Announcements

- ➢HW3 is out, due next Saturday.
- ➢3 weeks away from Project presentation!

✓Last few weeks: ML for Economic Problems

❖ From today: Economic Aspects of ML

CMSC 35401: The Interplay of Economics and ML (Winter 2024)

Bayesian Persuasion (a.k.a. Information Design)

Instructor: Haifeng Xu

- ML is about extracting information from data

- The next-step question: when you have information, how to use/exploit it? What’s the value of it?
  - Related to manipulate features to game a learning algorithm (later lectures)

Outline

- Introduction and Bayesian Persuasion
- Algorithms for Bayesian Persuasion
- Persuading Multiple Receivers

Two Primary Ways to Influence Agents' Behaviors

- Design/provide incentives
  - Auctions

You only pay the
second highest bid!

Two Primary Ways to Influence Agents' Behaviors

- Design/provide incentives
  - Auctions
  - Discounts/coupons

Two Primary Ways to Influence Agents' Behaviors

- Design/provide incentives
  - Auctions
  - Discounts/coupons
  - Job contract design

Bonus depends on performance, and is up to $1M!

Two Primary Ways to Influence Agents' Behaviors

- Design/provide incentives
  - Auctions
  - Discounts/coupons
  - Job contract design

Mechanism Design

- Influence agents' beliefs
  - Deception in wars/battles

All warfare is based on deception. Hence, when we are able to attack, we must seem unable; when using our forces, we must appear inactive…

-- Sun Tzu, *The Art of War*

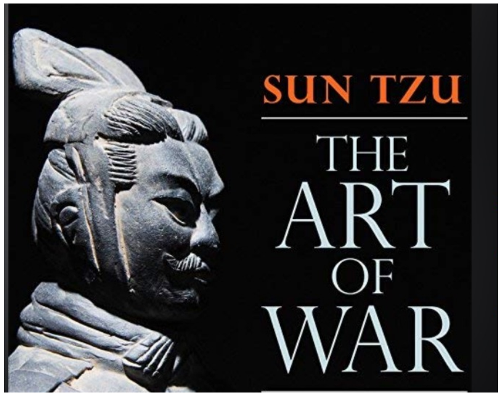

Two Primary Ways to Influence Agents' Behaviors

- Design/provide incentives
  - Auctions
  - Discounts/coupons
  - Job contract design

Mechanism Design

- Influence agents' beliefs
  - Deception in wars/battles
  - Strategic information disclosure

Strategic inventory information disclosure

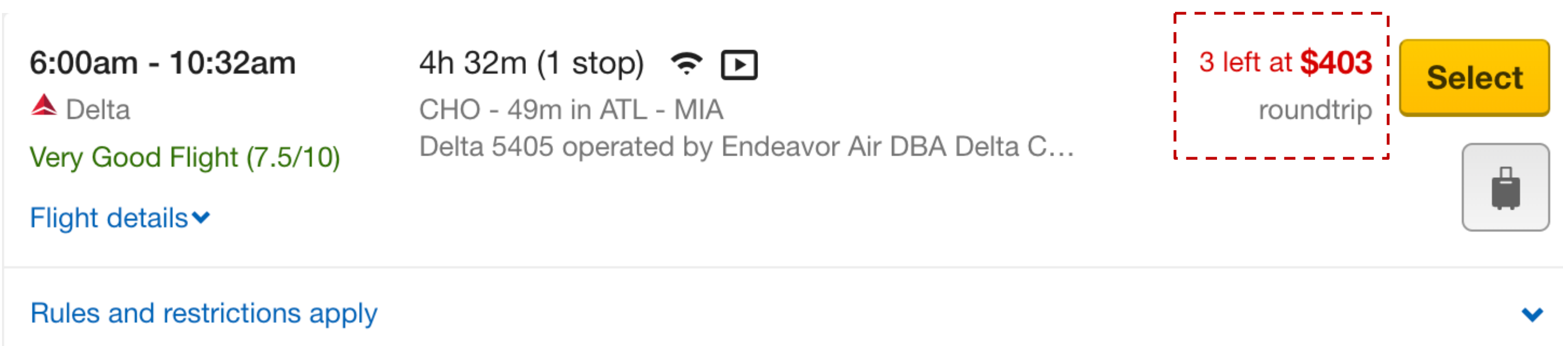

Two Primary Ways to Influence Agents' Behaviors

- Design/provide incentives
  - Auctions
  - Discounts/coupons
  - Job contract design

Mechanism Design

- Influence agents' beliefs
  - Deception in wars/battles
  - Strategic information disclosure

Strategic inventory information disclosure

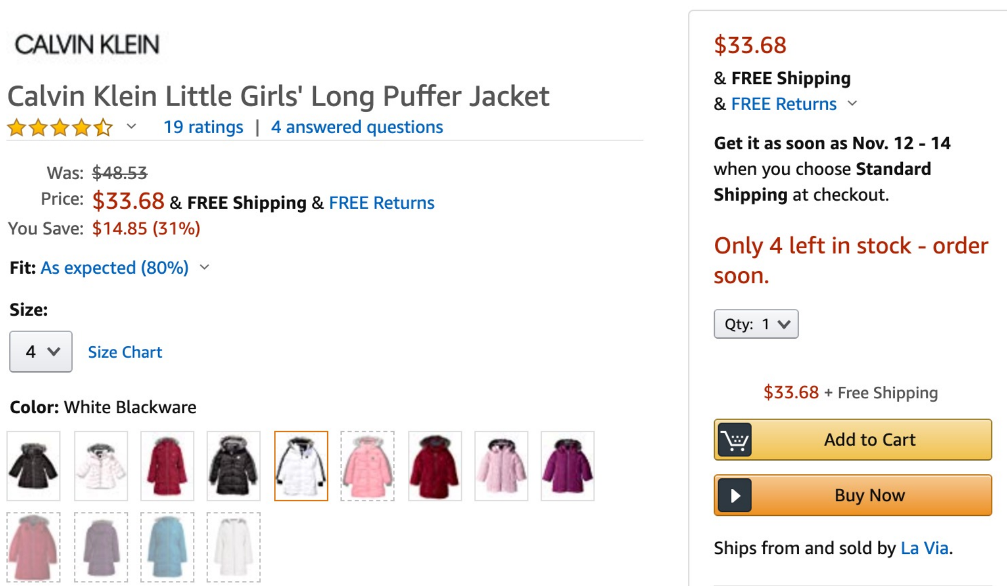

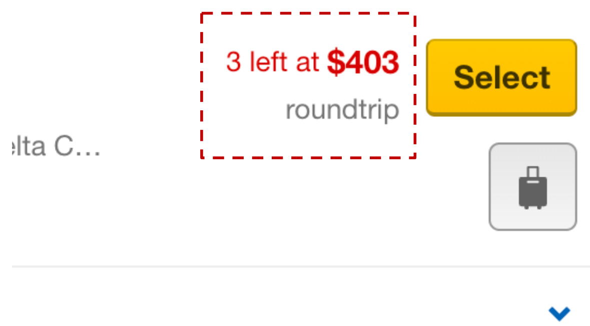

Two Primary Ways to Influence Agents' Behaviors

- Design/provide incentives
  - Auctions
  - Discounts/coupons
  - Job contract design

Mechanism Design

- Influence agents' beliefs
  - Deception in wars/battles
  - Strategic information disclosure
  - News articles, advertising, tweets, etc.

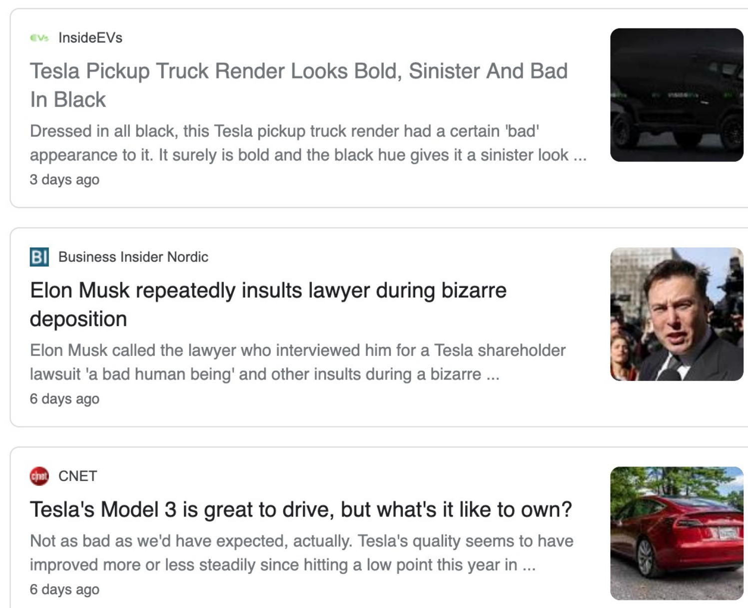

Two Primary Ways to Influence Agents' Behaviors

- Design/provide incentives
  - Auctions
  - Discounts/coupons
  - Job contract design

Mechanism Design

- Influence agents' beliefs
  - Deception in wars/battles
  - Strategic information disclosure
  - News articles, advertising, tweets …
  - In fact, most information you see is there with a purpose

Persuasion (information design)

A whole course from Booth on this topic

Persuasion is the act of exploiting an informational advantage in order to influence the decisions of others

- Intrinsic in human activities: advertising, negotiation, politics, security, marketing, financial regulation,…
- A large body of research

One Quarter of GDP Is Persuasion

By DONALD MCCLOSKEY AND ARJO KLAMER*

— The American Economic Review Vol. 85, No. 2, 1995.

Example: Recommendation Letters

- Advisor vs. recruiter
- 1/3 of the advisor's students are excellent; 2/3 are average
- A fresh graduate is randomly drawn from this population
- Recruiter
  - Utility $1 + \epsilon$ for hiring an excellent student; -1 for an average student
  - Utility 0 for not hiring
  - A-priori, only knows the advisor's student population

$$\underbrace{(1+\epsilon)\times 1/3 - 1\times 2/3}_{\textit{hiring}} \qquad < \qquad \underbrace{0}_{\textit{Not hiring}}$$

Example: Recommendation Letters

- Advisor vs. recruiter
- 1/3 of the advisor's students are excellent; 2/3 are average
- A fresh graduate is randomly drawn from this population
- Recruiter
  - Utility $1 + \epsilon$ for hiring an excellent student; -1 for an average student
  - Utility 0 for not hiring
  - A-priori, only knows the advisor's student population
- Advisor
  - Utility 1 if the student is hired, 0 otherwise
  - Knows whether the student is excellent or not

Example: Recommendation Letters

What is the advisor's optimal "recommendation strategy"?

- ➢ Attempt 1: always say "excellent" (equivalently, no information)
  - Recruiter ignores the recommendation
  - Advisor expected utility 0

Remark

Assume advisor "commits" to some policy, and recruiter is fully aware this policy and will best respond

Example: Recommendation Letters

What is the advisor’s optimal “recommendation strategy”?

➢ Attempt 2: honest recommendation (i.e., full information)

- Advisor expected utility 1/3

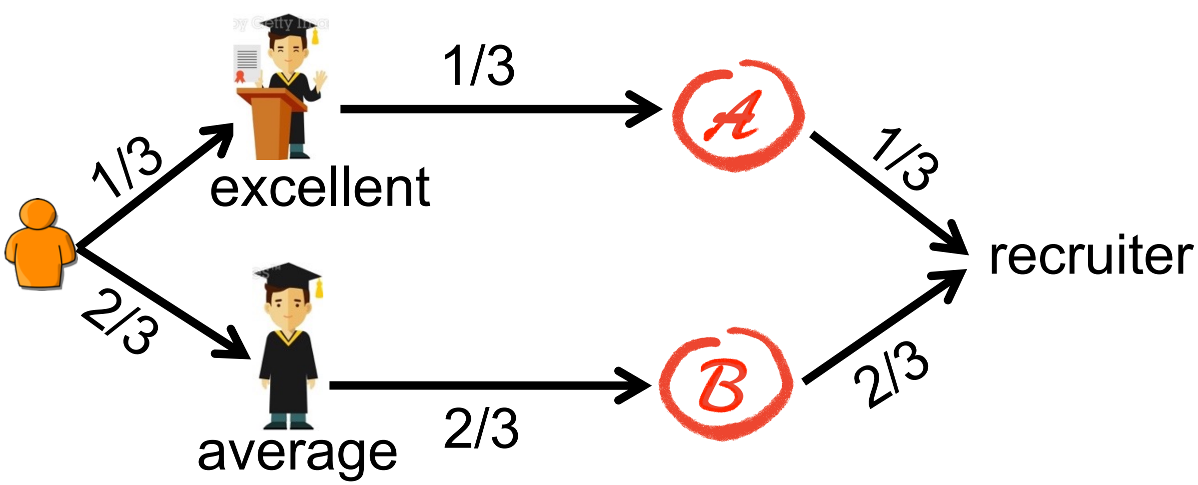

Example: Recommendation Letters

What is the advisor's optimal "recommendation strategy"?

- Attempt 3: noisy information → advisor expected utility 2/3

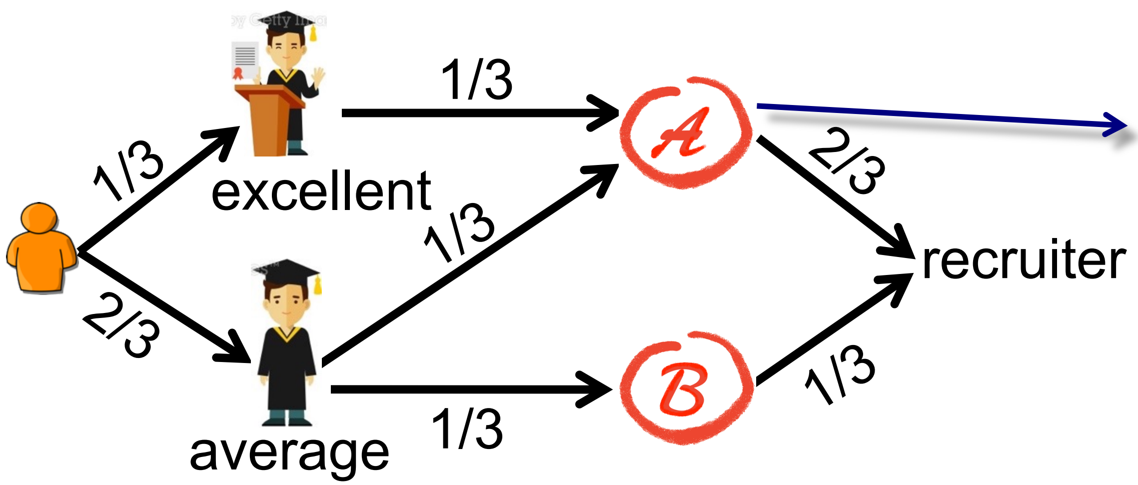

P(excellent | $\mathcal{A}$) = $1/2$

$(1+\epsilon - 1)/2 \quad > \quad 0$

Hiring Not hiring

Model of Bayesian Persuasion

- Two players: persuader (Sender, she), decision maker (Receiver he)
  - Previous example: advisor = sender, recruiter = receiver
- Receiver looks to take an action $i \in [n] = \{1, 2, \dots, n\}$
  - Receiver utility $r(i, \theta)$
  - Sender utility $s(i, \theta)$

$\theta \in \Theta$ is a random state of nature

- Both players know $\theta \sim prior\ dist.\ \mu$, but Sender has an informational advantage – she can observe realization of θ
- Sender wants to strategically reveal info about θ to "persuade" Receiver to take an action she likes
  - Concealing or revealing all info is not necessarily the best

Well…how to reveal partial information?

Revealing Information via Signaling

Definition: A signaling scheme is a mapping $\pi: \Theta \rightarrow \Delta_\Sigma$ where Σ is the set of all possible signals.

π is fully described by $\{\pi(\sigma, \theta)\}_{\theta\in\Theta, \sigma\in\Sigma}$ where $\pi(\sigma, \theta) =$ prob. of sending σ when observing θ (so $\sum_{\sigma\in\Sigma} \pi(\sigma, \theta) = 1$ for any θ)

Note:

- ✓ Statistically, π just creates a random variable σ that correlates with θ
- ✓ π is public knowledge, thus known by Receiver

Example

- ➢ $\Theta = \{Excellent, Average\}$, $\Sigma = \{A, B\}$
- ➢ $\pi(A, Average) = 1/2$

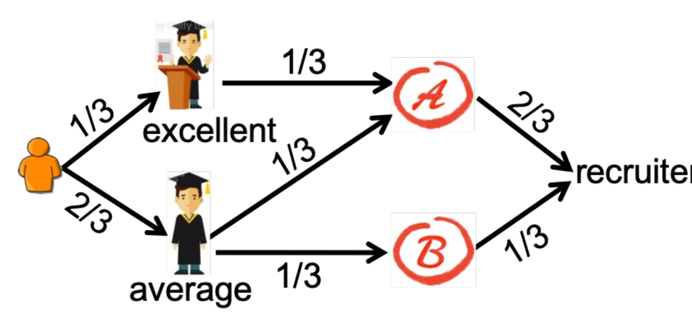

Revealing Information via Signaling

Definition: A signaling scheme is a mapping $\pi: \Theta \to \Delta_\Sigma$ where Σ is the set of all possible signals.

π is fully described by $\{\pi(\sigma,\theta)\}_{\theta\in\Theta,\sigma\in\Sigma}$ where $\pi(\sigma,\theta) =$ prob. of sending σ when observing θ (so $\sum_{\sigma\in\Sigma}\pi(\sigma,\theta) = 1$ for any θ)

What can Receiver infer about θ after receiving σ?

Bayes updating:

$$\Pr(\theta|\sigma) = \frac{\pi(\sigma,\theta)\cdot\mu(\theta)}{\sum_{\theta'}\pi(\sigma,\theta')\cdot\mu(\theta')}$$

$$\Pr(excellent|A) = 1/2$$

π is fully described by $\{\pi(\sigma,\theta)\}_{\theta\in\Theta,\sigma\in\Sigma}$ where $\pi(\sigma,\theta) =$ prob. of sending σ when observing θ (so $\sum_{\sigma\in\Sigma}\pi(\sigma,\theta) = 1$ for any θ)

Revealing Information via Signaling

Definition: A signaling scheme is a mapping $\pi: \Theta \rightarrow \Delta_\Sigma$ where Σ is the set of all possible signals.

π is fully described by $\{\pi(\sigma, \theta)\}_{\theta \in \Theta, \sigma \in \Sigma}$ where $\pi(\sigma, \theta) =$ prob. of sending σ when observing θ (so $\sum_{\sigma \in \Sigma} \pi(\sigma, \theta) = 1$ for any θ)

Would such noisy information benefit Receiver?

- Expected Receiver utility conditioned on σ:

$$R(\sigma) = \max_{i \in [n]} \left[\sum_{\theta \in \Theta} r(i, \theta) \cdot \frac{\pi(\sigma,\theta) \cdot \mu(\theta)}{\sum_{\theta'} \pi(\sigma,\theta') \cdot \mu(\theta')} \right]$$

- $\Pr(\sigma) = \sum_{\theta'} \pi(\sigma, \theta') \cdot \mu(\theta')$

Revealing Information via Signaling

Definition: A signaling scheme is a mapping $\pi: \Theta \to \Delta_\Sigma$ where Σ is the set of all possible signals.

π is fully described by $\{\pi(\sigma, \theta)\}_{\theta \in \Theta, \sigma \in \Sigma}$ where $\pi(\sigma, \theta) =$ prob. of sending σ when observing θ (so $\sum_{\sigma \in \Sigma} \pi(\sigma, \theta) = 1$ for any θ)

Would such noisy information benefit Receiver?

- Expected Receiver utility conditioned on σ:

$$R(\sigma) = \max_{i \in [n]} \left[\sum_{\theta \in \Theta} r(i, \theta) \cdot \frac{\pi(\sigma,\theta) \cdot \mu(\theta)}{\sum_{\theta'} \pi(\sigma,\theta') \cdot \mu(\theta')} \right]$$

- $\Pr(\sigma) = \sum_{\theta'} \pi(\sigma, \theta') \cdot \mu(\theta')$

Revealing Information via Signaling

Definition: A signaling scheme is a mapping $\pi: \Theta \to \Delta_\Sigma$ where Σ is the set of all possible signals.

π is fully described by $\{\pi(\sigma,\theta)\}_{\theta\in\Theta,\sigma\in\Sigma}$ where $\pi(\sigma,\theta) =$ prob. of sending σ when observing θ (so $\sum_{\sigma\in\Sigma}\pi(\sigma,\theta) = 1$ for any θ)

Would such noisy information benefit Receiver?

- Expected Receiver utility conditioned on σ:

$$R(\sigma) = \max_{i\in[n]} \left[\sum_{\theta\in\Theta} r(i,\theta) \cdot \frac{\pi(\sigma,\theta)\cdot\mu(\theta)}{\sum_{\theta'}\pi(\sigma,\theta')\cdot\mu(\theta')} \right]$$

- $\Pr(\sigma) = \sum_{\theta'}\pi(\sigma,\theta')\cdot\mu(\theta')$
  - $\Pr(\sigma)\cdot R(\sigma) = \max_i \sum_{\theta\in\Theta} r(i,\theta)\cdot\pi(\sigma,\theta)\cdot\mu(\theta)$ (a convex function of π)

- Expected Receiver utility under π: $\sum_\sigma \Pr(\sigma)\cdot R(\sigma)$

Revealing Information via Signaling

Fact. Receiver's expected utility (weakly) increases under any signaling scheme π.

Proof:

- Expected Receiver utility under π: $\sum_\sigma \Pr(\sigma) \cdot R(\sigma)$

Revealing Information via Signaling

Fact. Receiver's expected utility (weakly) increases under any signaling scheme π.

Proof:

- Expected Receiver utility under π: $\boxed{\sum_\sigma \Pr(\sigma) \cdot R(\sigma)}$

$$\sum_\sigma \Pr(\sigma) \cdot R(\sigma) = \sum_\sigma \max_{i \in [n]} \left[\sum_{\theta \in \Theta} r(i, \theta) \cdot \pi(\sigma, \theta) \cdot \mu(\theta)\right]$$

Revealing Information via Signaling

Fact. Receiver's expected utility (weakly) increases under any signaling scheme π.

Proof:

- Expected Receiver utility under π: $\boxed{\sum_{\sigma} \Pr(\sigma) \cdot R(\sigma)}$

$$\sum_{\sigma} \Pr(\sigma) \cdot R(\sigma) = \sum_{\sigma} \max_{i\in[n]} \left[\sum_{\theta\in\Theta} r(i,\theta) \cdot \pi(\sigma,\theta) \cdot \mu(\theta)\right]$$

$$\geq \max_{i\in[n]} \sum_{\sigma} \left[\sum_{\theta\in\Theta} r(i,\theta) \cdot \pi(\sigma,\theta) \cdot \mu(\theta)\right]$$

By HW3 problem 1

Revealing Information via Signaling

Fact. Receiver's expected utility (weakly) increases under any signaling scheme π.

Proof:

- Expected Receiver utility under π: $\boxed{\sum_\sigma \Pr(\sigma) \cdot R(\sigma)}$

$$\sum_\sigma \Pr(\sigma) \cdot R(\sigma) = \sum_\sigma \max_{i\in[n]} \left[\sum_{\theta\in\Theta} r(i,\theta) \cdot \pi(\sigma,\theta) \cdot \mu(\theta)\right]$$

$$\geq \max_{i\in[n]} \sum_\sigma \left[\sum_{\theta\in\Theta} r(i,\theta) \cdot \pi(\sigma,\theta) \cdot \mu(\theta)\right]$$

$$= \max_{i\in[n]} \sum_{\theta\in\Theta} r(i,\theta) \cdot \left(\sum_\sigma \pi(\sigma,\theta)\right) \cdot \mu(\theta)$$

$$= \max_{i\in[n]} \sum_{\theta\in\Theta} r(i,\theta) \cdot \mu(\theta)$$

Best expected receiver utility without information

Revealing Information via Signaling

Fact. Receiver's expected utility (weakly) increases under any signaling scheme π.

Remarks:

- Signaling scheme does increase Receiver's utility
- More (even noisy) information always helps a decision maker (DM)
  - Not true if multiple decision makers (will see examples later)

Corollary. Receiver's expected utility is maximized when Sender reveals full info, i.e., directly revealing the realized θ.

Because any other noisy scheme π can be improved by further revealing θ itself

Revealing Information via Signaling

Fact. Receiver's expected utility (weakly) increases under any signaling scheme π.

Remarks:

- Signaling scheme does increase Receiver's utility
- More (even noisy) information always helps a decision maker (DM)
  - Not true if multiple decision makers (will see examples later)

Corollary. Receiver's expected utility is maximized when Sender reveals full info, i.e., directly revealing the realized θ.

But this is not Sender's goal…

Sender Objective: maximize her own expected utility by picking π

Outline

- Introduction and Bayesian Persuasion
- Algorithms for Bayesian Persuasion
- Persuading Multiple Receivers

Q: What are obstacles when designing $\pi = \{\pi(\theta, \sigma)\}_{\theta \in \Theta, \sigma \in \Sigma}$?

- The set of all possible signals Σ is unclear and maybe too large
  - Too many possible signals to choose from (think about how many ways Amazon can reveal information to you)
- Key observation: a signal is mathematically nothing but a posterior distribution over Θ
  - Recall the Bayes updates: $\Pr(\theta|\sigma) = \frac{\pi(\sigma,\theta)\cdot\mu(\theta)}{\sum_{\theta'} \pi(\sigma,\theta')\cdot\mu(\theta')}$
- It turns out that n signals suffice

"Revelation Principle"

Fact. There always exists an optimal signaling scheme that uses at most n(= # receiver actions) signals, where signal σ_i induce optimal Receiver action i

- Conditioned on any signal σ
  - Receiver infers $\Pr(\theta|\sigma) = \frac{\pi(\sigma,\theta)\cdot\mu(\theta)}{\sum_{\theta'}\pi(\sigma,\theta')\cdot\mu(\theta')}$
  - Receiver takes optimal action $i^* = \arg\max_{i\in[n]} \sum_\theta \Pr(\theta|\sigma)\, r(i,\theta)$

- If two signal σ and σ' result in the same best action i^*, Sender can combine them as a single signal $\sigma_{i^*} = (\sigma, \sigma')$
  - Claim: i^* is still the optimal action conditioned on σ_{i^*}

$$\sum_\theta \Pr(\theta|\sigma)\, r(i^*,\theta) \ge \sum_\theta \Pr(\theta|\sigma)\, r(i,\theta), \quad \forall i \qquad \times p$$

$$\sum_\theta \Pr(\theta|\sigma')\, r(i^*,\theta) \ge \sum_\theta \Pr(\theta|\sigma')\, r(i,\theta), \quad \forall i \qquad \times q$$

$$\Rightarrow \quad \sum_\theta [\Pr(\theta|\sigma)p + \Pr(\theta|\sigma')q] r(i^*,\theta)$$
$$\ge \sum_\theta [\Pr(\theta|\sigma)p + \Pr(\theta|\sigma')q] r(i,\theta), \quad \forall i$$

Revelation Principle

Fact. There always exists an optimal signaling scheme that uses at most n(= # receiver actions) signals, where signal σ_i induce optimal Receiver action i

- Conditioned on any signal σ
  - Receiver infers $\Pr(\theta|\sigma) = \frac{\pi(\sigma,\theta)\cdot\mu(\theta)}{\sum_{\theta'}\pi(\sigma,\theta')\cdot\mu(\theta')}$
  - Receiver takes optimal action $i^* = \arg\max_{i\in[n]} \sum_\theta \Pr(\theta|\sigma)\, r(i,\theta)$
- If two signal σ and σ' result in the same best action i^*, Sender can combine them as a single signal $\sigma_{i^*} = (\sigma, \sigma')$
  - Claim: i^* is still the optimal action conditioned on σ_{i^*}

$$\sum_\theta \Pr(\theta|\sigma)\, r(i^*,\theta) \geq \sum_\theta \Pr(\theta|\sigma)\, r(i,\theta), \quad \forall i$$

$$\sum_\theta \Pr(\theta|\sigma')\, r(i^*,\theta) \geq \sum_\theta \Pr(\theta|\sigma')\, r(i,\theta), \quad \forall i$$

$$\Rightarrow \quad \sum_\theta [\Pr(\theta|\sigma)p + \Pr(\theta|\sigma')q] r(i^*,\theta) \geq \sum_\theta [\Pr(\theta|\sigma)p + \Pr(\theta|\sigma')q] r(i,\theta), \quad \forall i$$

$\Pr(\theta|\sigma_{i^*})$ is a linear combination of $\Pr(\theta|\sigma)$ and $\Pr(\theta|\sigma')$

Revelation Principle

Fact. There always exists an optimal signaling scheme that uses at most n(= # receiver actions) signals, where signal σ_i induce optimal Receiver action i

- Conditioned on any signal σ
  - Receiver infers $\Pr(\theta|\sigma) = \frac{\pi(\sigma,\theta)\cdot\mu(\theta)}{\sum_{\theta'}\pi(\sigma,\theta')\cdot\mu(\theta')}$
  - Receiver takes optimal action $i^* = \arg\max_{i\in[n]} \sum_\theta \Pr(\theta|\sigma)\ r(i,\theta)$
- If two signal σ and σ' result in the same best action i^*, Sender can combine them as a single signal $\sigma_{i^*} = (\sigma, \sigma')$
  - Claim: i^* is still the optimal action conditioned on σ_{i^*}
  - Both players' utilities did not change as receiver still takes i^* as Sender wanted
- Can merge all signals with optimal receiver action i^* as a single signal σ_{i^*}

Revelation Principle

Fact. There always exists an optimal signaling scheme that uses at most n(= # receiver actions) signals, where signal σ_i induce optimal Receiver action i

- Each σ_i can be viewed as an action recommendation of i (this should remind you correlated equilibrium)

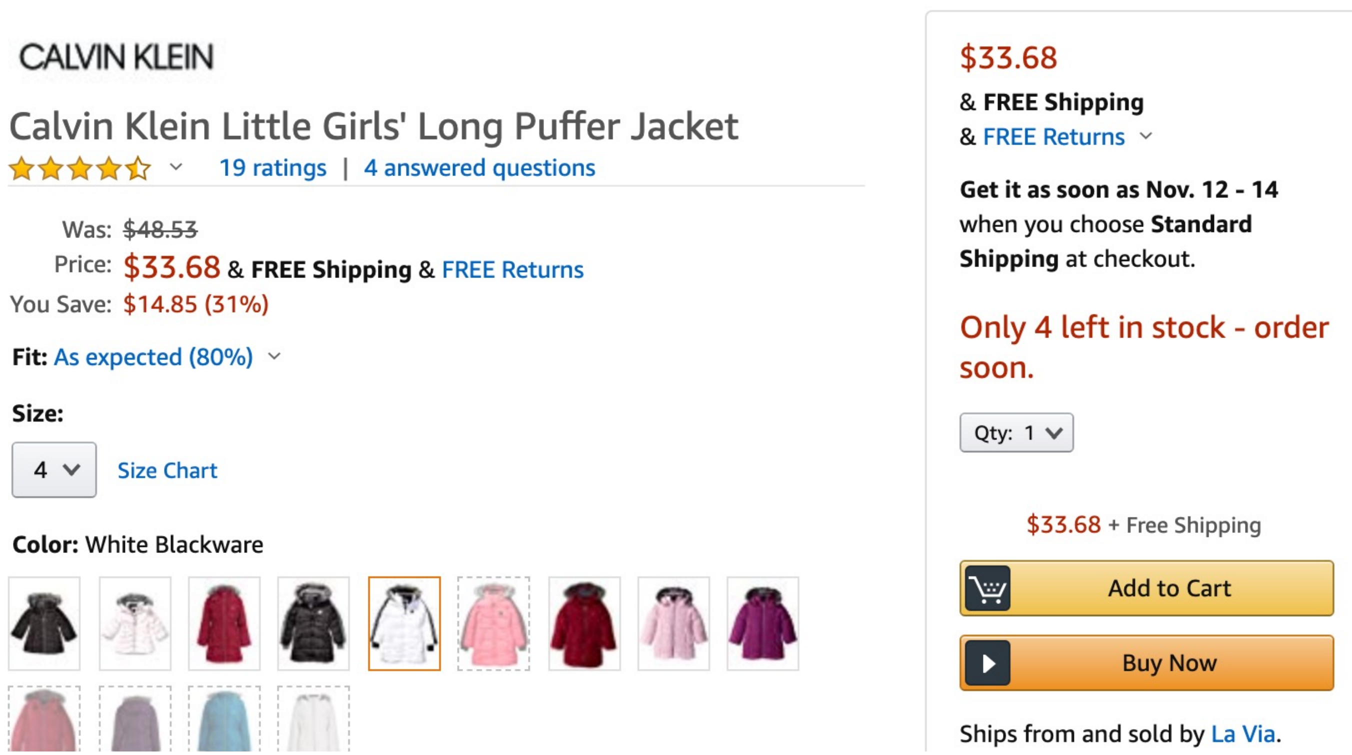

Optimal Persuasion via Linear Program

- Input: prior μ, sender payoff $s(i,\theta)$, receiver payoff $r(i,\theta)$
- Variables: $\pi(\sigma_i,\theta)$

σ_i indeed incentivizes Receiver best action i

$$
\begin{array}{lll}
\max & \sum_{\theta\in\Theta}\sum_{i=1}^{n} s(i,\theta)\cdot\pi(\sigma_i,\theta)\mu(\theta) & \\
\text{s.t.} & \boxed{\sum_{\theta\in\Theta} r(i,\theta)\cdot\pi(\sigma_i,\theta)\mu(\theta) \geq \sum_{\theta\in\Theta} r(j,\theta)\cdot\pi(\sigma_i,\theta)\mu(\theta),} & \text{for } i,j\in[n]. \\
& \sum_{i=1}^{n}\pi(\sigma_i,\theta) = 1, & \text{for } \theta\in\Theta. \\
& \pi(\sigma_i,\theta)\geq 0, & \text{for } \theta\in\Theta, i\in[n].
\end{array}
$$

Optimal Persuasion via Linear Program

- Input: prior μ, sender payoff $s(i,\theta)$, receiver payoff $r(i,\theta)$
- Variables: $\pi(\sigma_i,\theta)$

$$
\begin{array}{lll}
\max & \sum_{\theta\in\Theta}\sum_{i=1}^{n} s(i,\theta)\cdot\pi(\sigma_i,\theta)\mu(\theta) & \\
\text{s.t.} & \sum_{\theta\in\Theta} r(i,\theta)\cdot\pi(\sigma_i,\theta)\mu(\theta) \geq \sum_{\theta\in\Theta} r(j,\theta)\cdot\pi(\sigma_i,\theta)\mu(\theta), & \text{for } i,j\in[n]. \\
& \sum_{i=1}^{n}\pi(\sigma_i,\theta) = 1, & \text{for } \theta\in\Theta. \\
& \pi(\sigma_i,\theta)\geq 0, & \text{for } \theta\in\Theta, i\in[n].
\end{array}
$$

π is a valid signaling scheme

Optimal Persuasion via Linear Program

- Input: prior μ, sender payoff $s(i, \theta)$, receiver payoff $r(i, \theta)$
- Variables: $\pi(\sigma_i, \theta)$

Sender expected utility
(we know Receiver will take i at signal σ_i)

$$
\begin{aligned}
\max \quad & \boxed{\textstyle\sum_{\theta \in \Theta} \sum_{i=1}^{n} s(i,\theta) \cdot \pi(\sigma_i, \theta)\mu(\theta)} & \\
\text{s.t.} \quad & \textstyle\sum_{\theta \in \Theta} r(i,\theta) \cdot \pi(\sigma_i,\theta)\mu(\theta) \geq \sum_{\theta \in \Theta} r(j,\theta) \cdot \pi(\sigma_i,\theta)\mu(\theta), & \text{for } i, j \in [n]. \\
& \textstyle\sum_{i=1}^{n} \pi(\sigma_i, \theta) = 1, & \text{for } \theta \in \Theta. \\
& \pi(\sigma_i, \theta) \geq 0, & \text{for } \theta \in \Theta, i \in [n].
\end{aligned}
$$

Optimal Persuasion via Linear Program

- Input: prior μ, sender payoff $s(i,\theta)$, receiver payoff $r(i,\theta)$
- Variables: $\pi(\sigma_i,\theta)$

$$
\begin{aligned}
\max \quad & \textstyle\sum_{\theta\in\Theta}\sum_{i=1}^{n} s(i,\theta)\cdot\pi(\sigma_i,\theta)\mu(\theta) & \\
\text{s.t.} \quad & \textstyle\sum_{\theta\in\Theta} r(i,\theta)\cdot\pi(\sigma_i,\theta)\mu(\theta) \geq \sum_{\theta\in\Theta} r(j,\theta)\cdot\pi(\sigma_i,\theta)\mu(\theta), & \text{for } i,j\in[n]. \\
& \textstyle\sum_{i=1}^{n}\pi(\sigma_i,\theta)=1, & \text{for } \theta\in\Theta. \\
& \pi(\sigma_i,\theta)\geq 0, & \text{for } \theta\in\Theta, i\in[n].
\end{aligned}
$$

This should remind you the LP for correlated equilibria

Outline

- Introduction and Bayesian Persuasion
- Algorithms for Bayesian Persuasion
- Persuading Multiple Receivers

Recommendation Letter Example cont'd


- Advisor vs. two fellowship programs
- 1/3 of the advisor's students are excellent; 2/3 are average
- A fresh graduate is randomly drawn from this population
- Each fellowship:
  - Utility $1 + \epsilon$ for awarding excellent student; -1 for average student
  - Utility 0 for no award
  - A-priori, only knows the advisor's student population
  - Student can accept both fellowships
- Advisor
  - Utility 1 if student gets at least one fellowship, 0 otherwise
  - Knows whether the student is excellent or not

Recommendation Letter Example cont'd

What is the advisor's optimal "recommendation strategy"?

Well, we learned the lesson — noisy info!

Recommendation Letter Example cont'd

What is the advisor's optimal "recommendation strategy"?

➢ Optimal public scheme → advisor expected utility 2/3

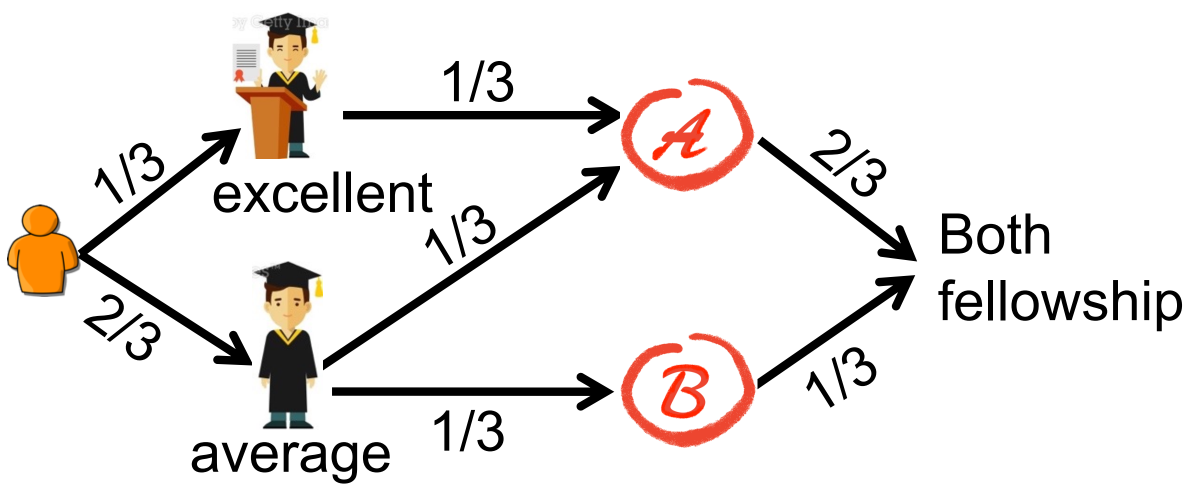

Recommendation Letter Example cont'd

What is the advisor's optimal "recommendation strategy"?

- Optimal private scheme → advisor expected utility 1

excellent

Recommendation Letter Example cont'd

What is the advisor's optimal "recommendation strategy"?

- Optimal private scheme → advisor expected utility 1
- Conditioned on "strong", excellent with prob ½
- Always at least one fellowship recommended "strong"

average

Recommendation Letter Example cont'd

Generalize this example to n fellowships:

advisor utility of optimal private scheme

$\geq \frac{n+1}{2}$ advisor utility of optimal pubic scheme

Conceptual Message

Being able to persuade privately may have a huge advantage

Remark: fellowship programs' utilities did not decrease

Thank You

Haifeng Xu

University of Chicago

haifengxu@uchicago.edu