payoff; a l.s.c. function
- Difference between $\overline{V}$ and $\underline{V}$:
  - equilibrium selection (multiple Receivers)
  - tie-breaking (single Receiver)
  - or something else entirely! (formally, we don't even require $\overline{V} \geq \underline{V}$)

Model: Sender's uncertainty

- Nature designs information structure

Model: Sender's uncertainty

- Nature designs information structure
  - $\pi : \Omega \times \mathcal{S} \to \Delta\mathcal{R}$

Model: Sender's uncertainty

- Nature designs information structure
  - $\pi : \Omega \times \mathcal{S} \to \Delta\mathcal{R}$
  - $\mathcal{R}$: signal realizations

Model: Sender's uncertainty

- Nature designs information structure
  - $\pi : \Omega \times \mathcal{S} \to \Delta\mathcal{R}$
  - $\mathcal{R}$: signal realizations
- Interpretation

Model: Sender's uncertainty

- Nature designs information structure
  - $\pi : \Omega \times \mathcal{S} \to \Delta\mathcal{R}$
  - $\mathcal{R}$: signal realizations
- Interpretation
  - The Sender does not know the distribution of Receivers' beliefs

Model: Sender's uncertainty

- Nature designs information structure
  - $\pi : \Omega \times \mathcal{S} \to \Delta\mathcal{R}$
  - $\mathcal{R}$: signal realizations
- Interpretation
  - The Sender does not know the distribution of Receivers' beliefs
  - Receivers could acquire additional information from other sources **after** seeing the Sender's signal

Model: Sender's uncertainty

- Nature designs information structure
  - $\pi : \Omega \times \mathcal{S} \to \Delta\mathcal{R}$
  - $\mathcal{R}$: signal realizations
- Interpretation
  - The Sender does not know the distribution of Receivers' beliefs
  - Receivers could acquire additional information from other sources **after** seeing the Sender's signal
  - Correlated noise (maximal concern for robustness)

Model: Sender's uncertainty

- Nature designs information structure
  - $\pi : \Omega \times \mathcal{S} \to \Delta\mathcal{R}$
  - $\mathcal{R}$: signal realizations
- Interpretation
  - The Sender does not know the distribution of Receivers' beliefs
  - Receivers could acquire additional information from other sources **after** seeing the Sender's signal
  - Correlated noise (maximal concern for robustness)
- Multiple Receivers

Model: Sender's uncertainty

- Nature designs information structure
  - $\pi : \Omega \times \mathcal{S} \to \Delta\mathcal{R}$
  - $\mathcal{R}$: signal realizations
- Interpretation
  - The Sender does not know the distribution of Receivers' beliefs
  - Receivers could acquire additional information from other sources **after** seeing the Sender's signal
  - Correlated noise (maximal concern for robustness)
- Multiple Receivers
  - discriminatory disclosures embedded into derivation of $\underline{V}(\mu)$

Model: Sender's uncertainty

- Nature designs information structure
  - $\pi : \Omega \times \mathcal{S} \to \Delta\mathcal{R}$
  - $\mathcal{R}$: signal realizations
- Interpretation
  - The Sender does not know the distribution of Receivers' beliefs
  - Receivers could acquire additional information from other sources **after** seeing the Sender's signal
  - Correlated noise (maximal concern for robustness)
- Multiple Receivers
  - discriminatory disclosures embedded into derivation of $\underline{V}(\mu)$
  - given common posterior μ, Nature provides (possibly private) signals to the agents and coordinates them on course of action most adversarial to Sender (among those consistent with assumed solution concept)

Model: Sender's uncertainty

- Nature designs information structure
  - $\pi : \Omega \times \mathcal{S} \to \Delta\mathcal{R}$
  - $\mathcal{R}$: signal realizations
- Interpretation
  - The Sender does not know the distribution of Receivers' beliefs
  - Receivers could acquire additional information from other sources **after** seeing the Sender's signal
  - Correlated noise (maximal concern for robustness)
- Multiple Receivers
  - discriminatory disclosures embedded into derivation of $\underline{V}(\mu)$
  - given common posterior μ, Nature provides (possibly private) signals to the agents and coordinates them on course of action most adversarial to Sender (among those consistent with assumed solution concept)
  - e.g., Bayes-correlated eq. given μ

Model: Sender's uncertainty

- Nature designs information structure
  - $\pi : \Omega \times \mathcal{S} \to \Delta\mathcal{R}$
  - $\mathcal{R}$: signal realizations
- Interpretation
  - The Sender does not know the distribution of Receivers' beliefs
  - Receivers could acquire additional information from other sources **after** seeing the Sender's signal
  - Correlated noise (maximal concern for robustness)
- Multiple Receivers
  - discriminatory disclosures embedded into derivation of $\underline{V}(\mu)$
  - given common posterior μ, Nature provides (possibly private) signals to the agents and coordinates them on course of action most adversarial to Sender (among those consistent with assumed solution concept)
  - e.g., Bayes-correlated eq. given μ
- Online Appendix: conditionally independent signals

Robust Solutions

Robust Solutions

- Sender's expected payoffs when
  - Sender selects signal q
  - Nature selects signal π

$$\underline{v}(q,\, \pi) \equiv \sum_{\Omega} \int_{\mathcal{S}} \int_{\mathcal{R}} \underline{V}(\mu_0^{s,r}) d\pi(r|\omega,\, s) dq(s|\omega) \mu_0(\omega)$$

$$\overline{v}(q,\, \pi) \equiv \sum_{\Omega} \int_{\mathcal{S}} \int_{\mathcal{R}} \overline{V}(\mu_0^{s,r}) d\pi(r|\omega,\, s) dq(s|\omega) \mu_0(\omega)$$

where $\mu_0^{s,r}$ is the common posterior obtained from the prior μ_0 given realization (s, r) of the signal (q, π)

Worst-case optimality

Definition

Signal q is worst-case optimal if, for all signals q',

$$\inf_{\pi} \underline{v}(q, \pi) \geq \inf_{\pi} \underline{v}(q', \pi).$$

Worst-case optimality

- Define the Sender's payoff from full disclosure of the state, conditional on some belief μ, under the adversarial selection:

$$\underline{V}_{\mathsf{full}}(\mu) \equiv \sum_{\Omega} \underline{V}(\delta_\omega)\mu(\omega)$$

where δ_ω is a Dirac measure assigning prob 1 to ω.

Worst-case optimality

- Define the Sender's payoff from full disclosure of the state, conditional on some belief μ, under the adversarial selection:

$$\underline{V}_{\text{full}}(\mu) \equiv \sum_{\Omega} \underline{V}(\delta_\omega)\mu(\omega)$$

where δ_ω is a Dirac measure assigning prob 1 to ω.

Remark

Since both Nature and Sender can reveal state, signal q is worst-case optimal iff

$$\inf_{\pi} \underline{v}(q, \pi) = \underline{V}_{\text{full}}(\mu_0)$$

- W: set of worst-case optimal signals
  - non-empty (full disclosure is worst-case optimal)

Robust Solutions

Definition

Signal q_{RS} *is a* **robust solution** *if it maximizes* $\overline{v}(q, \emptyset)$ *over* W*.*

Robust Solutions

Definition

Signal q_{RS} is a **robust solution** *if it maximizes $\overline{v}(q, \emptyset)$ over W.*

- Lexicographic preferences:

Robust Solutions

Definition

Signal q_{RS} *is a* **robust solution** *if it maximizes* $\overline{v}(q, \emptyset)$ *over* W*.*

- Lexicographic preferences:
  - The Sender first guarantees herself the highest payoff guarantee in the worst-case scenario

Robust Solutions

Definition

Signal q_{RS} *is a* **robust solution** *if it maximizes* $\overline{v}(q, \emptyset)$ *over* W.

- Lexicographic preferences:
  - The Sender first guarantees herself the highest payoff guarantee in the worst-case scenario
  - If multiple policies yield the same payoff guarantee, she breaks the tie by considering the best-case scenario

Robust Solutions

Definition

Signal q_{RS} *is a* **robust solution** *if it maximizes* $\overline{v}(q, \emptyset)$ *over* W*.*

- Lexicographic preferences:
    - The Sender first guarantees herself the highest payoff guarantee in the worst-case scenario
    - If multiple policies yield the same payoff guarantee, she breaks the tie by considering the best-case scenario
- Clearly, q_{RS} also maximizes $\sup_\pi \overline{v}(q, \pi)$ over W

Robust Solutions

Definition

Signal q_{RS} *is a* **robust solution** *if it maximizes* $\overline{v}(q, \emptyset)$ *over* W*.*

- Lexicographic preferences:
  - The Sender first guarantees herself the highest payoff guarantee in the worst-case scenario
  - If multiple policies yield the same payoff guarantee, she breaks the tie by considering the best-case scenario
- Clearly, q_{RS} also maximizes $\sup_{\pi} \overline{v}(q, \pi)$ over W
  - Conservative approach: Sender prefers to provide information herself rather than counting on Nature to do it

Robust Solutions

Lemma

Signal q_{RS} is a robust solution iff the distribution of posterior beliefs $\rho_{RS} \in \Delta\Delta\Omega$ that it induces maximizes

$$\int \overline{V}(\mu) d\rho(\mu)$$

over the set of distributions of posterior beliefs $\mathcal{W} \subset \Delta\Delta\Omega$ satisfying

- *Bayes plausibility*

$$\int \mu d\rho(\mu) = \mu_0$$

- *Worst-case optimality (WCO)*

$$\int \mathit{lco}(\underline{V})(\mu) d\rho(\mu) = \underline{V}_{\mathit{full}}(\mu_0)$$

Robust vs Bayesian Solutions

- **Bayesian solutions**:

- **Robust solutions:**

Robust vs Bayesian Solutions

- **Bayesian solutions**:
  - q_{BP} maximizes $\overline{v}(q, \emptyset)$ over Q (feasible signals)
- **Robust solutions:**

Robust vs Bayesian Solutions

- **Bayesian solutions**:
  - q_{BP} maximizes $\overline{v}(q, \emptyset)$ over Q (feasible signals)

- **Robust solutions:**
  - q_{RS} maximizes $\overline{v}(q, \emptyset)$ over $W \subset Q$ (worst-case optimal signals)

Robust vs Bayesian Solutions

- **Bayesian solutions**:
  - q_{BP} maximizes $\overline{v}(q, \emptyset)$ over Q (feasible signals)
  - $\rho_{BP} \in \Delta\Delta\Omega$ maximizes $\int \overline{V}(\mu) d\rho(\mu)$ over all distributions $\rho \in \Delta\Delta\Omega$ satisfying Bayes plausibility, $\int \mu d\rho(\mu) = \mu_0$
- **Robust solutions:**
  - q_{RS} maximizes $\overline{v}(q, \emptyset)$ over $W \subset Q$ (worst-case optimal signals)

Robust vs Bayesian Solutions

- **Bayesian solutions**:
  - q_{BP} maximizes $\overline{v}(q, \emptyset)$ over Q (feasible signals)
  - $\rho_{BP} \in \Delta\Delta\Omega$ maximizes $\int \overline{V}(\mu)d\rho(\mu)$ over all distributions $\rho \in \Delta\Delta\Omega$ satisfying Bayes plausibility, $\int \mu d\rho(\mu) = \mu_0$

- **Robust solutions:**
  - q_{RS} maximizes $\overline{v}(q, \emptyset)$ over $W \subset Q$ (worst-case optimal signals)
  - $\rho_{RS} \in \Delta\Delta\Omega$ maximizes $\int \overline{V}(\mu)d\rho(\mu)$ over all distributions $\rho \in \Delta\Delta\Omega$ satisfying Bayes plausibility, $\int \mu d\rho(\mu) = \mu_0$, and the WCO constraint

$$\int \text{lco}(\underline{V})(\mu)d\rho(\mu) = \underline{V}_{\text{full}}(\mu_0)$$

Separation Theorem

Separation Theorem

Theorem

There exists

$$\mathcal{F} \subseteq 2^{\Omega}$$

such that

$$\mathcal{W} = \{\rho \in \Delta\Delta\Omega : \; \rho \; \textit{satisfies BP} \; \textit{and} \; supp(\mu) \in \mathcal{F}, \forall\, \mu \in supp(\rho)\}\,.$$

Separation Theorem

Theorem

There exists

$$\mathcal{F} \subseteq 2^{\Omega}$$

such that

$$\mathcal{W} = \left\{\rho \in \Delta\Delta\Omega : \ \rho \ \textit{satisfies BP} \ \textit{and} \ supp(\mu) \in \mathcal{F}, \forall \, \mu \in supp(\rho)\right\}.$$

Therefore, $\rho_{RS} \in \Delta\Delta\Omega$ *is a robust solution iff* ρ_{RS} *maximizes*

$$\int \overline{V}(\mu) d\rho(\mu)$$

over all Bayes-plausible distributions over posterior beliefs $\rho \in \Delta\Delta\Omega$ *such that*

$$supp(\mu) \in \mathcal{F}, \forall \mu \in supp(\rho).$$

Separation Theorem

Theorem

Let

$$\mathcal{F} \equiv \{B \subseteq \Omega : \underline{V}|_{\Delta B} \geq \underline{V}_{\textit{full}}|_{\Delta B}\}.$$

Then,

$$\mathcal{W} = \{\rho \in \Delta\Delta\Omega : \ \rho \ \textit{ satisfies BP } \ \textit{and } \ supp(\mu) \in \mathcal{F}, \forall\, \mu \in supp(\rho)\}\,.$$

Therefore, $\rho_{RS} \in \Delta\Delta\Omega$ *is a robust solution iff* ρ_{RS} *maximizes*

$$\int \overline{V}(\mu) d\rho(\mu)$$

over all Bayes-plausible distributions over posterior beliefs $\rho \in \Delta\Delta\Omega$ *such that*

$$supp(\mu) \in \mathcal{F}, \forall \mu \in supp(\rho).$$

Separation Theorem

Theorem

Let

$$\mathcal{F} \equiv \{B \subseteq \Omega : \underline{V}(\mu) \geq \underline{V}_{\mathit{full}}(\mu), \forall \mu \in \Delta B\}.$$

Then,

$$\mathcal{W} = \{\rho \in \Delta\Delta\Omega : \ \rho \ \textit{satisfies BP} \ \textit{and} \ supp(\mu) \in \mathcal{F}, \forall \, \mu \in supp(\rho)\} \, .$$

Therefore, $\rho_{RS} \in \Delta\Delta\Omega$ *is a robust solution iff* ρ_{RS} *maximizes*

$$\int \overline{V}(\mu) d\rho(\mu)$$

over all Bayes-plausible distributions over posterior beliefs $\rho \in \Delta\Delta\Omega$ *such that*

$$supp(\mu) \in \mathcal{F}, \forall \mu \in supp(\rho).$$

Separation Theorem

- Idea:
  - Suppose Sender induces posterior μ with $supp(\mu) = B$ for which there exists $\eta \in \Delta B$ s.t. $\underline{V}(\eta) < \underline{V}_{\text{full}}(\eta)$.
  - Starting from μ, Nature can induce η with strictly positive probability.
  - Starting from μ, Nature can bring Sender's payoff strictly below $\underline{V}_{\text{full}}(\mu)$.
  - This is because Nature can respond to any other posterior $\mu' \in supp(\rho)$ by fully disclosing the state,

$$\int \text{lco}(\underline{V})(\tilde{\mu})d\rho(\tilde{\mu}) < \underline{V}_{\text{full}}(\mu_0)$$

  - Hence, Sender's policy inducing such μ cannot be worst-case optimal.

Separation Theorem

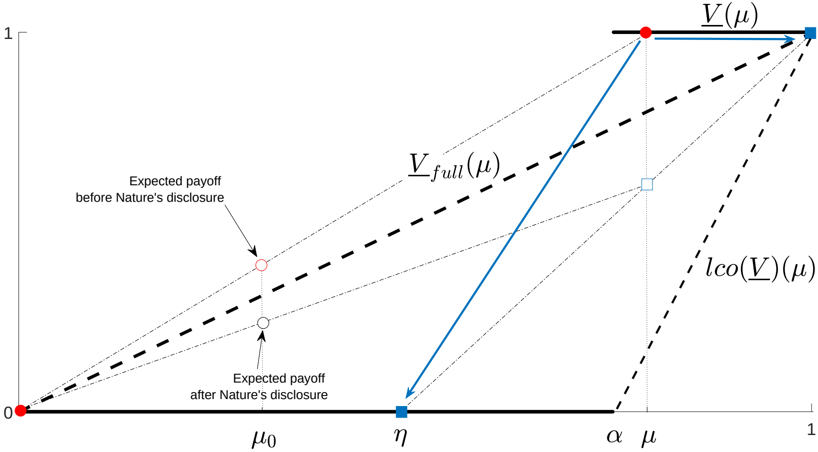

Figure: Prosecutor example

Properties of Robust Solutions

Existence

Corollary

A robust solution always exists.

Existence

Corollary

A robust solution always exists.

- Existence follows because the WCO constraint is only a constraint on feasible supports (compactness is preserved).
- Existence guaranteed by possibility for Nature to condition on realization of Sender's signal.

State separation

Corollary

Suppose there exist $\omega, \omega' \in \Omega$ *and* $\lambda \in (0,1)$ *s.t.*

$$\underline{V}(\lambda\delta_\omega + (1-\lambda)\delta_{\omega'}) < \lambda\underline{V}(\delta_\omega) + (1-\lambda)\underline{V}(\delta_{\omega'}),$$

Then any robust solution must separate ω *and* ω'*.*

State separation

Corollary

Suppose there exist $\omega, \omega' \in \Omega$ *and* $\lambda \in (0,1)$ *s.t.*

$$\underline{V}(\lambda\delta_\omega + (1-\lambda)\delta_{\omega'}) < \lambda\underline{V}(\delta_\omega) + (1-\lambda)\underline{V}(\delta_{\omega'}),$$

Then any robust solution must separate ω *and* ω'*.*

- Assumption: there exists **some belief** supported on $\{\omega, \omega'\}$ under which Sender's payoff below full disclosure
- Conclusion: **ALL** posterior beliefs must separate ω and ω'.

Full disclosure vs No restriction

Corollary (Full disclosure)

Full disclosure is the unique robust solution if $\mathcal{F} = \Omega$*, meaning that any pair of states must be separated under any worst-case optimal distribution.*

Full disclosure vs No restriction

Corollary (Full disclosure)

Full disclosure is the unique robust solution if $\mathcal{F} = \Omega$, meaning that any pair of states must be separated under any worst-case optimal distribution.

Corollary (No restrictions)

All feasible distributions are worst-case optimal if, and only if, $\Omega \in \mathcal{F}$, meaning that no pair of states must be separated under any worst-case optimal distribution. Then, the set of robust solutions coincides with the set of Bayesian solutions.

Robustness of Bayesian Solutions

Corollary

Bayesian solution ρ_{BP} is robust iff for any $\mu \in supp(\rho_{BP})$ and any $\eta \in \Delta\Omega$ s.t. $supp(\eta) \subset supp(\mu)$,

$$\underline{V}(\eta) \geq \underline{V}_{\textit{full}}(\eta).$$

Robustness of Bayesian Solutions

Corollary

Bayesian solution ρ_{BP} is robust iff for any $\mu \in supp(\rho_{BP})$ and any $\eta \in \Delta\Omega$ s.t. $supp(\eta) \subset supp(\mu)$,

$$\underline{V}(\eta) \geq \underline{V}_{\mathit{full}}(\eta).$$

- Corollary for the binary-state case: Any robust solution is either
  - full disclosure, or
  - a Bayesian solution.

Worst-case optimality preserved under more disclosure

Corollary

$\mathcal{W}$ is closed under Blackwell dominance: If $\rho' \in \mathcal{W}$, and ρ Blackwell dominates ρ', then $\rho \in \mathcal{W}$.

Informativeness of Robust vs Bayesian solutions

Corollary

Given any Bayesian solution ρ_{BP}, there exists robust solution ρ_{RS} s.t. either ρ_{RS} and ρ_{BP} not comparable in Blackwell order, or ρ_{RS} Blackwell dominates ρ_{BP}.

Informativeness of Robust vs Bayesian solutions

Corollary

Given any Bayesian solution ρ_{BP}, there exists robust solution ρ_{RS} s.t. either ρ_{RS} and ρ_{BP} not comparable in Blackwell order, or ρ_{RS} Blackwell dominates ρ_{BP}.

- Proof: If Bayesian solution ρ_{BP} is Blackwell more informative than robust solution ρ_{RS}, then ρ_{BP} also robust.

Informativeness of Robust vs Bayesian solutions

Corollary

Given any Bayesian solution ρ_{BP}, there exists robust solution ρ_{RS} s.t. either ρ_{RS} and ρ_{BP} not comparable in Blackwell order, or ρ_{RS} Blackwell dominates ρ_{BP}.

- Proof: If Bayesian solution ρ_{BP} is Blackwell more informative than robust solution ρ_{RS}, then ρ_{BP} also robust.
- Reason why robustness calls for more disclosure:

Informativeness of Robust vs Bayesian solutions

Corollary

Given any Bayesian solution ρ_{BP}, there exists robust solution ρ_{RS} s.t. either ρ_{RS} and ρ_{BP} not comparable in Blackwell order, or ρ_{RS} Blackwell dominates ρ_{BP}.

- Proof: If Bayesian solution ρ_{BP} is Blackwell more informative than robust solution ρ_{RS}, then ρ_{BP} also robust.
- Reason why robustness calls for more disclosure:
    - It is **not** because Sender worries that Nature fully discloses the state if she does not.

Informativeness of Robust vs Bayesian solutions

Corollary

Given any Bayesian solution ρ_{BP}, there exists robust solution ρ_{RS} s.t. either ρ_{RS} and ρ_{BP} not comparable in Blackwell order, or ρ_{RS} Blackwell dominates ρ_{BP}.

- Proof: If Bayesian solution ρ_{BP} is Blackwell more informative than robust solution ρ_{RS}, then ρ_{BP} also robust.
- Reason why robustness calls for more disclosure:
  - It is **not** because Sender worries that Nature fully discloses the state if she does not.
  - Concealing information gives Nature more room for adversarial design.

Informativeness of Robust vs Bayesian solutions

Corollary

Given any Bayesian solution ρ_{BP}, there exists robust solution ρ_{RS} s.t. either ρ_{RS} and ρ_{BP} not comparable in Blackwell order, or ρ_{RS} Blackwell dominates ρ_{BP}.

- Proof: If Bayesian solution ρ_{BP} is Blackwell more informative than robust solution ρ_{RS}, then ρ_{BP} also robust.
- Reason why robustness calls for more disclosure:
  - It is **not** because Sender worries that Nature fully discloses the state if she does not.
  - Concealing information gives Nature more room for adversarial design.
- If Bayesian solution ρ_{BP} is not robust and is strictly Blackwell dominated by robust solution ρ_{RS}, then ρ_{RS} separates states that ρ_{BP} does not.

Informativeness of Robust vs Bayesian solutions

Corollary

Given any Bayesian solution ρ_{BP}, there exists robust solution ρ_{RS} s.t. either ρ_{RS} and ρ_{BP} not comparable in Blackwell order, or ρ_{RS} Blackwell dominates ρ_{BP}.

- Proof: If Bayesian solution ρ_{BP} is Blackwell more informative than robust solution ρ_{RS}, then ρ_{BP} also robust.
- Reason why robustness calls for more disclosure:
  - It is **not** because Sender worries that Nature fully discloses the state if she does not.
  - Concealing information gives Nature more room for adversarial design.
- If Bayesian solution ρ_{BP} is not robust and is strictly Blackwell dominated by robust solution ρ_{RS}, then ρ_{RS} separates states that ρ_{BP} does not.
- Conclusion not true with conditionally independent signals!

Concavification

- Let $v_{\text{low}} := \min_{\omega \in \Omega} \overline{V}(\delta_\omega) - 1$
- Auxiliary function

$$\overline{V}_{\mathcal{F}}(\mu) = \begin{cases} \overline{V}(\mu) & \text{if } supp(\mu) \in \mathcal{F} \text{ and } \overline{V}(\mu) \geq v_{\text{low}} \\ v_{\text{low}} & \text{otherwise} \end{cases}$$

Corollary

A feasible distribution $\rho \in \Delta\Delta\Omega$ is robust iff

$$\int \overline{V}_{\mathcal{F}}(\mu) d\rho(\mu) = co(\overline{V}_{\mathcal{F}})(\mu_0).$$

Concavification

- Let $v_{\text{low}} := \min_{\omega\in\Omega} \overline{V}(\delta_\omega) - 1$
- Auxiliary function

$$\overline{V}_{\mathcal{F}}(\mu) = \begin{cases} \overline{V}(\mu) & \text{if } supp(\mu) \in \mathcal{F} \text{ and } \overline{V}(\mu) \geq v_{\text{low}} \\ v_{\text{low}} & \text{otherwise} \end{cases}$$

Corollary

A feasible distribution $\rho \in \Delta\Delta\Omega$ is robust iff

$$\int \overline{V}_{\mathcal{F}}(\mu) d\rho(\mu) = co(\overline{V}_{\mathcal{F}})(\mu_0).$$

- Corollary: We need at most $|\Omega|$ signals in a robust solution.

Weighted objective function

Weighted objective function

- Suppose that instead of the lexicographic approach, the Sender maximizes

$$\sup_{q \in Q} \left\{ \lambda \inf_{\pi \in \Pi} \underline{v}(q, \pi) + (1 - \lambda) \overline{v}(q, \emptyset) \right\},$$

for some $\lambda \in [0, 1]$.

Weighted objective function

- Suppose that instead of the lexicographic approach, the Sender maximizes

$$\sup_{q \in Q} \left\{ \lambda \inf_{\pi \in \Pi} \underline{v}(q, \pi) + (1 - \lambda)\overline{v}(q, \emptyset) \right\},$$

  for some $\lambda \in [0, 1]$.

- Possible interpretation: The Sender is Bayesian and λ is the assessed probability of Nature being adversarial.

Weighted objective function

- Suppose that instead of the lexicographic approach, the Sender maximizes

$$\sup_{q \in Q} \left\{ \lambda \inf_{\pi \in \Pi} \underline{v}(q, \pi) + (1 - \lambda) \overline{v}(q, \emptyset) \right\},$$

  for some $\lambda \in [0, 1]$.

- Possible interpretation: The Sender is Bayesian and λ is the assessed probability of Nature being adversarial.

- Related concepts in other settings: Hurwicz (1951), Gul and Pesendorfer (2015), and Grant et al. (2020).

Weighted objective function

Under a regularity condition on the objective function:

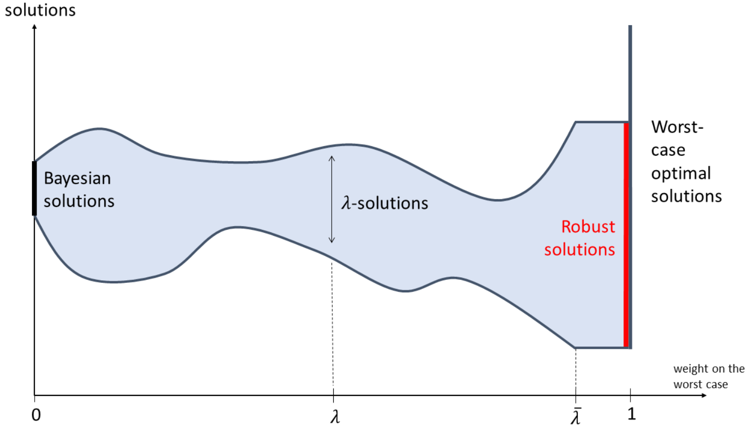

Weighted objective function

Let d denote the Chebyshev metric on $\Delta\Omega$: $d(\mu, \eta) = \max_{\omega \in \Omega} |\mu(\omega) - \eta(\omega)|$.

Weighted objective function

Let d denote the Chebyshev metric on $\Delta\Omega$: $d(\mu, \eta) = \max_{\omega \in \Omega} |\mu(\omega) - \eta(\omega)|$.

Definition

The function $\overline{V}$ is regular if there exist positive constants K and L such that for every non-degenerate $\mu \in \Delta\Omega$ and every $\omega \in$ supp(μ), there exists $\eta \in \Delta\Omega$ with supp$(\eta) \subseteq$ supp$(\mu) \setminus \{\omega\}$ such that $d(\mu, \eta) \leq K\mu(\omega)$ and

$$\overline{V}(\mu) - \overline{V}(\eta) \leq Ld(\mu, \eta).$$

Weighted objective function

Let d denote the Chebyshev metric on $\Delta\Omega$: $d(\mu, \eta) = \max_{\omega \in \Omega} |\mu(\omega) - \eta(\omega)|$.

Definition

The function $\overline{V}$ is regular if there exist positive constants K and L such that for every non-degenerate $\mu \in \Delta\Omega$ and every $\omega \in$ supp(μ), there exists $\eta \in \Delta\Omega$ with supp$(\eta) \subseteq$ supp$(\mu) \setminus \{\omega\}$ such that $d(\mu, \eta) \leq K\mu(\omega)$ and

$$\overline{V}(\mu) - \overline{V}(\eta) \leq Ld(\mu, \eta).$$

Regularity requires that, for any μ and any $\omega \in \text{supp}(\mu)$, there exists a nearby belief supported on $\text{supp}(\mu) \setminus \{\omega\}$ that is not much worse for the designer under the favorable selection $\overline{V}$.

Weighted objective function

Examples of regular functions:

- Lipschitz continuous $\overline{V}$; but this is weaker because the Lipschitz condition is required to hold:

Weighted objective function

Examples of regular functions:

- Lipschitz continuous $\overline{V}$; but this is weaker because the Lipschitz condition is required to hold:
  - only for beliefs μ that attach vanishing probability to some state ω,

Weighted objective function

Examples of regular functions:

- Lipschitz continuous $\overline{V}$; but this is weaker because the Lipschitz condition is required to hold:
  - only for beliefs μ that attach vanishing probability to some state ω,
  - only for *some* belief η in the neighborhood of a given μ,

Weighted objective function

Examples of regular functions:

- Lipschitz continuous $\overline{V}$; but this is weaker because the Lipschitz condition is required to hold:
  - only for beliefs μ that attach vanishing probability to some state ω,
  - only for *some* belief η in the neighborhood of a given μ,
  - only in one direction (the condition rules out functions $\overline{V}(\mu)$ that *decrease* at an infinite rate as $\mu(\omega)$ approaches 0);

Weighted objective function

Examples of regular functions:

- Lipschitz continuous $\overline{V}$; but this is weaker because the Lipschitz condition is required to hold:
  - only for beliefs μ that attach vanishing probability to some state ω,
  - only for *some* belief η in the neighborhood of a given μ,
  - only in one direction (the condition rules out functions $\overline{V}(\mu)$ that *decrease* at an infinite rate as $\mu(\omega)$ approaches 0);
- $\overline{V}(\mu) = v(\mathbb{E}_\mu[\omega])$, for some real-valued function v;

Weighted objective function

Examples of regular functions:

- Lipschitz continuous $\overline{V}$; but this is weaker because the Lipschitz condition is required to hold:
  - only for beliefs μ that attach vanishing probability to some state ω,
  - only for *some* belief η in the neighborhood of a given μ,
  - only in one direction (the condition rules out functions $\overline{V}(\mu)$ that *decrease* at an infinite rate as $\mu(\omega)$ approaches 0);
- $\overline{V}(\mu) = v(\mathbb{E}_\mu[\omega])$, for some real-valued function v;
- $\overline{V}(\mu) = \sum_{i=1}^{k} a_i \mathbf{1}_{\{\mu \in A_i\}}$ for some partition $(A_1, ..., A_k)$ of $\Delta\Omega$

Weighted objective function

Let $S(\lambda)$ denote the set of solutions to the weighted problem with weight λ.

Weighted objective function

Let $S(\lambda)$ denote the set of solutions to the weighted problem with weight λ.

Theorem

Suppose that $\overline{V}$ is regular. Then, there exists $\overline{\lambda} < 1$ such that for all $\lambda \in (\overline{\lambda}, 1)$, $S(\lambda)$ coincides with the set of robust solutions.

Weighted objective function

Let $S(\lambda)$ denote the set of solutions to the weighted problem with weight λ.

Theorem

Suppose that $\overline{V}$ is regular. Then, there exists $\overline{\lambda} < 1$ such that for all $\lambda \in (\overline{\lambda}, 1)$, $S(\lambda)$ coincides with the set of robust solutions.

Without regularity: Any limit of λ-solutions as $\lambda \nearrow 1$ is a robust solution (but not the other way around).

Weighted objective function

Lemma

There exists a constant $\delta > 0$ such that for any μ such that supp$(\mu) \notin \mathcal{F}$,

$$\underline{V}_{\mathit{full}}(\mu) - \mathit{lco}(\underline{V})(\mu) \geq \delta \cdot \max_{B \subseteq \mathit{supp}(\mu),\, B \notin \mathcal{F}} \min_{\omega \in B} \{\mu(\omega)\}.$$

Weighted objective function

Lemma

There exists a constant $\delta > 0$ such that for any μ such that supp$(\mu) \notin \mathcal{F}$,

$$\underline{V}_{\mathit{full}}(\mu) - \mathit{lco}(\underline{V})(\mu) \geq \delta \cdot \max_{B \subseteq \mathit{supp}(\mu),\, B \notin \mathcal{F}} \min_{\omega \in B} \{\mu(\omega)\}.$$

Lemma

For a regular function $\overline{V}$, there exists a constant $\Delta > 0$ such that for any μ such that supp$(\mu) \notin \mathcal{F}$,

$$\overline{V}(\mu) - \mathit{co}(\overline{V}_{\mathcal{F}})(\mu) \leq \Delta \max_{B \subseteq \mathit{supp}(\mu),\, B \notin \mathcal{F}} \min_{\omega \in B} \{\mu(\omega)\}.$$

Weighted objective function

Lemma

There exists a constant $\delta > 0$ such that for any μ such that supp$(\mu) \notin \mathcal{F}$,

$$\underline{V}_{\textit{full}}(\mu) - \textit{lco}(\underline{V})(\mu) \geq \delta \cdot \max_{B \subseteq \textit{supp}(\mu),\, B \notin \mathcal{F}} \min_{\omega \in B} \{\mu(\omega)\}.$$

Lemma

For a regular function $\overline{V}$, there exists a constant $\Delta > 0$ such that for any μ such that supp$(\mu) \notin \mathcal{F}$,

$$\overline{V}(\mu) - \textit{co}(\overline{V}_{\mathcal{F}})(\mu) \leq \Delta \max_{B \subseteq \textit{supp}(\mu),\, B \notin \mathcal{F}} \min_{\omega \in B} \{\mu(\omega)\}.$$

Lemma

Suppose that $\overline{V}$ is regular. There exists $\overline{\lambda} < 1$ such that, for all $\lambda \in (\overline{\lambda}, 1]$, if ρ is a λ-solution, then ρ cannot assign positive probability to μ such that supp$(\mu) \notin \mathcal{F}$.

Applications

Privately Informed Receiver

- **Guo and Shmaya (ECMA, 2019)**
- State ω is the value to a buyer
- Exogenous price $p \in (0, 1)$
- Seller's payoff is 1 if trade, 0 otherwise
- Buyer's exogenous private information given by $f(t|\omega)$, ordered by MLRP
- A Bayesian solution has an *interval structure*: each buyer's type t is induced to trade on an interval of states, and less optimistic types trade on smaller intervals

Privately Informed Receiver

- **Guo and Shmaya (ECMA, 2019)**
- State ω is the value to a buyer
- Exogenous price $p \in (0, 1)$
- Seller's payoff is 1 if trade, 0 otherwise
- Buyer's exogenous private information given by $f(t|\omega)$, ordered by MLRP
- A Bayesian solution has an *interval structure*: each buyer's type t is induced to trade on an interval of states, and less optimistic types trade on smaller intervals

Proposition

Any robust solution separates states $\omega \leq p$ from states $\omega' > p$.

Limits to Price Discrimination

- **Bergemann, Brooks, Morris (AER, 2015)**
  - The designer segments the market to maximize either producer or consumer surplus.

Limits to Price Discrimination

- **Bergemann, Brooks, Morris (AER, 2015)**
  - The designer segments the market to maximize either producer or consumer surplus.

Proposition

The BBM solution is robust.

Limits to Price Discrimination

- **Bergemann, Brooks, Morris (AER, 2015)**
  - The designer segments the market to maximize either producer or consumer surplus.

Proposition

The BBM solution is robust.

- The solution is robust even though it is very intricate.

Lemons problem

- Seller's value: ω (known to seller, unknown to buyer)
- Buyer's value: $\omega + \Delta$, with $\Delta > 0$ (Δ is a constant)
- Exogenous price p drawn from $U[0, 1]$
- Trade if *(i)* $p \geq \omega$ and *(ii)* $\mathbb{E}_\mu[\tilde{\omega}|\tilde{\omega} \leq p] + \Delta > p$
- Seller designs info structure

Lemons problem

- Seller's value: ω (known to seller, unknown to buyer)
- Buyer's value: $\omega + \Delta$, with $\Delta > 0$ (Δ is a constant)
- Exogenous price p drawn from $U[0, 1]$
- Trade if *(i)* $p \geq \omega$ and *(ii)* $\mathbb{E}_\mu[\tilde{\omega}|\tilde{\omega} \leq p] + \Delta > p$
- Seller designs info structure

Proposition

Under any robust solution ρ_{RS}, for any $\mu, \mu' \in$ supp(ρ_{RS}), $diam($supp$(\mu)), diam($supp$(\mu')) \leq \Delta$ but $diam($supp$(\mu) \cup$ supp$(\mu')) > \Delta$.

Lemons problem

- Seller's value: ω (known to seller, unknown to buyer)
- Buyer's value: $\omega + \Delta$, with $\Delta > 0$ (Δ is a constant)
- Exogenous price p drawn from $U[0, 1]$
- Trade if *(i)* $p \geq \omega$ and *(ii)* $\mathbb{E}_\mu[\tilde{\omega}|\tilde{\omega} \leq p] + \Delta > p$
- Seller designs info structure

Proposition

Under any robust solution ρ_{RS}, for any $\mu, \mu' \in$ supp(ρ_{RS}), $diam($supp$(\mu)), diam($supp$(\mu')) \leq \Delta$ but $diam($supp$(\mu) \cup$ supp$(\mu')) > \Delta$.

- Robust solutions are minimally informative among those that eliminate adverse selection

Supermodular Games

- Continuum of Receivers
- $a_i = \{0, 1\}$; $A \in [0, 1]$: aggregate "attack"
- Payoff from not attacking normalized to 0; payoff from attacking

$$\begin{cases} g > 0 & \text{if} \quad A \geq \omega \\ b < 0 & \text{if} \quad A < \omega \end{cases}$$

- Designer's payoff: $1 - A$
- Bayesian solution under best rationalizable profile: Upper censorship
  - Reveals each $\omega < 0$ w.p. $\gamma_{BP} \in (0, 1)$ (w.p. $1 - \gamma_{BP}$, reveals nothing)
  - Conceals all $\omega > 0$

Supermodular Games

- Continuum of Receivers
- $a_i = \{0, 1\}$; $A \in [0, 1]$: aggregate "attack"
- Payoff from not attacking normalized to 0; payoff from attacking

$$\begin{cases} g > 0 & \text{if} \quad A \geq \omega \\ b < 0 & \text{if} \quad A < \omega \end{cases}$$

- Designer's payoff: $1 - A$
- Bayesian solution under best rationalizable profile: Upper censorship
  - Reveals each $\omega < 0$ w.p. $\gamma_{BP} \in (0, 1)$ (w.p. $1 - \gamma_{BP}$, reveals nothing)
  - Conceals all $\omega > 0$

Proposition

Suppose that $\overline{V}$ and $\underline{V}$ capture the payoff in the best and worst rationalizable profile for the Sender, respectively. Then, a robust solution reveals $\omega < 0$ w.p. $\gamma^* > \gamma_{BP}$, conceals all $\omega \in [0, 1]$, reveals all $\omega > 1$ with certainty.

Conditionally Independent Signals

Conditionally-independent Robust Solutions

- Nature cannot condition on the realization of Sender's signal
  - $\pi : \Omega \to \Delta\mathcal{R}$
  - so far: $\pi : \Omega \times \mathcal{S} \to \Delta\mathcal{R}$

Existence of CI-Robust Solutions

- A robust solution may fail to exist

Existence of CI-Robust Solutions

- A robust solution may fail to exist
- A robust solution exists if $\underline{V}$ is continuous

Existence of CI-Robust Solutions

- A robust solution may fail to exist
- A robust solution exists if $\underline{V}$ is continuous

Existence of CI-Robust Solutions

- A robust solution may fail to exist
- A robust solution exists if $\underline{V}$ is continuous

Definition

A feasible distribution $\rho \in \Delta\Delta\Omega$ *is a* **weak CI-robust solution** *if it maximizes*

$$\int \overline{V}(\mu) d\rho(\mu)$$

over $cl(W_{CI})$*, where* $cl(W_{CI})$ *is closure (in weak** * *topology) of set of CI-worst-case solutions.*

Existence of CI-Robust Solutions

- A robust solution may fail to exist
- A robust solution exists if $\underline{V}$ is continuous

Definition

A feasible distribution $\rho \in \Delta\Delta\Omega$ is a **weak CI-robust solution** *if it maximizes*

$$\int \overline{V}(\mu) d\rho(\mu)$$

over $cl(W_{CI})$, where $cl(W_{CI})$ is closure (in weak* topology) of set of CI-worst-case solutions.

Theorem

A weak solution exists no matter $\underline{V}$.

Separation under CI-Robust Solutions

- Sufficient conditions for state separation under CI-robust solutions
  - weaker than those for robust solutions

Separation under CI-Robust Solutions

- Sufficient conditions for state separation under CI-robust solutions
  - weaker than those for robust solutions
  - intuitively, states ω and ω' must be separated if $\underline{V}(\mu)$ lies strictly below $\underline{V}_{\text{full}}(\mu)$ for μ in the neighborhood of δ_ω or $\delta_{\omega'}$.

Separation under CI-Robust Solutions

- Sufficient conditions for state separation under CI-robust solutions
  - weaker than those for robust solutions
  - intuitively, states ω and ω' must be separated if $\underline{V}(\mu)$ lies strictly below $\underline{V}_{\text{full}}(\mu)$ for μ in the neighborhood of δ_ω or $\delta_{\omega'}$.
  - whenever ω and ω' must be separated under CI-robust solutions, they must be separated under robust solutions

CI-robust solutions: Binary state

- Unlike robust solutions, CI-robust solutions for binary states need not coincide with either
  - Bayesian solutions, or
  - full disclosure.

Blackwell Informativeness of CI-robust solutions

- Unlike robust solutions, CI-robust solutions need not be Blackwell more informative than Bayesian solution

- Example in which unique Bayesian solution is Blackwell strictly more informative than all CI-robust solutions

  - Nature cannot engineer MPS **conditional** each on s separately
  - Thus, any additional disclosure by Nature moves **all** posterior beliefs.
  - It is possible that a less informative signal ρ_{RS} is worst-case optimal, but a more informative signal ρ_{BP} is not.

Conclusions

- Bayesian persuasion when Sender uncertain about
  - Receivers' information
  - strategy selection
- Robust solutions
  - best-case optimal among worst-case optimal ones
- Separation theorem
  - any pair of states over which Nature can construct beliefs yielding less than the full-information payoff are separated
- Robustness $\implies$ more disclosure
  - but only through more separation (not a MPS over the same support)
- Future work:
  - Implications for applications, especially ones where tractability is an issue
  - Robust discriminatory disclosure
  - Other notions of robustness

Conclusions

Thank you