Multi-Sender Persuasion
A Computational Perspective

Safwan Hossain*
Harvard University

Tonghan Wang*
Harvard University

Tao Lin*
Harvard University

Yiling Chen
Harvard University

David C. Parkes
Harvard University

Haifeng Xu
University of Chicago

Information Design

How can an agent with informational advantage, strategically reveal this information to another agent to influence their behaviour?

1. Bayesian Persuasion

- Two player game between a sender, who gets to observe a world state $\theta \in \Theta$, and a receiver who gets to take an action [1].
- The utility of both players depend on this action along with the world state.
  - Complete Information - sender knows receiver utility
- Both players share a common prior belief μ about the possible world states θ.
- The sender can commit to strategically revealing her knowledge of the world state through signaling.

- θ - student quality $\{$good, bad$\}$
- $a \in \{$hire, not hire$\}$

- Utility $u(a, \theta)$: +1 if student is hired
- Utility $w(a, \theta)$: +1 if hiring good student or not hiring bad ones

Sender - Professor

Receiver - Hiring Manager

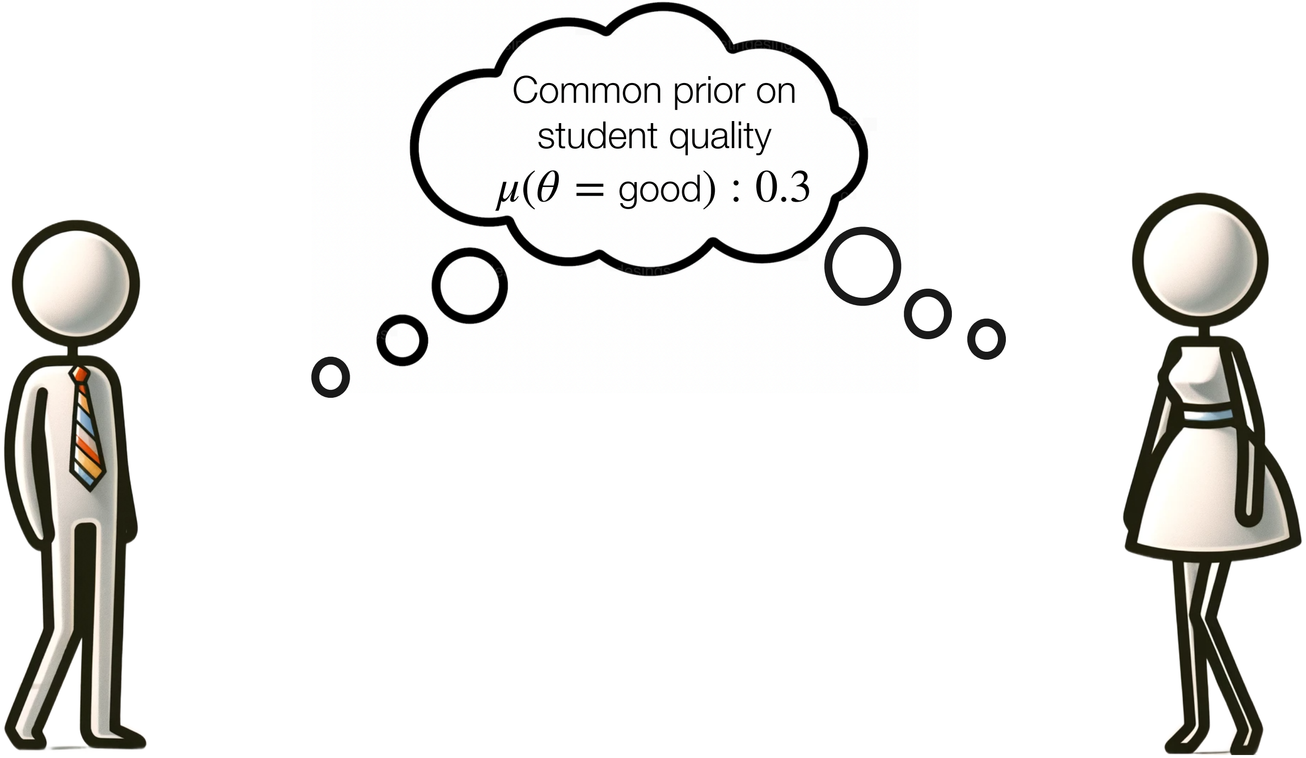

Sender - Professor

$u(a, \theta)$: +1 if student hired

Receiver - Hiring Manager

$w(a, \theta)$: +1 if right decision is made

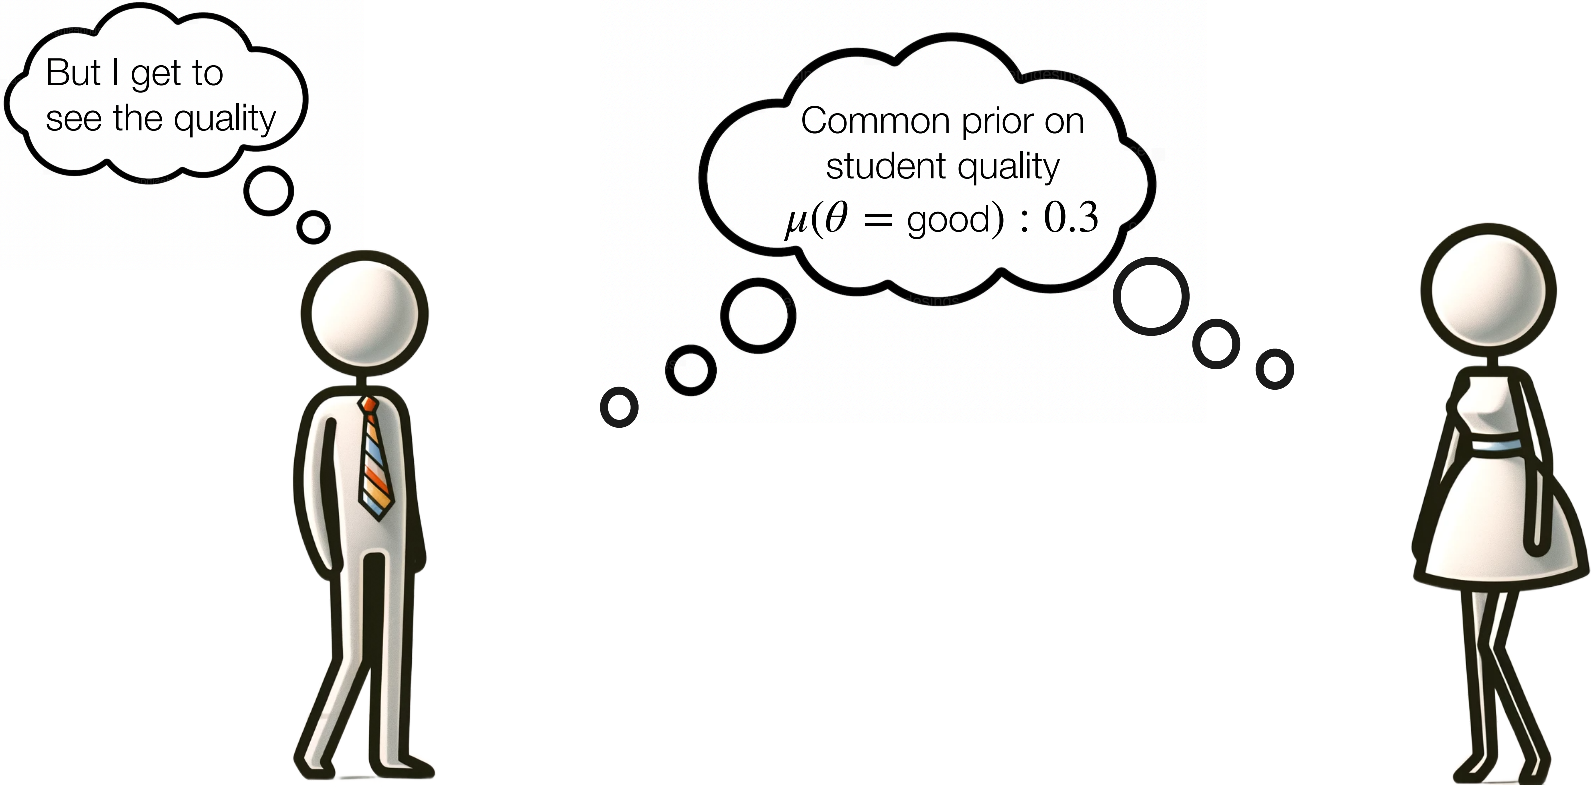

Sender - Professor

$u(a, \theta)$: +1 if student hired

Receiver - Hiring Manager

$w(a, \theta)$: +1 if right decision is made

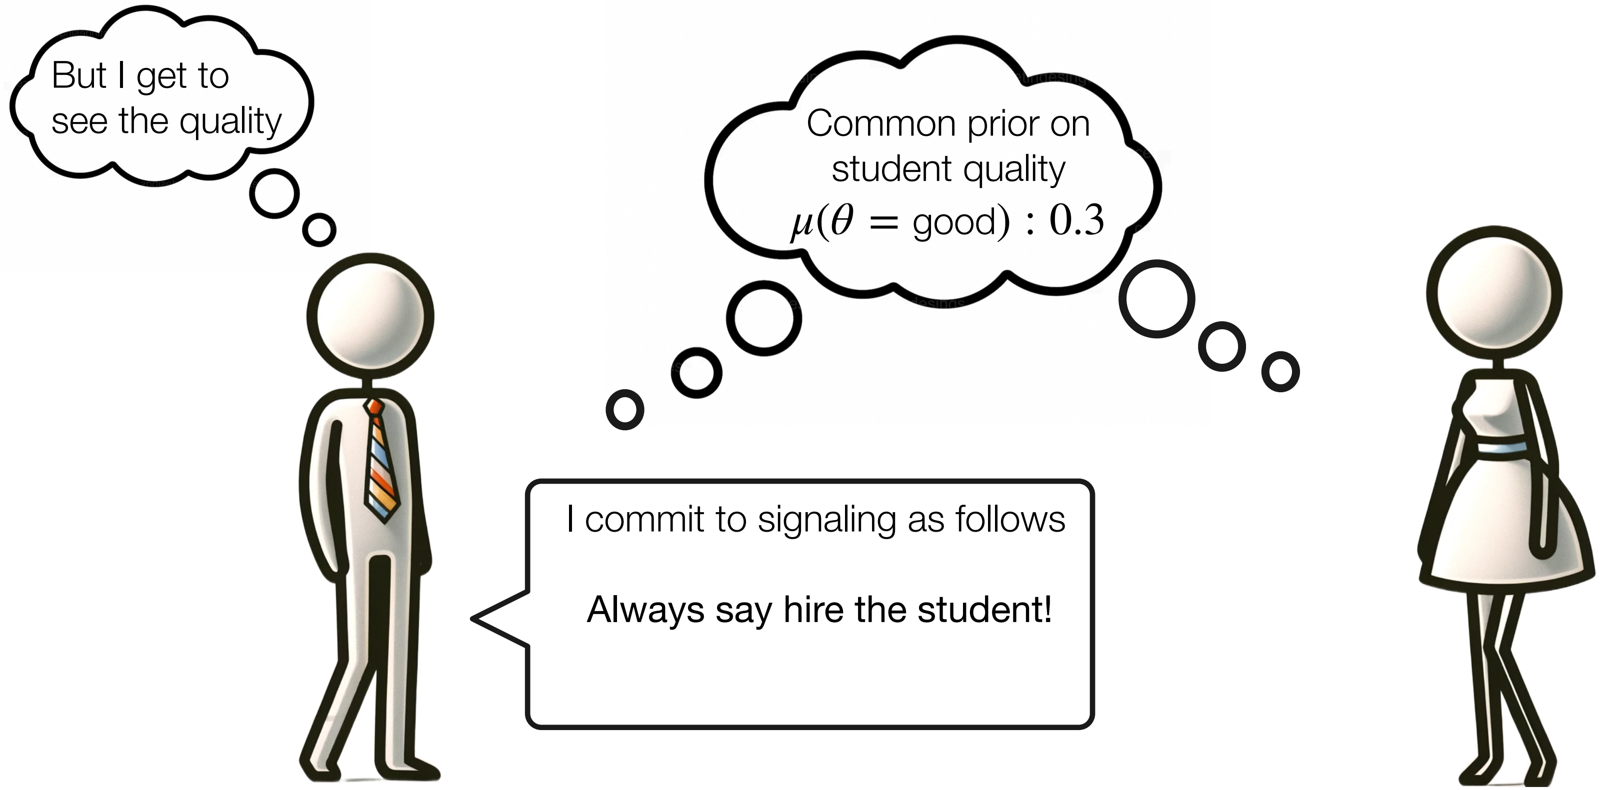

Sender - Professor

$u(a, \theta)$: +1 if student hired

Receiver - Hiring Manager

$w(a, \theta)$: +1 if right decision is made

But I get to see the quality

Common prior on student quality
$\mu(\theta = \text{good}) : 0.3$

That's totally uninformative!

$\mathbb{E}[u(a^*, \theta)] = 0$

I commit to signaling as follows

Always say hire the student!

$a^* = \text{not hire}$

$\mathbb{E}[w(a^*, \theta)] = 0.7$

Sender - Professor

$u(a, \theta)$: +1 if student hired

Receiver - Hiring Manager

$w(a, \theta)$: +1 if right decision is made

But I get to see the quality

Common prior on student quality
$\mu(\theta = \text{good}) : 0.3$

That's fully informative!

$\mathbb{E}[u(a^*, \theta)] = 0.3$

I commit to signaling as follows

$\theta = \text{good} \implies \text{say hire}$
$\theta = \text{bad} \implies \text{don't hire}$

$a* =$ hire if s=hire
$a* =$! hire if s= ! hire
$\mathbb{E}[w(a^*, \theta)] = 1.0$

Sender - Professor
$u(a, \theta)$: +1 if student hired

Receiver - Hiring Manager
$w(a, \theta)$: +1 if right decision is made

But I get to see the quality

Common prior on student quality $\mu(\theta = \text{good}) : 0.3$

I commit to signaling as follows

$\theta = \text{good} \implies$ say hire

$\theta = \text{bad} \implies$ say hire 42% of the time

Sender - Professor

$u(a, \theta)$: +1 if student hired

Receiver - Hiring Manager

$w(a, \theta)$: +1 if right decision is made

But I get to see the quality

Common prior on student quality

$\mu(\theta = \text{good}) : 0.3$

I commit to signaling as follows

$\pi(s = \text{hire} \,|\, \theta = \text{good}) = 1$

$\pi(s = \text{hire} \,|\, \theta = \text{bad}) = 0.42$

Sender - Professor

$u(a, \theta)$: +1 if student hired

Receiver - Hiring Manager

$w(a, \theta)$: +1 if right decision is made

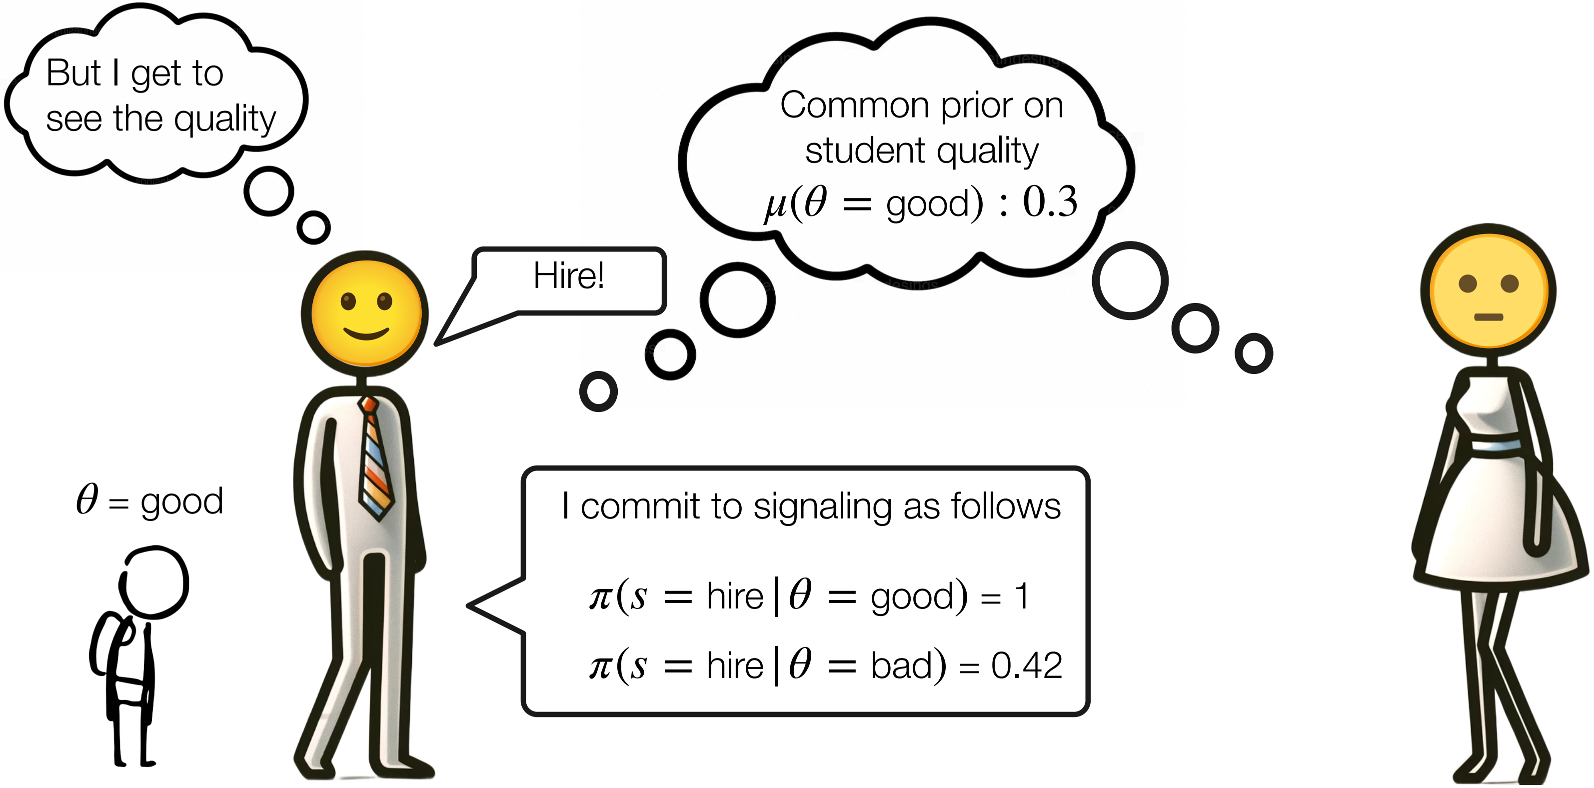

Sender - Professor
$u(a, \theta)$: +1 if student hired

Receiver - Hiring Manager
$w(a, \theta)$: +1 if right decision is made

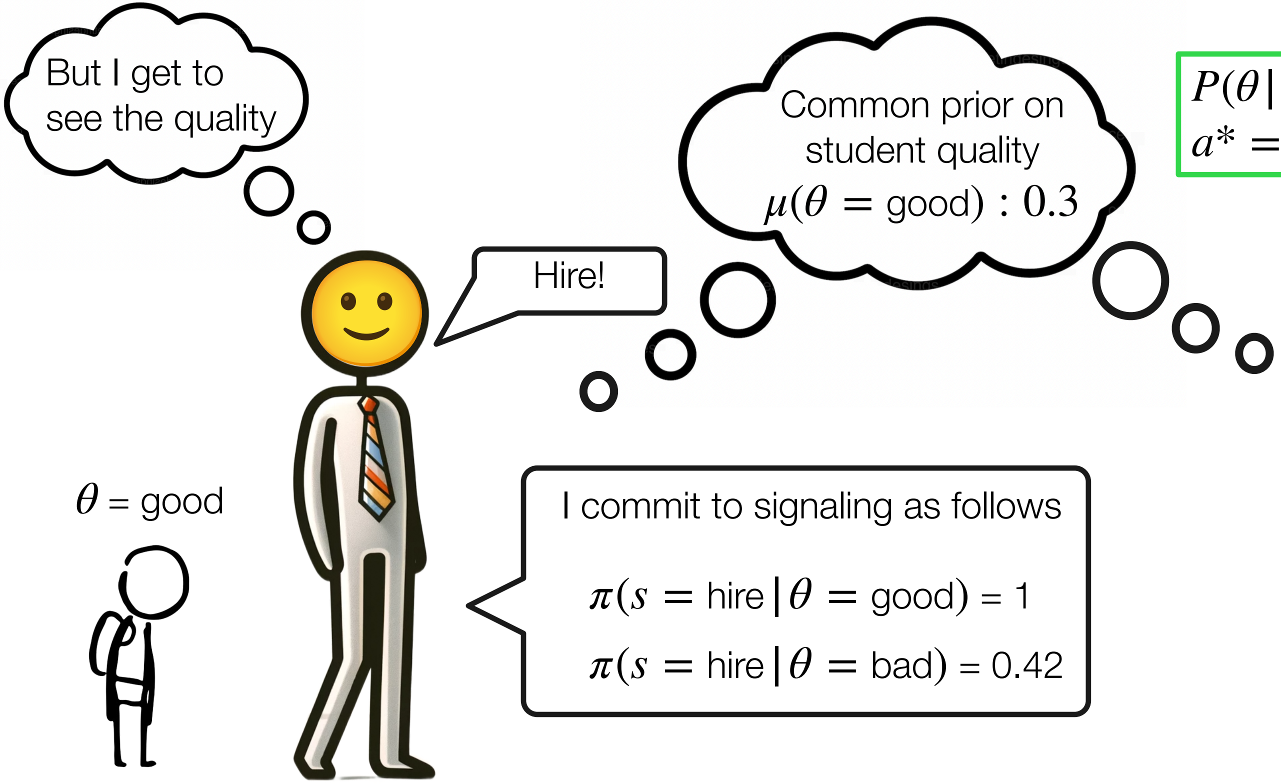

$$P(\theta \mid s = \text{hire}) \propto \pi(s = \text{hire} \mid \theta)\mu(\theta)$$
$$a^* = \text{argmax}_a w(a, \theta) P(\theta \mid s = \text{hire})$$

Sender - Professor

$u(a, \theta)$: +1 if student hired

Receiver - Hiring Manager

$w(a, \theta)$: +1 if right decision is made

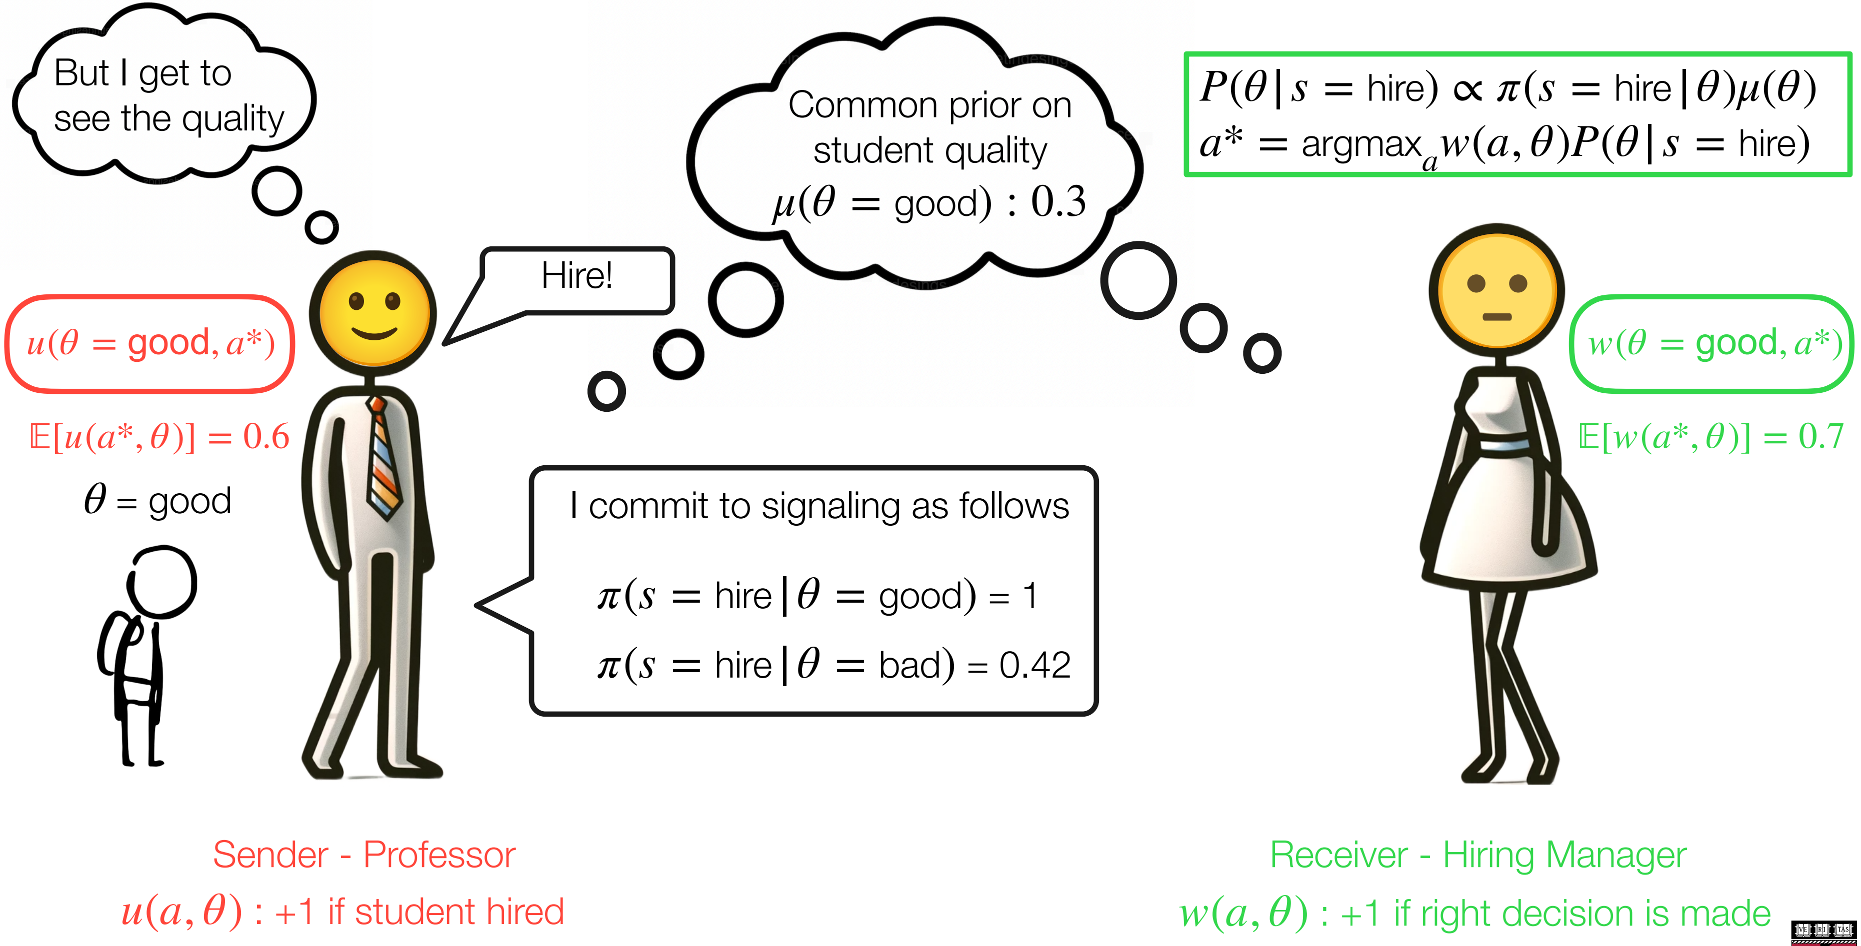
But I get to see the quality
Common prior on student quality $\mu(\theta = \text{good}) : 0.3$
$P(\theta | s = \text{hire}) \propto \pi(s = \text{hire} | \theta)\mu(\theta)$
$a^* = \text{argmax}_a w(a, \theta)P(\theta | s = \text{hire})$
Hire!
$u(\theta = \text{good}, a^*)$
$\mathbb{E}[u(a^*, \theta)] = 0.6$
θ = good
$w(\theta = \text{good}, a^*)$
$\mathbb{E}[w(a^*, \theta)] = 0.7$
I commit to signaling as follows
$\pi(s = \text{hire} | \theta = \text{good}) = 1$
$\pi(s = \text{hire} | \theta = \text{bad}) = 0.42$
Sender - Professor
$u(a, \theta)$: +1 if student hired
Receiver - Hiring Manager
$w(a, \theta)$: +1 if right decision is made

1. Bayesian Persuasion - Details

- Sender must commit to signaling scheme before realization

- When sender is designing/choosing signaling scheme they have no more information than receiver.
  - Chooses a scheme to maximize expected ex-ante utility

- In the standard setting, under mild assumptions optimal scheme can be solved using a linear program [2].

2. Motivations

- In many settings involving persuasion, there maybe multiple informationally advantageous senders looking to persuade a single agent.

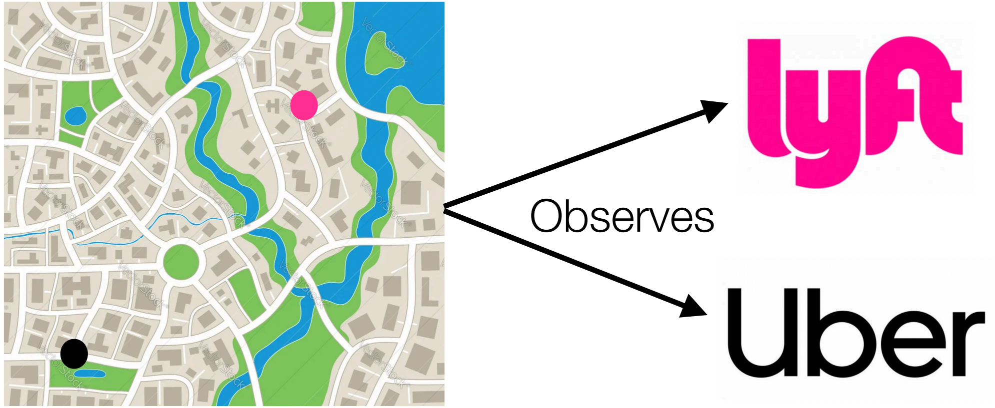

Real time traffic
and demand data

2. Motivations

- In many settings involving persuasion, there maybe multiple informationally advantageous senders looking to persuade a single agent.

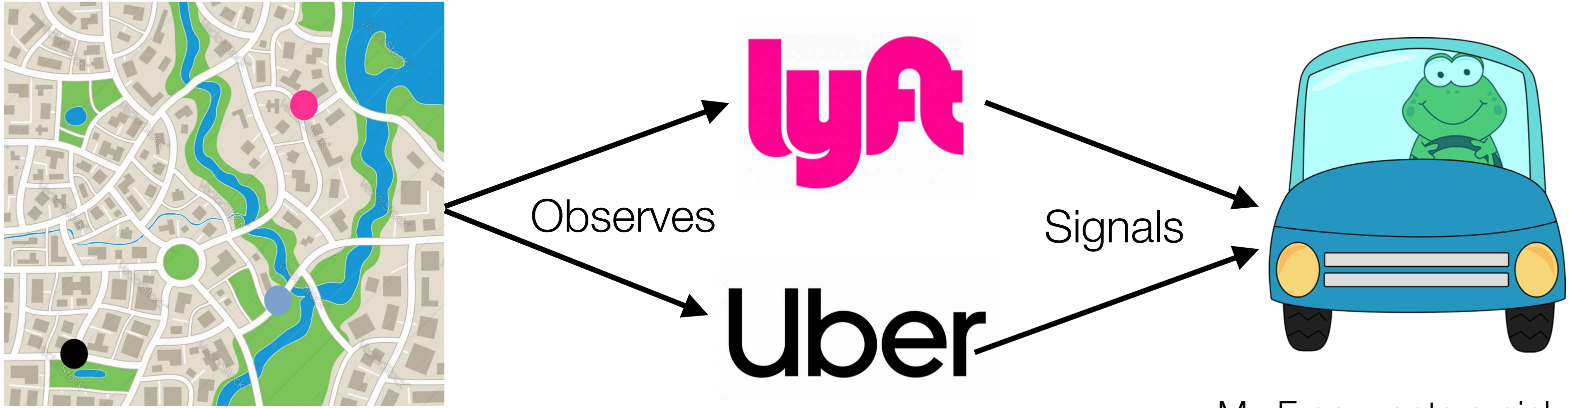

Real time traffic
and demand data

Mr. Frog wants a pick-up
near the marsh

2. Multi-Sender Persuasion Model

- n senders with utility $u_i(\theta, a)$, and 1 receiver with utility $w(\theta, a)$
- All senders and receiver have a common prior $\mu(\theta)$
- Senders jointly observe a realization $\theta \sim \mu$,
- Sender simultaneously commit to signaling $\pi_i(s \mid \theta)$; receiver observes $\bar{s} \sim \prod \pi_i$
    - Receiver updates belief based on joint signal and takes their optimal action.

Multi-leader single follower Stackelberg game.

Simultaneous interaction between senders; sequential interaction with receiver

2. Multi-Sender Persuasion Model

- Senders jointly observe a realization $\theta \sim \mu$,
- Sender simultaneously commit to signaling $\pi_i(s \mid \theta)$; receiver observes $\bar{s} \sim \prod \pi_i$
  - Receiver updates belief based on joint signal and takes their optimal action.
  - Sender utility: $u_i(\pi_i, \pi_{-i}) = \sum_{\theta} \sum_{\bar{s} \in S^n} \bar{\pi}(\bar{s} \mid \theta) u_i(\theta, a^*)$

Nash equilibrium between sender's signaling strategy given receiver best responds.

$$\forall i \, \forall \pi_i \; u_i(\pi_i^{eq}, \bar{\pi}_{-i}^{eq}) \geq u_i(\pi_i, \bar{\pi}_{-i}^{eq})$$

2. Related Works

[7] studies multi-sender setting wherein senders can arbitrarily correlate their signals by conditioning on others realizations - leads to some simplifications.
[8] studies multi-sender persuasion where senders commit to schemes sequentially. Different solution concept: sub-game perfect equilibrium.
[9] Study a simultaneous and independent model, but with zero sum sender utilities.

- We consider senders committing to simultaneous and independent schemes, but with no restriction on utility or state/action/signal structure.
- Prior works provide little computational insights on the problem.

2. Equilibrium Characterization

- [6] notes that under two conditions an equilibrium exists that fully reveals the realized state to the receiver, allowing them to achieve maximal utility:
  - (1) there is a unique optimal action for the receiver at each state.
  - (2) The signal space $|S| = |\Theta|$, the state space.

- (Equilibrium): Any sender chooses to deterministically map state to signal.

2. Equilibrium Characterization

- [6] notes that under two conditions a full information equilibrium exists:
  - (1) there is a unique optimal action for the receiver at each state.
  - (2) The signal space $|S| = |\Theta|$, the state space.

- We prove that condition 2 can be significantly relaxed.

- Theorem: A full information equilibrium exists if:
  - (1) there is a unique optimal action for the receiver at each state.
  - (2') The signal space $|S| \geq min(|\Theta|^{1/n-1}, |A|^{1/n-1})$

2. Equilibrium Characterization

- Theorem: A full information equilibrium exists if:
  - (1) there is a unique optimal action for the receiver at each state.
  - (2') The signal space $|S| \geq min(|\Theta|^{1/n-1}, |A|^{1/n-1})$

- Map each state to a deterministic set of signals such that:
  - each state is uniquely identifiable by $n-1$ signals.
  - Leverages an idea from error correcting codes.

- If $|\Theta| > |A|$, can interpret full information as revealing optimal action at each state.

2. Equilibrium Characterization

- Theorem: A full information equilibrium exists if:
  - (1) there is a unique optimal action for the receiver at each state.
  - (2') The signal space $|S| \geq min(|\Theta|^{1/n-1}, |A|^{1/n-1})$

- We have relaxed the 2nd conditions. But how about the first?

- Theorem: If condition (1) does not hold, then under some tie-breaking rule, computing the Nash equilibrium is PPAD-Hard.
  - Reduction from finding equilibrium in two-player games with binary utilities.

2. Best Response

- In single sender persuasion, optimal signaling scheme (i.e. sender's best strategy) can be expressed as a linear program.

- In multi-sender persuasion, a sender i's best response is their optimal signaling scheme for a fixed set of scheme of other senders, $\bar{\pi}_{-i}$.

Theorem: Even for 2 senders, computing a single sender's best response is NP-Hard. Further, it is NP-Hard to even approximate the best response.

2. Best Response Hardness

- An involved reduction from the hardness of public persuasion with k receivers.

Public Persuasion

- 1 sender k receivers; public scheme
- Common prior $\mu(\theta)$ with $|\Theta|$ states.
- Binary action receivers with utility $w_j(\theta, a)$
- Sender utility depends on receiver $u_j(\theta, a)$

Equivalent Multi-Sender Persuasion

- 2 senders and 1 receiver
- $\Theta' = \Theta \cup \{\bar{\theta}_1, \ldots \bar{\theta}_k\}$ states
- The receiver action space $A = A_+ \cup A_- \cup a_\infty$ where $|A_+| = |A_-| = k$

2. Best Response Hardness

- $\Theta' = \Theta \cup \{\bar{\theta}_1, \ldots \bar{\theta}_k\}$ states:
  - When $\theta \in \Theta$ happens, single receiver utility $w(a_j^+, \theta) = w_j(a^+, \theta)$
  - When $\bar{\theta}_j$ happens, very bad for receiver to take anything except $\{a_j^+, a_j^-\}$

- Non-best responding sender's signaling scheme uses k signals such that:
  - Realization s_j implies receiver will take action $\{a_j^+, a_j^-\}$

$\forall\, k$ possible signal realization of non-best responding sender, the single receiver's plausible actions mimic that of the k^{th} receiver in public persuasion

2. Deep Learning Approach

Given the difficulty of best-response and non trivial equilibrium in this setting, we propose deep learning methods for finding local equilibria.

- Challenging due to the non-convex discontinuous utility landscape

2. Deep Learning Approach

- Established the difficulty of computing equilibrium in the general setting, even with access to complete information.

- Local Equilibrium: No profitable deviation in a local neighborhood. Popular notion in deep learning for game.

- Can we use deep learning methods to find local equilibrium of this game with only sample access to utility?
  - Key Challenge: Joint utility landscape is non-convex and discontinuous.

2. Experimental Results

Propose novel architecture to approximate this utility with sample access.

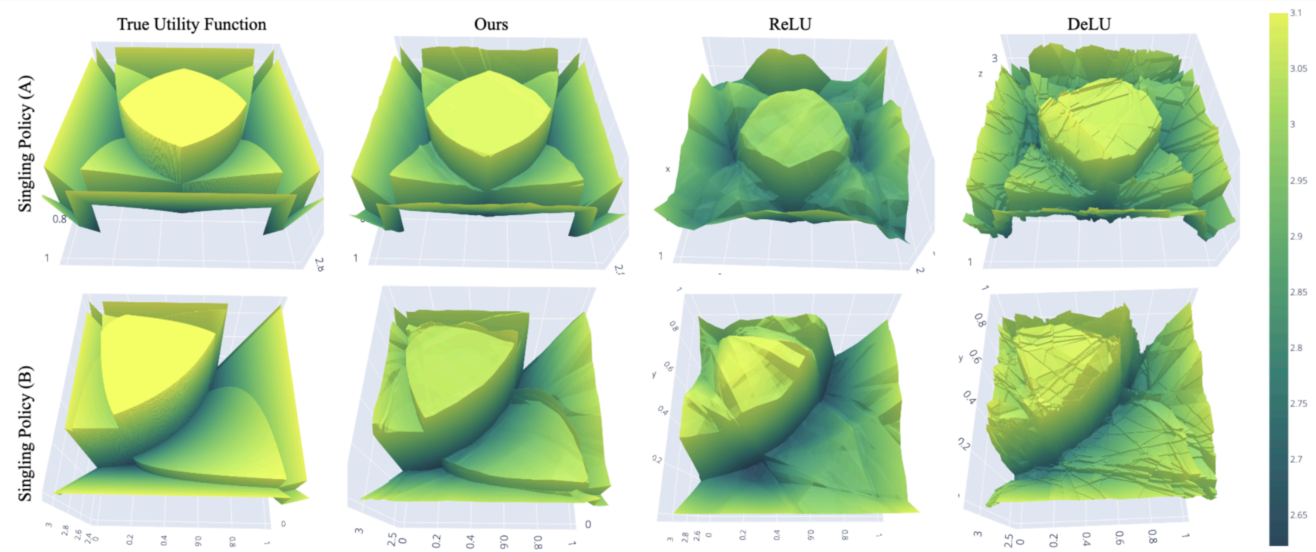

2. Experimental Results

When paired with extra-gradient algorithm, we find local equilibria that outperforms other methods as well as full revelation equilibrium.

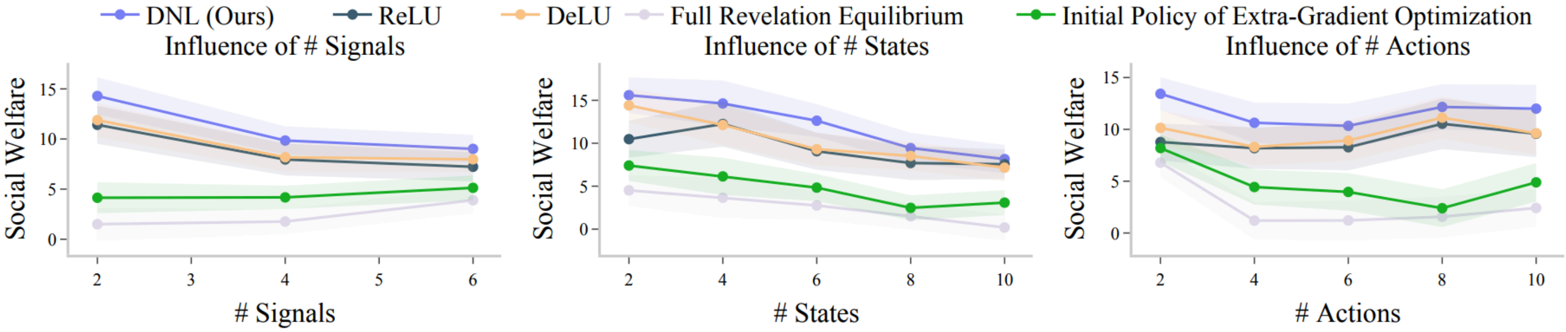

Figure 5: Our method achieves higher social welfare compared against baselines and full-revelation solutions in games with 4 senders.

Thank you!

References:

[7]: Matthew Gentzkow, Emir Kamenica. **Bayesian persuasion with multiple senders and rich signal spaces.** Games and Economic Behavior. 2017

[8]: Fei Li, Peter Norman. **Sequential persuasion.** Theoretical Economics. 2021

[9]: Dilip Ravindran, Zhihan Cui. **Competing Persuaders in Zero-Sum Games.** Preprint. 2020