Persuasion with risk-conscious agents: A geometric approach

Jerry Anunrojwong

Agoda | MIT | Chulalongkorn University

Conference on Web and Internet Economics (WINE) 2019
Based on joint work with **Krishnamurthy Iyer** and **David Lingenbrink**

Motivation

Motivation

Information design/Bayesian Persuasion: How should a platform share information with its users to achieve better outcomes?

Lots of theoretical work.

Lots of recent applications.

Motivation

Information design/Bayesian Persuasion: How should a platform share information with its users to achieve better outcomes?

Lots of theoretical work.

- Pioneers: Rayo and Segal (2010), Kamenica and Gentzkow (2011)
- Econ: Bergemann and Morris (2016, 2018), Kolotilin et al. (2017), Taneva (2019), Doval and Ely (2016), ...
- CS: Dughmi and Xu (2016, 2017), Gan et al. (2019), ...

Lots of recent applications.

Motivation

Information design/Bayesian Persuasion: How should a platform share information with its users to achieve better outcomes?

Lots of theoretical work.

Lots of recent applications.

- waiting times in queues — Lingenbrink and Iyer (2017)
- crowdsourcing exploration — Papanastasiou et al. (2017)
- content quality in social networks
  Candogan and Drakopoulos (2017), Candogan (2019)
- demand and inventory signaling – Jain et al. (2018)
- security games – Xu et al. (201*)
- ad auctions – Badanidiyuru et al. (2018)
- voting – Alonso and Camara (2016)
- warning against risks — Alizamir et al. (2019)

Motivation

Information design/Bayesian Persuasion: How should a platform share information with its users to achieve better outcomes?

Lots of theoretical work.

Lots of recent applications.

Central Assumption: Users are **expected utility** (EU) maximizers.

Motivation

Extensive empirical evidence that human behavior is not adequately modeled by EU maximization.

Motivation

Extensive empirical evidence that human behavior is not adequately modeled by EU maximization.

Service systems: customers are typically averse to the uncertainty in how long they have to wait for service.

- uncertain wait-times perceived to be **longer** (Maister 2005).
- **Variance** of the wait-time impacts the decision to avail service.

Motivation

Extensive empirical evidence that human behavior is not adequately modeled by EU maximization.

Service systems: customers are typically averse to the uncertainty in how long they have to wait for service.

- uncertain wait-times perceived to be **longer** (Maister 2005).
- **Variance** of the wait-time impacts the decision to avail service.

Finance: Mean-variance, risk-measures, Value-at-risk, CVar, etc.

Motivation

Extensive empirical evidence that human behavior is not adequately modeled by EU maximization.

Service systems: customers are typically averse to the uncertainty in how long they have to wait for service.

- uncertain wait-times perceived to be **longer** (Maister 2005).
- **Variance** of the wait-time impacts the decision to avail service.

Finance: Mean-variance, risk-measures, Value-at-risk, CVar, etc.

In this talk

Extend the methodology of information design to incorporate **general** models of human behavior

Model

Bayesian persuasion: standard framework

Unknown state of the world: $\omega \in \Omega$ "waiting time"

Bayesian persuasion: standard framework

Unknown state of the world: $\omega \in \Omega$

Receiver: Bayesian and expected-utility maximizer

- Prior μ^* about ω
- Must choose an action $a \in A$

“whether to join or leave the queue”

Bayesian persuasion: standard framework

Unknown state of the world: $\omega \in \Omega$

Receiver: Bayesian and expected-utility maximizer

- Prior μ^* about ω
- Must choose an action $a \in A$

Sender: Bayesian, can **commit** to a *signaling scheme* $\pi : \Omega \mapsto S$

- **observes** ω, then sends a (perhaps random) signal $s = \pi(\omega) \in S$ to receiver before receiver acts

Bayesian persuasion: standard framework

Unknown state of the world: $\omega \in \Omega$

Receiver: Bayesian and expected-utility maximizer

- Prior μ^* about ω
- Must choose an action $a \in A$

Sender: Bayesian, can **commit** to a *signaling scheme* $\pi : \Omega \mapsto S$

- **observes** ω, then sends a (perhaps random) signal $s = \pi(\omega) \in S$ to receiver before receiver acts

Receiver updates belief to μ, with optimal action $a^*(\mu)$.

Bayesian persuasion: standard framework

Unknown state of the world: $\omega \in \Omega$

Receiver: Bayesian and expected-utility maximizer

- Prior μ^* about ω
- Must choose an action $a \in A$

Sender: Bayesian, can **commit** to a *signaling scheme* $\pi : \Omega \mapsto S$

- **observes** ω, then sends a (perhaps random) signal $s = \pi(\omega) \in S$ to receiver before receiver acts

Receiver updates belief to μ, with optimal action $a^*(\mu)$.

Sender's goal: Maximize $\mathbb{E}[\texttt{SenderUtil}(\omega, a^*(\mu))]$

Bayesian persuasion: our framework

Unknown state of the world: $\omega \in \Omega$

Receiver: Bayesian and ~~expected-utility maximizer~~

- Prior μ^* about ω
- Must choose an action $a \in A$

Sender: Bayesian, can **commit** to a *signaling scheme* $\pi : \Omega \mapsto S$

- **observes** ω, then sends a (perhaps random) signal $s = \pi(\omega) \in S$ to receiver before receiver acts

Receiver updates belief to μ, with optimal action $a^*(\mu)$.

Sender's goal: Maximize $\mathbb{E}[\texttt{SenderUtil}(\omega, a^*(\mu))]$

Bayesian persuasion: our framework

Unknown state of the world: $\omega \in \Omega$

Receiver: Bayesian and ~~expected-utility maximizer~~ **“risk-conscious”**

- Prior μ^* about ω
- Must choose an action $a \in A$

Sender: Bayesian, can **commit** to a *signaling scheme* $\pi : \Omega \mapsto S$

- **observes** ω, then sends a (perhaps random) signal $s = \pi(\omega) \in S$ to receiver before receiver acts

Receiver updates belief to μ, with optimal action $a^*(\mu)$.

Sender’s goal: Maximize $\mathbb{E}[\texttt{SenderUtil}(\omega, a^*(\mu))]$

Risk-conscious agents: definition

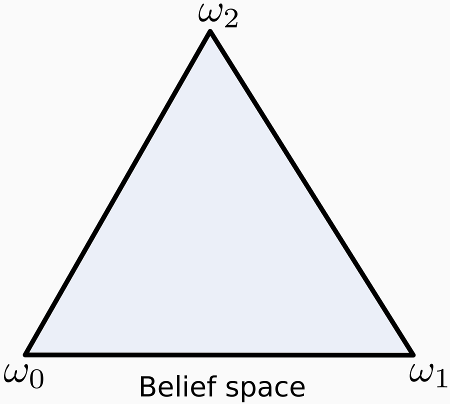

Risk-conscious agents: definition

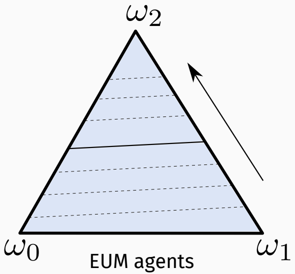

Expected utility maximizers: utility is **linear** in belief.

Risk-conscious agents: definition

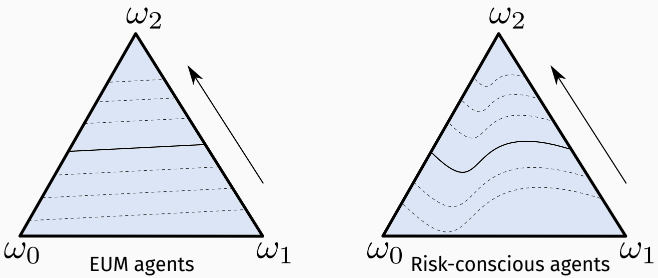

Expected utility maximizers: utility is **linear** in belief.

Risk-conscious agents: utility is a **non-linear** function of belief.

Challenges

Revelation principle holds for EUM, fails for risk-conscious

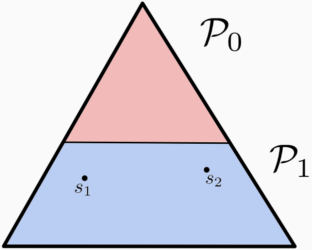

Suppose two signals s_1 and s_2 induce receiver beliefs with same optimal action.

Revelation principle holds for EUM, fails for risk-conscious

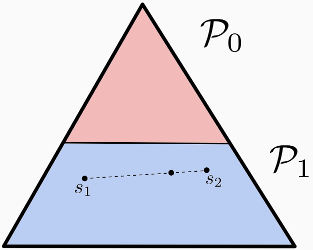

Revealing only that the signal is in $\{s_1, s_2\}$ would induce beliefs on the line segment joining the individual beliefs.

Revelation principle holds for EUM, fails for risk-conscious

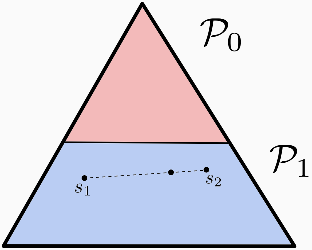

Coalescing signals does not change optimal action.

Revelation principle holds for EUM, fails for risk-conscious

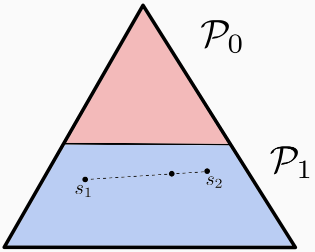

Coalescing signals does not change optimal action.

Coalesce all signals inducing same optimal action into a single **action recommendation**

Revelation principle holds for EUM, fails for risk-conscious

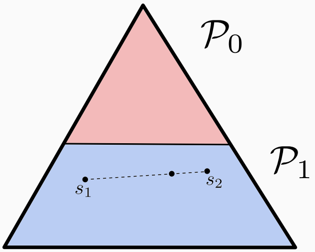

$\Longrightarrow$ One signal per action suffices (**action recommendation**)

Revelation principle holds for EUM, fails for risk-conscious

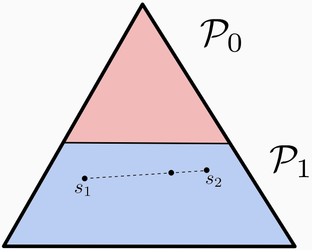

$\Longrightarrow$ One signal per action suffices (**action recommendation**)

$\Longrightarrow$ Sender's problem: Linear program with **obedience** constraints

Revelation principle holds for EUM, fails for risk-conscious

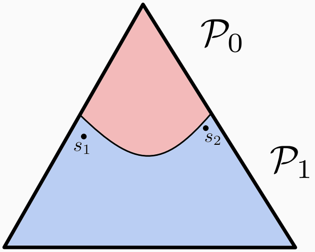

Revelation principle holds for EUM, fails for risk-conscious

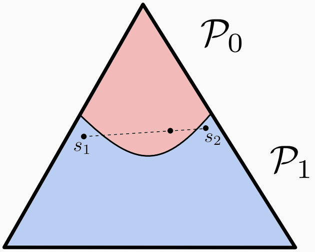

Coalescing signals inducing same action can change the receiver's optimal action

Revelation principle holds for EUM, fails for risk-conscious

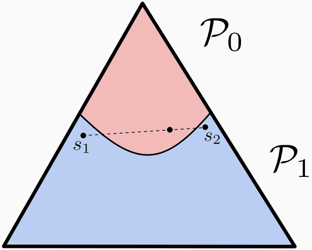

Coalescing signals inducing same action can change the receiver's optimal action

$\implies$ Revelation principle **fails.**

$\implies$ action recommendations may not suffice

$\implies$ No LP formulation with obedience constraints

Main Results

Theorem

The sender's optimal signaling mechanism can be found by solving a ***convex*** *optimization program.*

Theorem

At most $|\Omega|$ ***signals per action*** *suffice to optimally persuade.*

Distribution of beliefs $\longrightarrow$ mean-posterior over induced action

Belief recommendation $\Rightarrow$ optimize $\eta \in \Delta(\Delta(\Omega))$ subject to Bayes-plausibility: expectation of posteriors equals prior.

Let $\mathcal{P}_a$ be the set of beliefs that induce receiver action a.

We can decompose $\eta \in \Delta(\Delta(\Omega))$ as the weight b_a on each action a, and the corresponding distribution η_a over $\mathcal{P}_a$.

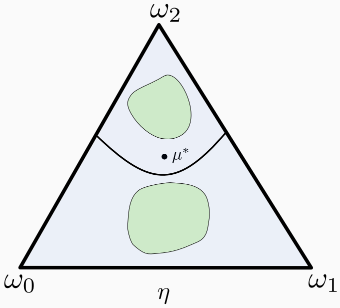

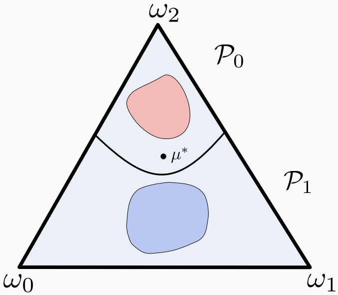

Distribution of beliefs → mean-posterior over induced action

Let $m_a = \mathbf{E}\eta_a$ denote **mean-posterior** for action a under η.

Main insight: Instead of optimizing over η, optimize directly over the mean posteriors m_a and multipliers b_a.

- Sender only cares about which action is taken → only b_a.
- Bayes-plausibility means the only constraint on valid η is that the expectation over posteriors equals the prior: $\mu^* = \sum_a b_a m_a$.
- For any m_a there exists a corresponding η_a iff $m_a \in \text{CHull}(\mathcal{P}_a)$.

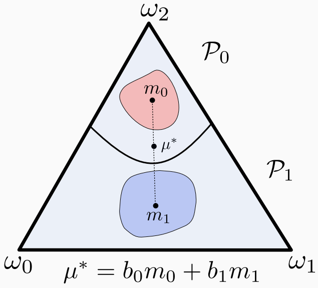

Main results

Theorem

*The sender's optimal signaling mechanism can be found by solving a **convex** optimization program.*

$$\max_{\{t_a : a \in A\}} \sum_{\omega \in \Omega} \sum_{a \in A} t_a(\omega) v(\omega, a)$$

$$\text{subject to} \quad \sum_{a \in A} t_a(\omega) = \mu^*(\omega), \quad \text{for each } \omega \in \Omega.$$

$$t_a \in \text{CHull}(\mathcal{P}_a \cup \{\mathbf{0}\}) \quad \text{for each } a \in A.$$

Here $t_a(\omega) = b_a m_a(\omega)$.

Main results

Theorem

*At most $|\Omega|$ **signals per action** suffice to optimally persuade.*

Proof: $m_a \in \mathrm{CHull}(\mathcal{P}_a)$ can always be represented as the convex combination of $|\Omega|$ points in $\mathcal{P}_a$. (**Caratheodory's theorem**).

Main results

Theorem

*The sender's optimal signaling mechanism can be found by solving a **convex** optimization program.*

$$\max_{\{t_a : a \in A\}} \sum_{\omega \in \Omega} \sum_{a \in A} t_a(\omega) v(\omega, a)$$

$$\text{subject to} \quad \sum_{a \in A} t_a(\omega) = \mu^*(\omega), \quad \text{for each } \omega \in \Omega.$$

$$t_a \in \text{CHull}(\mathcal{P}_a \cup \{\mathbf{0}\}) \quad \text{for each } a \in A.$$

Here $t_a(\omega) = b_a m_a(\omega)$. (Note: receiver enters this through $\mathcal{P}_a$ only.)

Binary persuasion

Binary persuasion: definition and characterization

Receiver must choose between action 0 or 1.

- join the queue or leave

Sender prefers action 1 over 0 in all states.

Assumption: receiver's utility is convex in her beliefs.

Binary persuasion: definition and characterization

Receiver must choose between action 0 or 1.

- join the queue or leave

Sender prefers action 1 over 0 in all states.

Assumption: receiver's utility is convex in her beliefs.

Results:

- Linear programming formulation
- Canonical set of signals:
  - either reveal the state or induce uncertainty between two states.

Binary persuasion: join/leave queue example

Customers must decide whether to wait in a queue for a service

- Utility $= \tau - \mathbf{E}[T] - \beta\sqrt{\mathbf{Var}(T)}$ $\quad$ (T : waiting time)
- This is convex in belief $\Rightarrow$ binary persuasion.

Binary persuasion: join/leave queue example

Customers must decide whether to wait in a queue for a service

- Utility $= \tau - \mathbf{E}[T] - \beta\sqrt{\mathbf{Var}(T)}$ (T : waiting time)
- This is convex in belief $\Rightarrow$ binary persuasion.

The optimal signaling mechanism has a "sandwich structure."

- Join signals are ordered from high-mean, high-variance to low-mean, low-variance, all with equal utility.

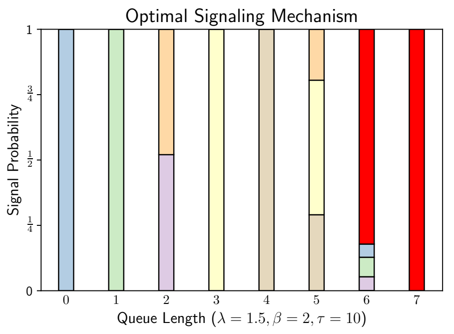

Full persuasion

Full persuasion

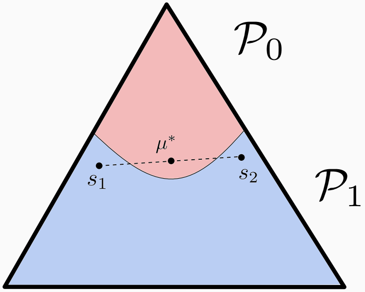

Under prior μ^*, the receiver chooses action 0.

With optimal persuasion, the receiver's belief moves to s_1 or s_2.

$\implies$ A risk-conscious receiver can be **fully** persuaded to action 1.

This is not possible for an EUM receiver.

Beyond persuading one risk-conscious receiver

Other interpretations: public persuasion and robust persuasion

Instead of taking receiver's utility as primitive, one can choose $\mathcal{P}_a$ as a model primitive.

The key is that our "receiver" has one posterior.

$$
\begin{aligned}
&\max_{\{t_a : a \in A\}} \sum_{\omega \in \Omega} \sum_{a \in A} t_a(\omega) v(\omega, a) \\
&\text{subject to} \quad \sum_{a \in A} t_a(\omega) = \mu^*(\omega), \quad \text{for each } \omega \in \Omega. \\
&\qquad\qquad t_a \in \mathrm{CHull}(\mathcal{P}_a \cup \{\mathbf{0}\}) \quad \text{for each } a \in A.
\end{aligned}
$$

Here $t_a(\omega) = b_a m_a(\omega)$. (Note: receiver enters this through $\mathcal{P}_a$ only.)

Other interpretations: public persuasion and robust persuasion

Instead of taking receiver's utility as primitive, one can choose $\mathcal{P}_a$ as a model primitive.

The key is that our "receiver" has one posterior.

A "receiver" could be a group of agents choosing an action collectively, and sender seeks to **publicly** persuade.

Other interpretations: public persuasion and robust persuasion

Instead of taking receiver's utility as primitive, one can choose $\mathcal{P}_a$ as a model primitive.

The key is that our "receiver" has one posterior.

A "receiver" could be a group of agents choosing an action collectively, and sender seeks to **publicly** persuade.

Receiver has private type, and the sender seeks to **robustly** persuade: sender seeks to persuade each type of receiver.

Conclusion

We extend the theory of persuasion to be include more realistic models of human behavior.

More sophisticated methods of persuasion are necessary.

- Convex program
- Bound the number of signals

We find structural regularity in binary persuasion.

- Linear program
- Canonical set of signals – between at most two states
- Sandwich structure numerically for the queue problem

We can do public persuasion and robust persuasion.