Non-Bayesian Information Design:

Learning and LLM-Based Approaches

Tao Lin

Microsoft Research (2025) → CUHK-Shenzhen (2026)

2026/1/9, Huawei

Information Design

An economic model about ***information asymmetry***: *one player ("sender") strategically reveals information to influence the decision of another player ("receiver").*

Examples:

- Advertising
  - Seller reveals product information to buyers
- School designs letter grading scheme
- Professor writing recommendation letter
- ...

Information Design is a form of "Persuasion"

One Quarter of GDP Is Persuasion

By Donald McCloskey and Arjo Klamer*

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

Home › Publications › Economic Roundup Issue 1, 2013 › Persuasion is now 30 per cent of US GDP

Persuasion is now 30 per cent of US GDP

Gerry Antioch[1]

Date: 06 June 2013

Classical Information Design Models

- Many classical models for information design:
  - "Information Disclosure Games" (Grossman, 1981; Milgrom 1981)
  - "Cheap Talk" (Crawford & Sobel, 1982)
  - "Bayesian Persuasion" (Kamenica & Gentzkow, 2011)
  - …
- **Common modeling approach:**
  - **Abstract signal space:** The information transmitted from sender to receiver is modeled by a random variable s correlated with the state of the world ω
  - **Bayesian receiver:** The receiver does Bayes update after receiving s
- Importantly, *how the signal s is communicated (e.g., wording)* doesn't matter.
- Our work: *non-Bayesian* information design, via "learning + LLM" approaches.

Outline

- Background on a Classical Information Design Model: "Bayesian Persuasion" [Kamenica & Gentzkow, 2011]
- Information Design with a Learning Receiver
- Information Design with Large Language Models

Example of **Bayesian Persuasion**: Recommendation Letter

Sender

Prior μ

$p = 0.2$

0.8

State of the world ω

suitable

not suitable

Signal s

"suitable"

"not suitable"

?

Receiver

Utility:

$u(\text{hire}, \omega) = 1$

$u(\text{not hire}, \omega) = 0$

Q: What's the professor's optimal recommendation strategy?

Utility:

$v(\text{hire}, \mathit{suitable}) = 1$

$v(\text{hire}, \mathit{not\ suitable}) = -1$

$v(\text{not hire}, \omega) = 0$

Bayesian reasoning:

Hire if $\Pr(\mathit{suitable} \mid s) \geq 0.5$

Example of **Bayesian Persuasion**: Recommendation Letter

Sender

Prior μ

State of the world ω

Signal s

Receiver

$p = 0.2$ → *suitable* → *"suitable"*

0.8 → *not suitable* → *"suitable"*

"not suitable"

Utility:

$u(\text{hire}, \omega) = 1$

$u(\text{not hire}, \omega) = 0$

Q: What's the professor's optimal recommendation strategy?

- **Always say "suitable"**: recruiter does not hire; professor gets 0

Utility:

$v(\text{hire}, \mathit{suitable}) = 1$

$v(\text{hire}, \mathit{not\ suitable}) = -1$

$v(\text{not hire}, \omega) = 0$

Bayesian reasoning:

Hire if $\Pr(\mathit{suitable} \mid s) \geq 0.5$

Example of **Bayesian Persuasion**: Recommendation Letter

Sender

Prior μ

State of the world ω

Signal s

Receiver

$p = 0.2$ *suitable* $\xrightarrow{0.2}$ *"suitable"*

0.8 *not suitable* $\xrightarrow{0.8}$ *"not suitable"*

Utility:

$u(\text{hire}, \omega) = 1$

$u(\text{not hire}, \omega) = 0$

Q: What's the professor's optimal recommendation strategy?

- **Always say "suitable"**: recruiter does not hire; professor gets 0

- **Hoest Recommendation:** professor gets 0.2

Utility:

$v(\text{hire}, \mathit{suitable}) = 1$

$v(\text{hire}, \mathit{not\ suitable}) = -1$

$v(\text{not hire}, \omega) = 0$

Bayesian reasoning:

Hire if $\Pr(\mathit{suitable} \mid s) \geq 0.5$

Example of **Bayesian Persuasion**: Recommendation Letter

Sender

Prior μ

State of the world ω

Signal s

Receiver

$p = 0.2$ → *suitable* $\xrightarrow{0.2}$ *"suitable"*

0.8 → *not suitable* $\xrightarrow{0.2}$ *"suitable"*

not suitable $\xrightarrow{0.6}$ *"not suitable"*

Utility:

$u(\text{hire}, \omega) = 1$

$u(\text{not hire}, \omega) = 0$

Q: What's the professor's optimal recommendation strategy?

- **The optimal strategy** (partial info revelation):
  - *if suitable, say "suitable";*
  - *if not, say "suitable" w.p.* 25%

  Professor gets 0.4

Utility:

$v(\text{hire}, \textit{suitable}) = 1$

$v(\text{hire}, \textit{not suitable}) = -1$

$v(\text{not hire}, \omega) = 0$

Bayesian reasoning:

Hire if $\Pr(\textit{suitable} \mid s) \geq 0.5$

Key Assumptions in Classical BP Theory

Learning

- **Commitment:**
  - Sender can commit to a randomized mapping ("signaling scheme") $\pi: \Omega \to \Delta(S)$ before state realization.
- **Bayesian receiver:**
  - The receiver knows the prior μ and the sender's signaling scheme π, and does Bayes update after receiving signal s (and best responds)

- **Abstract signal space:**
  - Language doesn't matter – only the correlation between signal and state matters.

LLM

Outline

- Background on a Classical Information Design Model: “Bayesian Persuasion” [Kamenica & Gentzkow, 2011]
- Information Design with a Learning Receiver
- Information Design with Large Language Models

Generalized Principal-Agent Problem with a Learning Agent

Tao Lin

Yiling Chen

Harvard University

ICLR (International Conference on Learning Representations), 2025

Quantitative Economics, 2026

Learning in games has a long history

- Adaptive Dynamics & Fictitious Play: Brown (1951), Robinson (1951), Shapley (1953)
- The Theory of Learning in Games: Fudenberg & Levine (1991)
- No-regret learning and correlated equilibrium:
    - Hart & Mas-Colell (2000); Blum & Mansour (2007)
- Prediction, Learning, and Games: Cesa-Bianchi & Lugosi (2006)
- ……

Our work:

- *Replaces the Bayesian receiver with a learning receiver in information design problems*
- *Studies whether the learning outcome matches the classical outcome.*

Receiver's Learning Problem: *Contextual Multi-Armed Bandit*

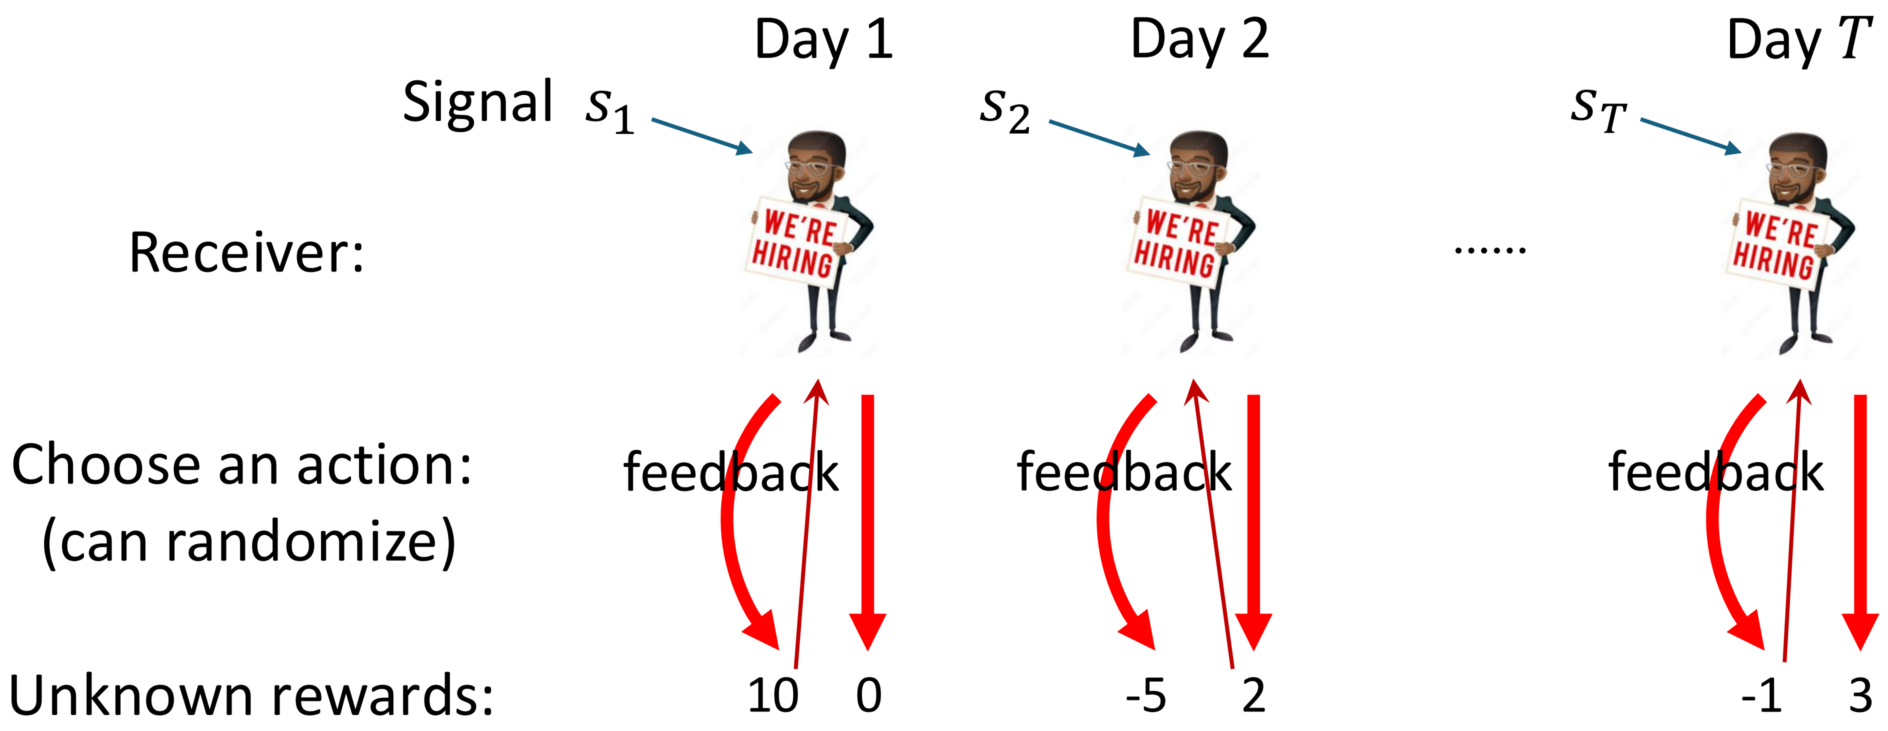

(Contextual) No-Regret Property

For any sequence of unknown rewards, after T rounds,

$\mathbb{E}$[**Total reward obtained**] $\geq$ **Total reward of the best signal-to-action mapping** $- O(\sqrt{T})$

No-regret learning algorithms exist; most are based on "smoothed best response to history".

Information Design with a Learning Receiver

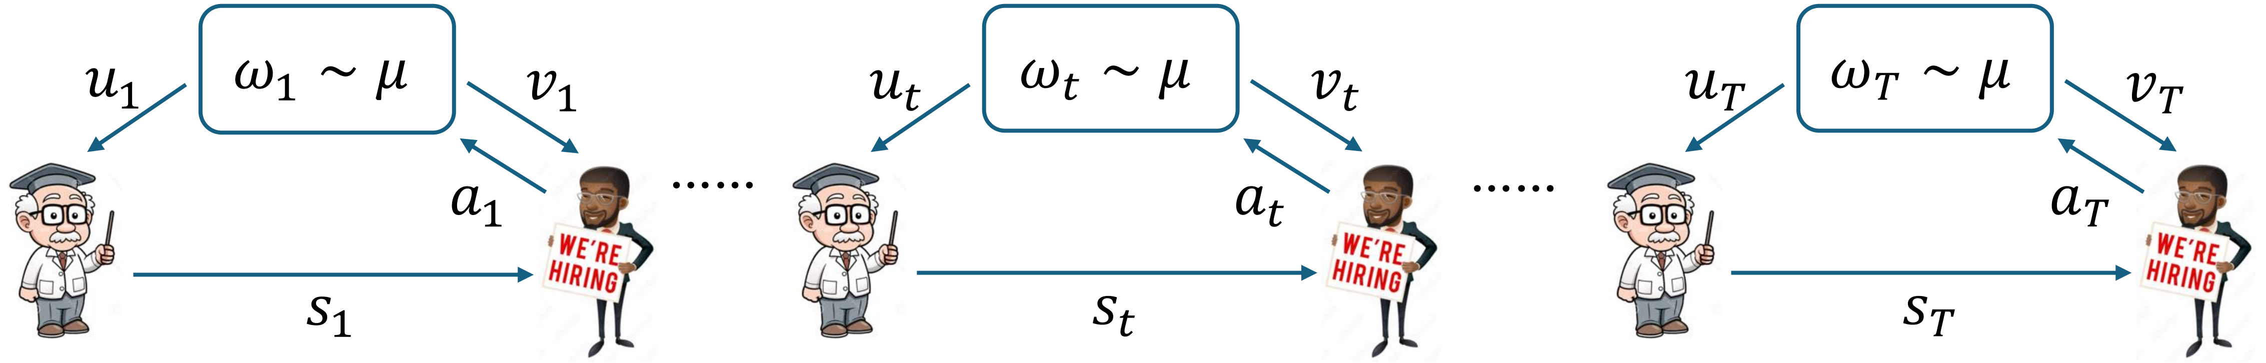

- Two players: sender and receiver
- Sender knows the state distribution $\mu \in \Delta(\Omega)$, *which the receiver doesn't need to know*
- At each round t:
    - The receiver uses a **Contextual MAB** algorithm to decide, for each possible signal, what action to choose: $\rho_t : S \to \Delta(A)$ *(based on history)*
    - State $\omega_t \sim \mu$ is realized
    - Sender sends signal $s_t \sim \pi_t(\cdot|\omega_t)$
    - Receiver takes action $a_t \sim \rho_t(\cdot|s_t)$
    - The two players obtain utilities $u(a_t, \omega_t)$, $v(a_t, \omega_t)$

No commitment: "Bayesian Persuasion" = "Cheap Talk" (Crawford & Sobel, 1982)

Information Design with a Learning Receiver

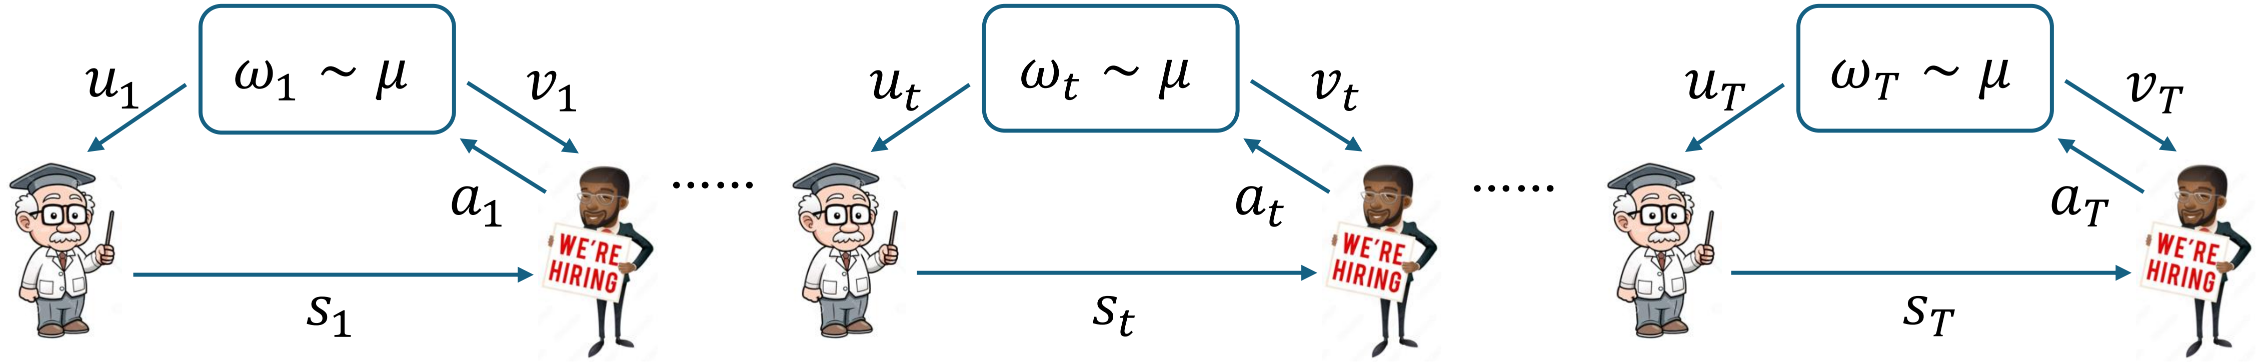

Our Questions:

With a learning receiver,

- Can the sender still achieve the classical outcome (with commitment and Bayesian receiver)?

$$U_{\text{sender}}(\text{learning receiver}) \geq U^*_{\text{sender}}(\text{Bayesian receiver})$$

- Can the sender ***do better than*** the classical outcome?

$$U_{\text{sender}}(\text{learning receiver}) > U^*_{\text{sender}}(\text{Bayesian receiver})$$

Main Contributions

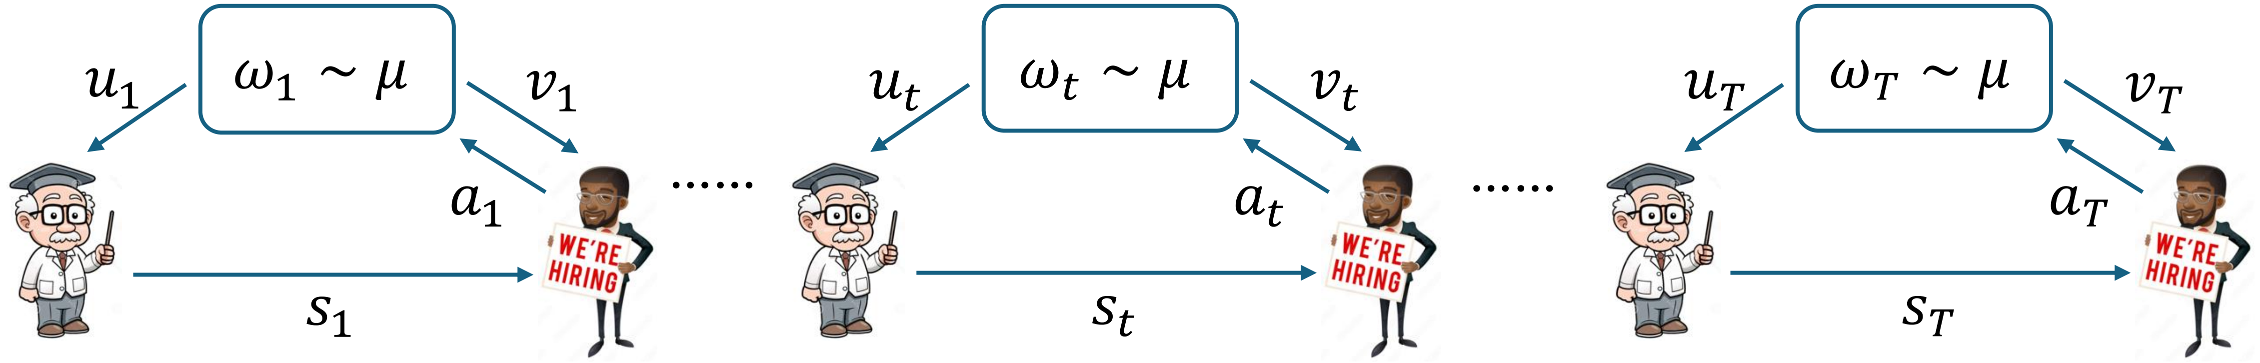

Result 1:

With a learning receiver,

- The sender can achieve the classical outcome:
  - $U_{\text{sender}}(\text{learning receiver}) \geq U^*_{\text{sender}}(\text{Bayesian receiver}) - O\left(\sqrt{\text{Reg}(T)}\right)$
  - **How?** Just use the optimal signaling scheme π^* in the classical setting. The receiver will learn to best respond as $T \to \infty$
  - **Why $O\left(\sqrt{\text{Reg}(T)}\right)$?** The receiver may take $\sqrt{\text{Reg}(T)}$-sub-optimal action in $\sqrt{\text{Reg}(T)}$ fraction of time, causing a total loss of $\sqrt{\text{Reg}(T)}$ to the sender.

Main Contributions

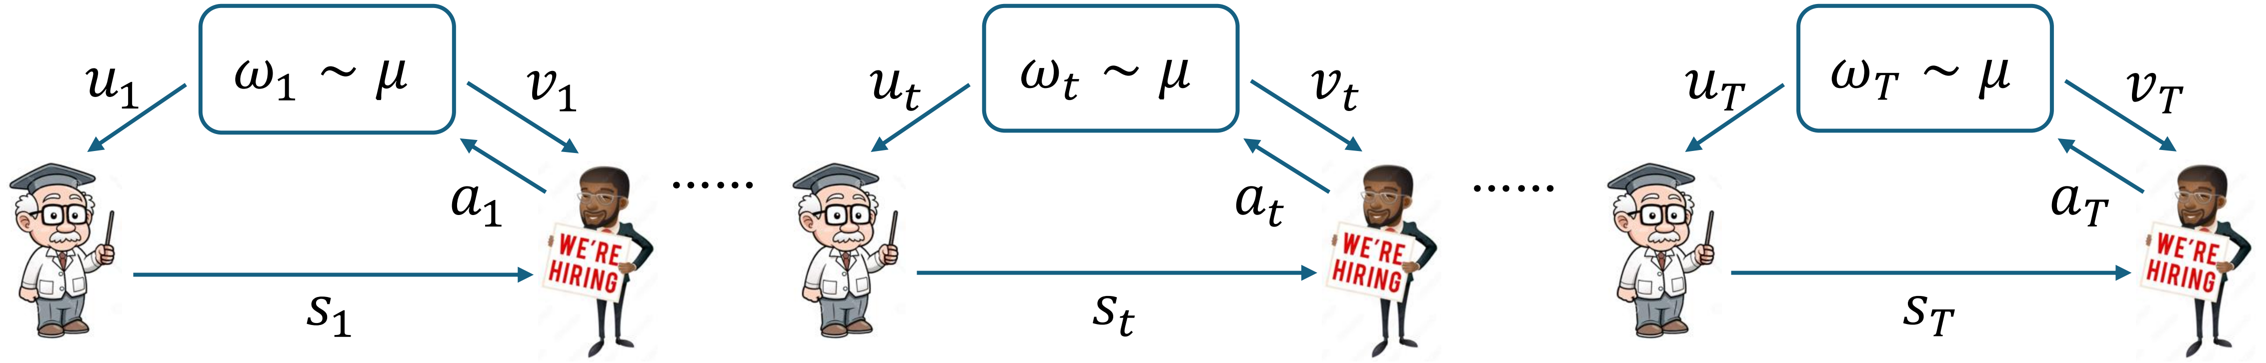

Result 2:

With a learning receiver,

- The sender can achieve the classical outcome.
- Can the sender ***do better than*** the classical outcome?
  - **Yes**, for all "smoothly-best-responding" no-regret learning receivers: $\exists$ instance,

    $$U_{\text{sender}}(\text{learning receiver}) > U^*_{\text{sender}}(\text{Bayesian receiver}) + \text{Const}$$

  - **No**, for all "no-swap-regret" learning receivers.

Intuition for why ***doing better*** is possible: Dynamic Strategy

First,
honest recommendation:

suitable (0.2) → suitable →

not s. (0.8) → not s. →

→ hire

→ not hire

gets 0.2

Then, always
say "suitable":

suitable (0.2) → suitable →

not s. (0.8) → suitable →

→ hire

gets 1

After some time, the
receiver will realize that
the signal is not truthful:

suitable (0.2) → suitable →

not s. (0.8) → suitable →

→ not hire

gets 0

Average utility > 0.4 ?

No-Swap-Regret Learning Algorithms

No-Regret

For any sequence of reward functions, after T rounds,

$$\mathbb{E}[\sum_{t=1}^{T} v_t(a_t)] \geq \max_{a \in A} \mathbb{E}[\sum_{t=1}^{T} v_t(a)] - O(\sqrt{T}).$$

Many no-regret MAB algorithms do "smoothed best response to history".

No-Swap-Regret

For any sequence of reward functions, after T rounds,

$$\mathbb{E}[\sum_{t=1}^{T} v_t(a_t)] \geq \max_{\phi: A \to A} \mathbb{E}\left[\sum_{t=1}^{T} v_t\big(\phi(a_t)\big)\right] - O(\sqrt{T}).$$

No-swap-regret MAB algorithms exist: [Hart & Mas-Colell, 2000] [Blum & Mansour, 2007]

Why can't the sender exploit a no-swap-regret learning receiver?

- Consider the signal-action pair (s_t, a_t) as a *joint signal* from some signaling scheme $\tilde{\pi}$.
- No-swap-regret guarantees approximate best response to $\tilde{\pi}$.

Our & Previous Work on Learning in Principal-Agent Games

- "Smoothly-best-responding" no-regret learning agents are exploitable in many games [1] [2]
- If the agent does **no-swap-regret** learning, then the principal
  - *cannot exploit* the agent in the games in [1] [2]: $U(\text{learning}) < U^*(\text{rational}) + o(1)$
  - *can exploit* the agent in some other games [3]: $U(\text{learning}) > U^*(\text{rational}) + const$

No-Swap-Regret is Not Exploitable

Games with **complete information**

[1] *Bimatrix Stackelberg Games*

[2] *Contract Design*

……

?

No-Swap-Regret is Exploitable

Games where the **agent** has **private information**

[3] *Auctions*

……

[1] Deng, Schneider, Sivan (2019). Strategizing against No-regret Learners.

[2] Guruganesh, Kolumbus, Schneider, Talgam-Cohen, Vlatakis-Gkaragkounis, Wang, Weinberg (2024). Contracting with a Learning Agent.

[3] Braverman, Mao, Schneider, Weinberg (2018). Selling to a No-Regret Buyer.

Our & Previous Work on Learning in Principal-Agent Games

- "Smoothly-best-responding" no-regret learning agents are exploitable in many games [1] [2]
- If the agent does **no-swap-regret** learning, then the principal
  - *cannot exploit* the agent in the games in [1] [2]: $U(\text{learning}) < U^*(\text{rational}) + o(1)$
  - *can exploit* the agent in some other games [3]: $U(\text{learning}) > U^*(\text{rational}) + const$

No-Swap-Regret is Not Exploitable

Games with **complete information**

[1] *Bimatrix Stackelberg Games*
[2] *Contract Design*
……

[Our work]

Games where the **principal** has **private information**, while the agent does not:

e.g., Information Design

No-Swap-Regret is Exploitable

Games where the **agent** has **private information**

[3] *Auctions*
……

[1] Deng, Schneider, Sivan (2019). Strategizing against No-regret Learners.
[2] Guruganesh, Kolumbus, Schneider, Talgam-Cohen, Vlatakis-Gkaragkounis, Wang, Weinberg (2024). Contracting with a Learning Agent.
[3] Braverman, Mao, Schneider, Weinberg (2018). Selling to a No-Regret Buyer.

Outline

- Background on a Classical Information Design Model: "Bayesian Persuasion" [Kamenica & Gentzkow, 2011]
- Information Design with a Learning Receiver
- Information Design with Large Language Models

Information Design with Large Language Models

Paul Dütting
Google Research

Safwan Hossain[†]
Harvard University

Tao Lin
Harvard University

Renato Paes Leme
Google Research

Sai Srivatsa Ravindranath
Harvard University

Haifeng Xu
University of Chicago

Song Zuo
Google Research

Working paper (arXiv 2025)

Key Assumptions in Classical BP Theory

Learning

- **Commitment**:
  - Sender can commit to a randomized mapping ("signaling scheme") $\pi: \Omega \to \Delta(S)$ before state realization.
- **Bayesian receiver:**
  - Knowing the prior μ and signaling scheme π, the receiver does Bayes update after receiving signal s (and then best responds)

- **Abstract signal space:**
  - Language doesn't matter – only the correlation between signal and state matters.

We aim to capture the linguistic aspect of persuasion

Example 1: Framing Effect (Tversky & Kahneman, 1981)

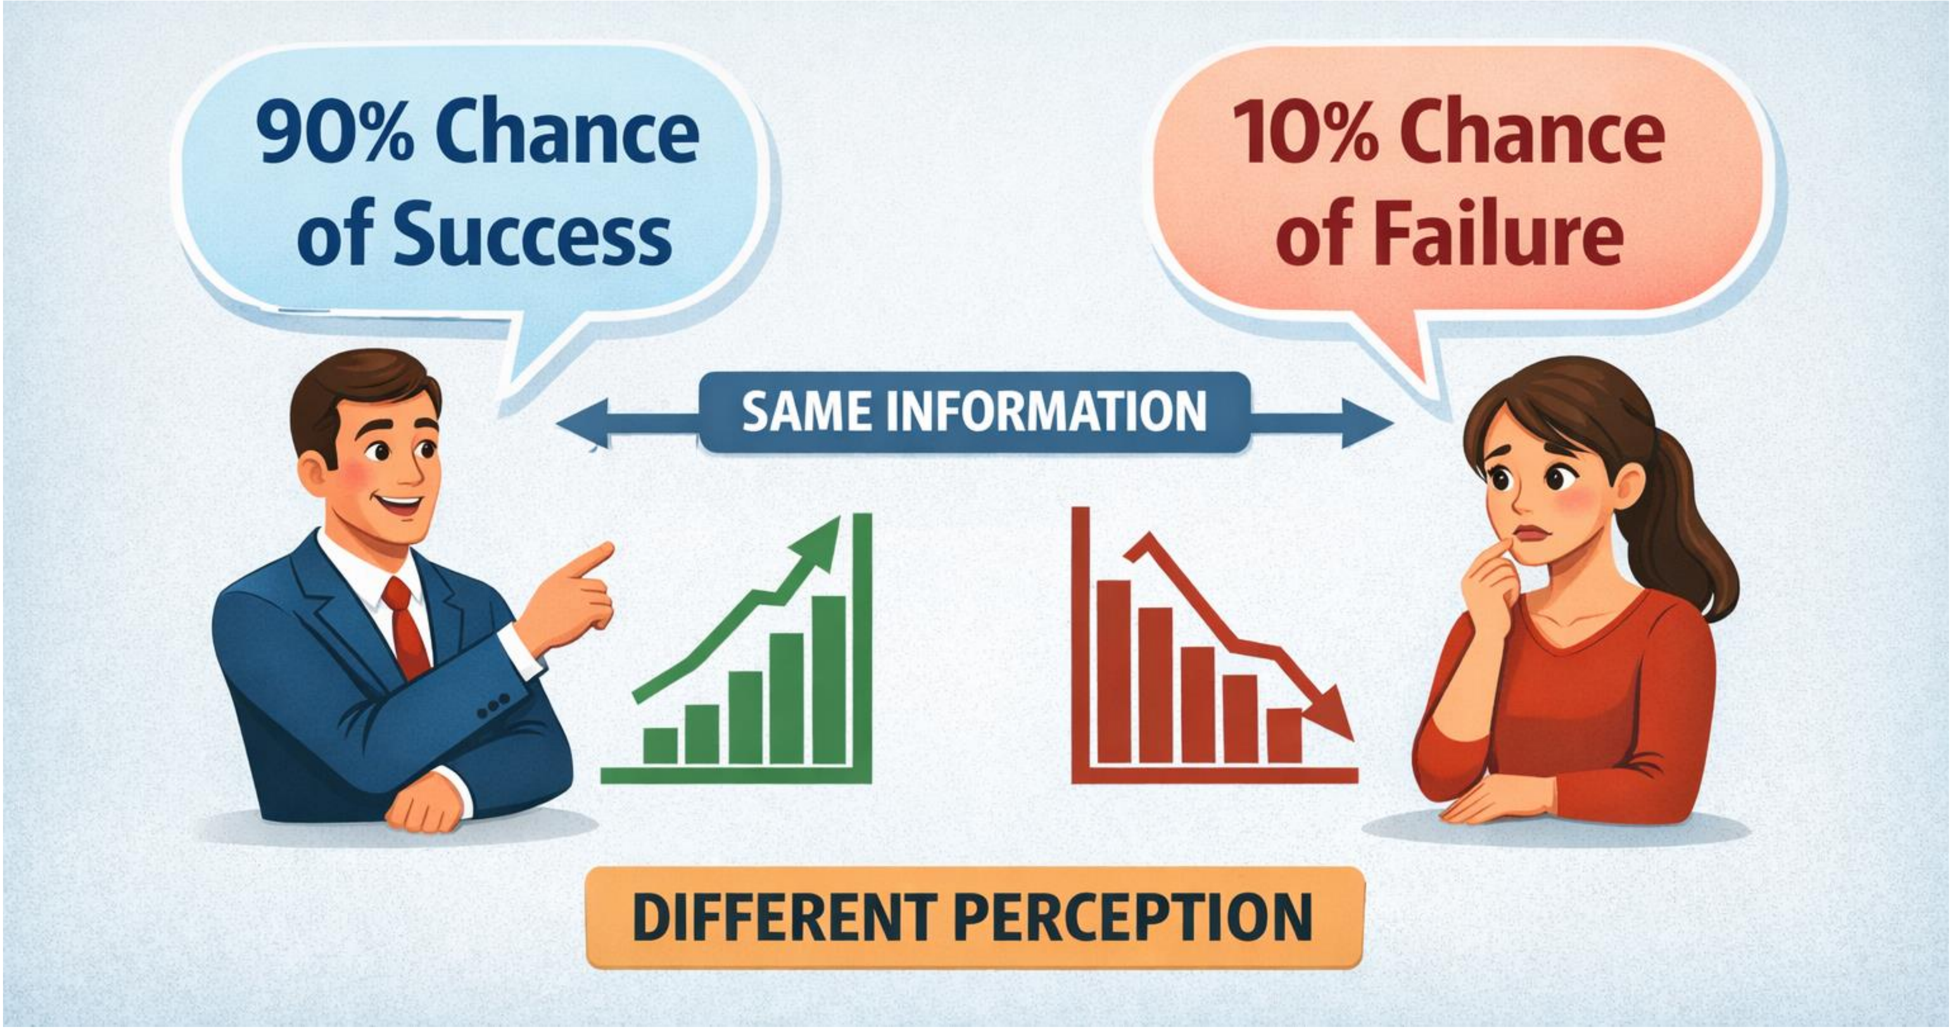

Example 2: Slogan/Logo of a Brand

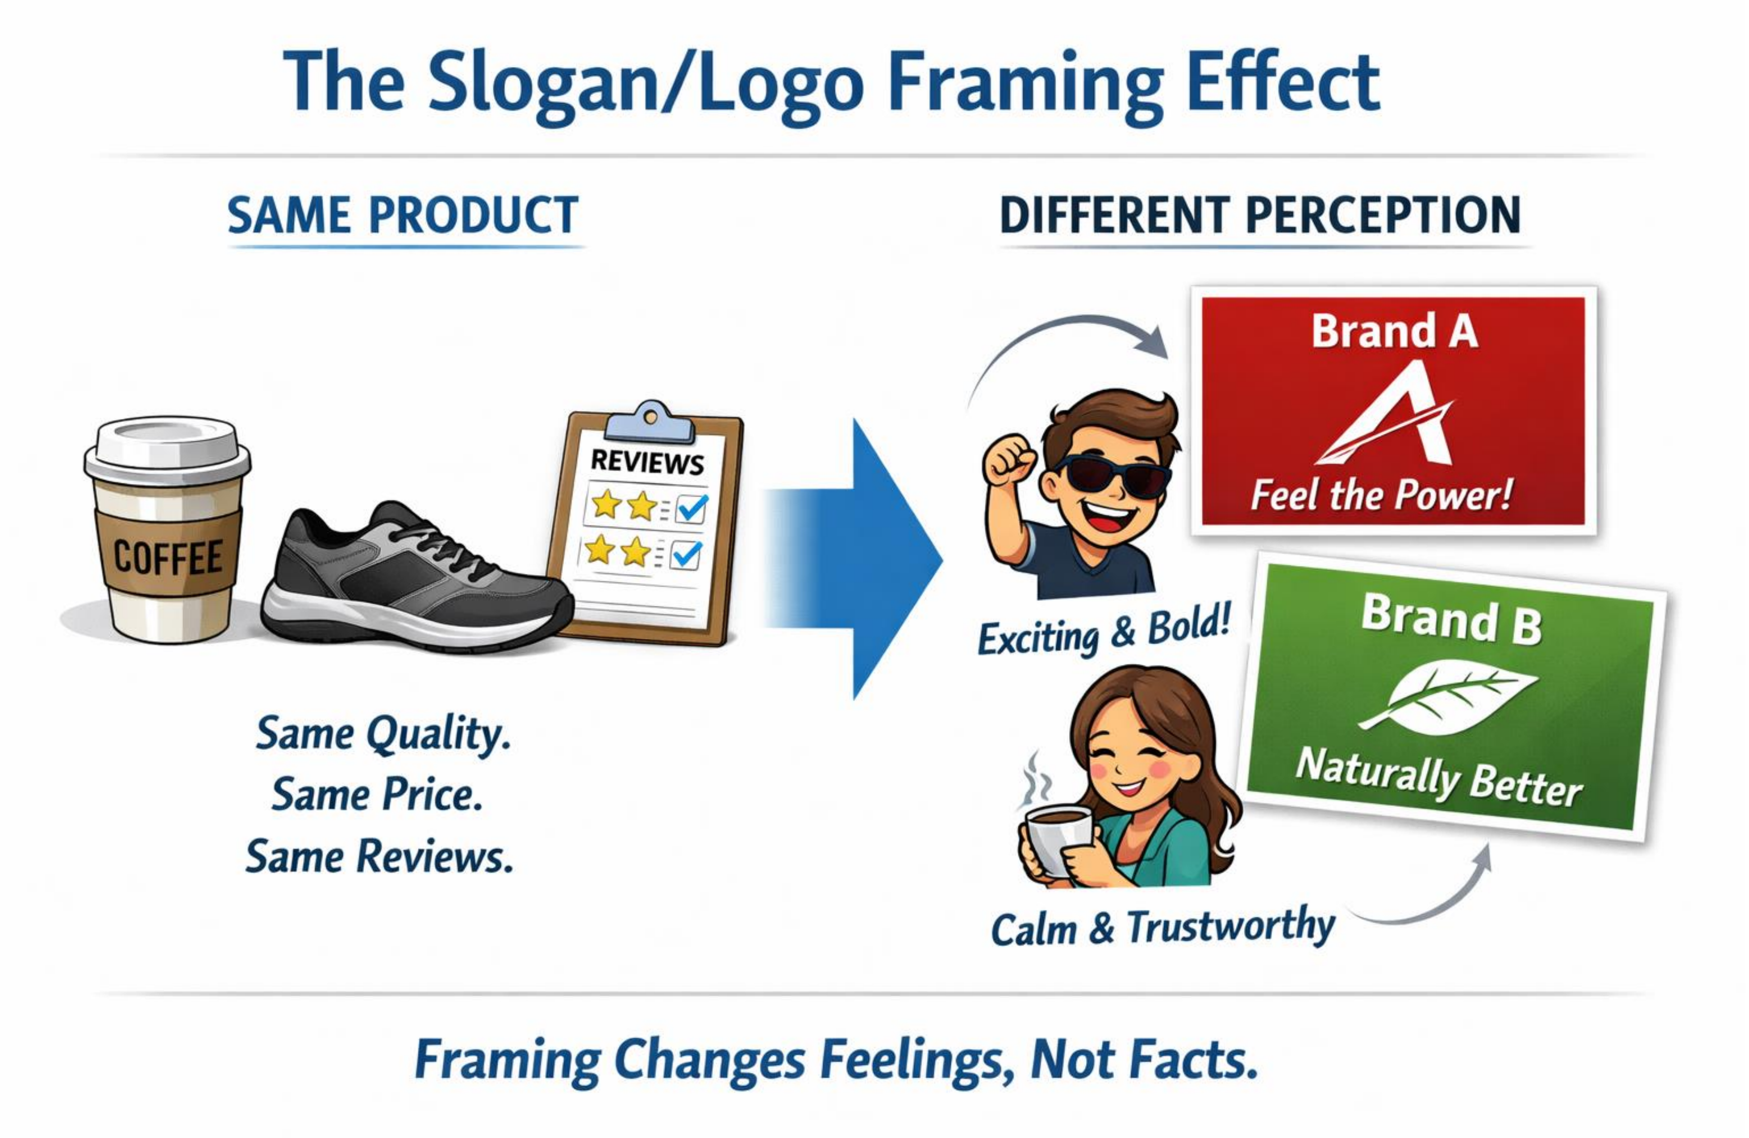

Our Contributions

1) We propose a **theoretical model** for “Information Design with Framing Effect”.

2) We use **Large Language Models** to
   - simulate real-world framing effect, and
   - optimize framing.

A **Theoretical Model** for "Persuasion with Framing Effect"

- Two players: sender  receiver 

- Sender chooses a framing c from a set of framings C
  - Sender has prior belief $\mu_0 \in \Delta(\Omega)$ for the state
  - The framing c shapes the receiver's *prior belief* to be $\mu_c = \ell(c)$
  - $\ell: C \to \Delta(\Omega)$ is a "*belief oracle*"
- With the receiver's prior belief being μ_c, Bayesian Persuasion game happens:
  - Sender designs a signaling scheme $\pi: \Omega \to \Delta(S)$, and sends signal $s \sim \pi(\cdot \mid \omega)$
  - After receiving s, the receiver obtains posterior belief $\mu_c(\cdot \mid s, \pi)$ by Bayes-updating from μ_c, and chooses an optimal action $a^*_{s,\pi}(\mu_c) \in \mathrm{argmax}_{a \in A} \sum_{\omega \in \Omega} \mu_c(\omega | s, \pi) v(a, \omega)$
  - Sender obtains utility $u\left(a^*_{s,\pi}(\mu_c), \omega\right)$

Framing c can be though of a "context":

- does not depend on the state ω, but still affects the receiver's prior belief

(non-Bayesian effect)

We study two sub-problems

- Two players: sender  receiver 

- Sender chooses a framing c from a set of framings C
  - Sender has prior belief $\mu_0 \in \Delta(\Omega)$ for the state
  - The framing c shapes the receiver's *prior belief* to be $\mu_c = \ell(c)$
  - $\ell: C \to \Delta(\Omega)$ is a "*belief oracle*"
- With the receiver's prior belief being μ_c, Bayesian Persuasion game happens:
  - Sender designs a signaling scheme $\pi: \Omega \to \Delta(S)$, and sends signal $s \sim \pi(\cdot \mid \omega)$
  - After receiving s, the receiver obtains posterior belief $\mu_c(\cdot \mid s, \pi)$ by Bayes-updating from μ_c, and chooses an optimal action $a^*_{s,\pi}(\mu_c) \in \operatorname{argmax}_{a \in A} \sum_{\omega \in \Omega} \mu_c(\omega|s,\pi) v(a,\omega)$

Problem 1: Framing-Only Optimization:

Fix π, find $\max_{c \in C} \mathbb{E}_{\omega \sim \mu_0,\, s \sim \pi(\cdot|\omega)}\left[\, u\left(a^*_{s,\pi}(\mu_c), \omega\right) \,\right]$

Problem 2: Joint Optimization:

$$\max_{c \in C,\, \pi: \Omega \to \Delta(S)} \mathbb{E}_{\omega \sim \mu_0,\, s \sim \pi(\cdot|\omega)}\left[\, u\left(a^*_{s,\pi}(\mu_c), \omega\right) \,\right]$$

Main Theoretical Finding:
Joint Optimization *is easier than* **Framing-Only Optimization**

Theorem 1:

Computing the optimal framing c^* is NP-hard

Theorem 2:

There exists a $\text{poly}\left(|\Omega|^{\frac{\log|A|}{\varepsilon^2}} \right)$ time algorithm to compute an ε-optimal (c^*, π^*) pair (under some oracle assumptions)

Problem 1: Framing-Only Optimization:

Fix π, find $\max_{c \in C} \mathbb{E}_{\omega \sim \mu_0,\, s \sim \pi(\cdot|\omega)} \left[u\left(a^*_{s,\pi}(\mu_c), \omega\right) \right]$

Problem 2: Joint Optimization:

$$\max_{c \in C,\, \pi : \Omega \to \Delta(S)} \mathbb{E}_{\omega \sim \mu_0,\, s \sim \pi(\cdot|\omega)} \left[u\left(a^*_{s,\pi}(\mu_c), \omega\right) \right]$$

Main Theoretical Finding:
Joint Optimization *is easier than* **Framing-Only Optimization**

Intuitions:

- Optimizing framing c is equivalent to optimizing prior belief $\mu_c \in B = \{\ell(c): c \in C\}$
- Write the sender's objective as a function of μ_c and π:

$$U(\mu_c, \pi) = \mathbb{E}_{\omega \sim \mu_0,\, s \sim \pi(\cdot|\omega)}\left[\, u\left(a^*_{s,\pi}(\mu_c), \omega\right) \right]$$

- **Observation 1:** *Fixing π, $U(\mu_c, \pi)$ is a* ***discontinuous*** *function of μ_c*
  - Small change in c (small change in μ_c) $\rightarrow$ Small change in posterior belief $\rightarrow$ Sudden change in receiver's action $\rightarrow$ Large change in sender's utility
- **Observation 2:** $U^*(\mu_c) = \max_{\pi:\Omega \to \Delta(S)} U(\mu_c, \pi)$ *is a* ***continuous*** *function of μ_c*

Problem 1: Framing-Only Optimization:

Fix π, find $\max_{c \in C} \mathbb{E}_{\omega \sim \mu_0,\, s \sim \pi(\cdot|\omega)}\left[\, u\left(a^*_{s,\pi}(\mu_c), \omega\right) \right]$

Problem 2: Joint Optimization:

$$\max_{c \in C,\, \pi:\Omega \to \Delta(S)} \mathbb{E}_{\omega \sim \mu_0,\, s \sim \pi(\cdot|\omega)}\left[\, u\left(a^*_{s,\pi}(\mu_c), \omega\right) \right]$$

Our Contributions

1) We propose a **theoretical model** for “Information Design with Framing Effect.

2) We use **Large Language Models** to
   - simulate real-world framing effect, and
   - optimize framing.

Framing-Signaling Joint Optimization using LLM

We use LLM to do two things:

- Simulate the framing-to-belief oracle ℓ: $c \mapsto \mu_c$
- Optimize framing c

(Iterative Prompt Optimization)

e.g., receiver information

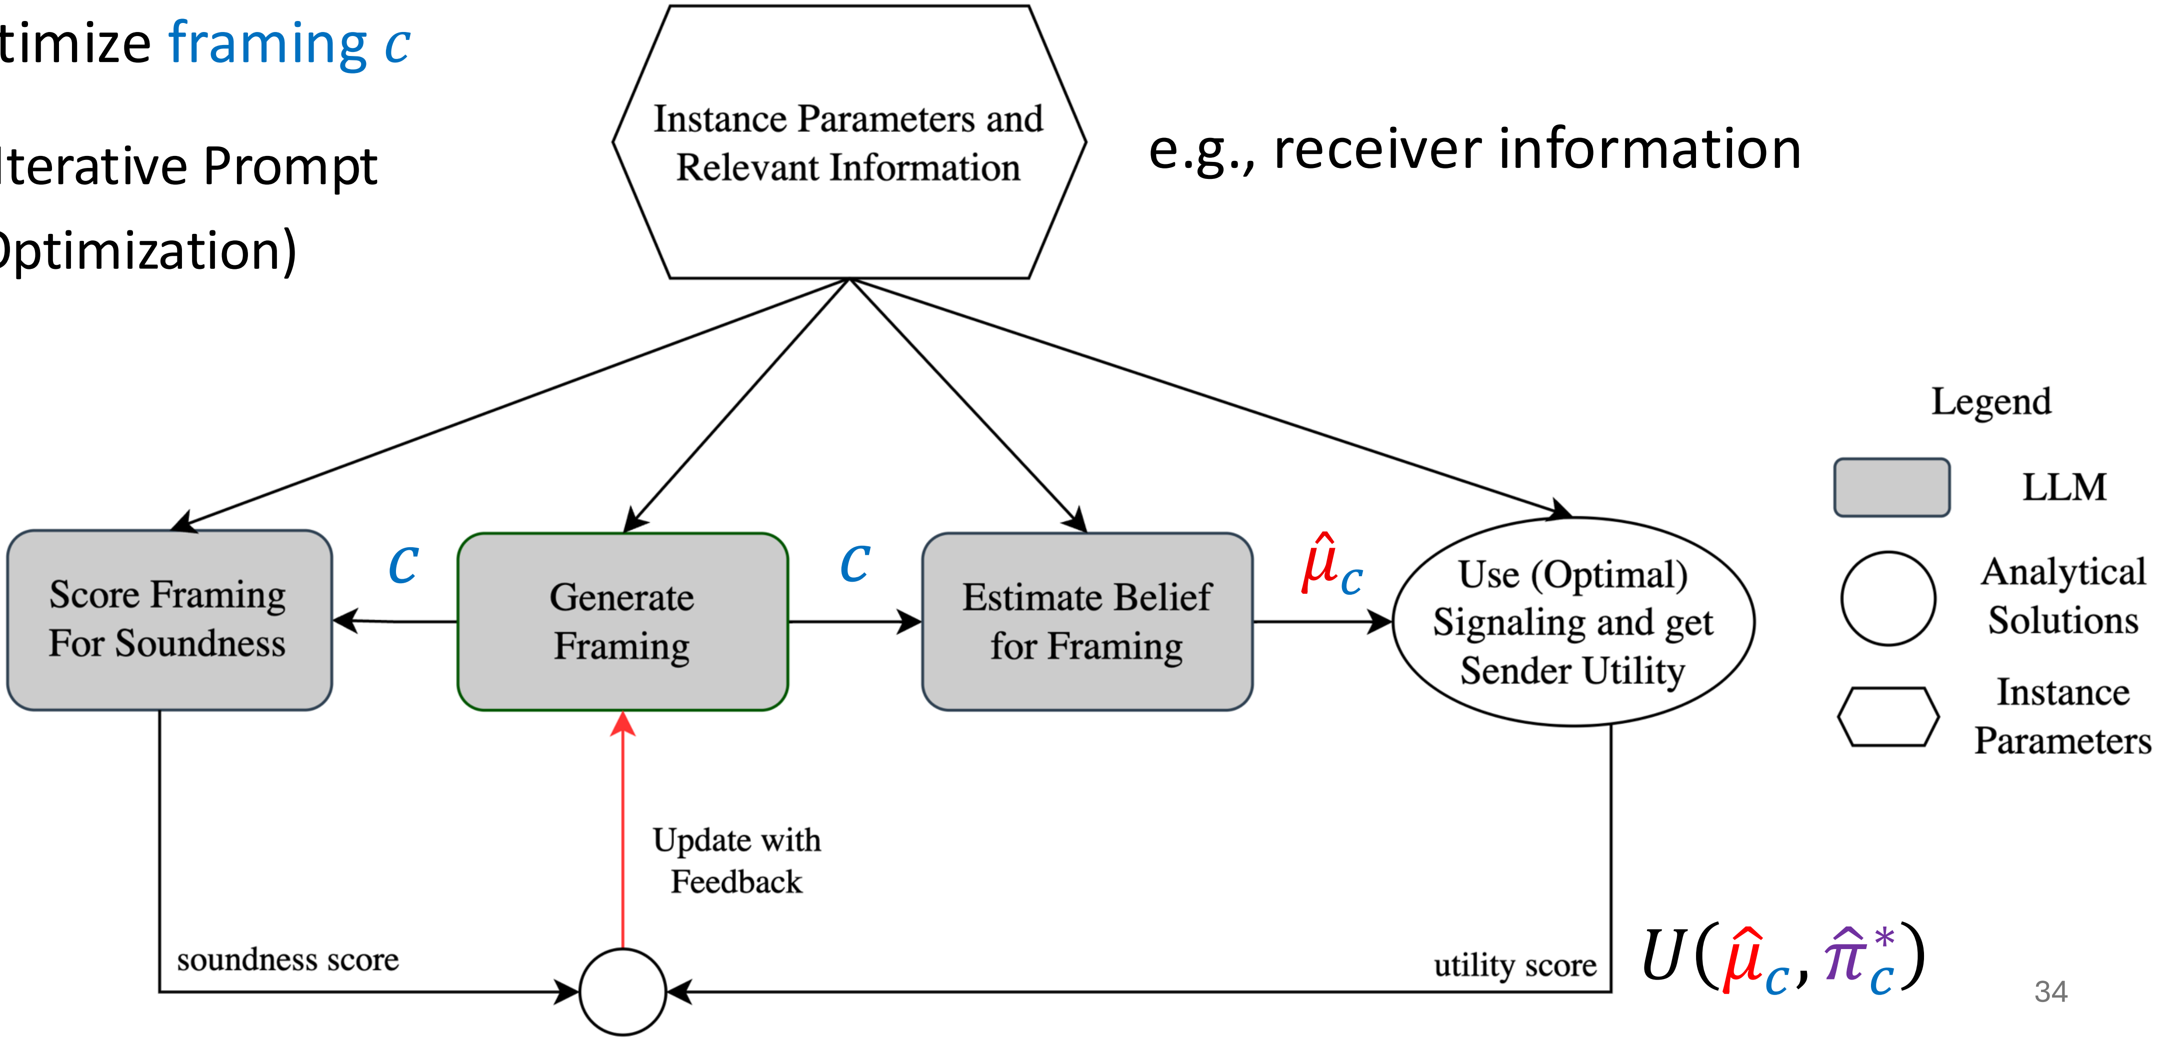

Case Study: House Buying

Sender: a realtor
(house-selling agent)

State: quality of a house

Framing c: description of the realtor:

*"Meet **Jeremy Hammond**, a dedicated realtor with over 8 years of experience, specializing in finding the perfect homes <u>for outdoor enthusiasts like you</u>.... Trust Jeremy to help you discover a home that <u>complements your active lifestyle while staying within your budget</u>."*

State-independent!

Signal (recommendation) s: "buy" or "not buy"

Receiver: a potential house-buyer

Henry lives in Boston and is an avid <u>outdoor person</u> who enjoys hiking and being in nature. For him, a "good" house has <u>low maintenance</u>, affords <u>easy access to trails, biking, running etc</u>, and far from the main city. He is single and <u>doesn't like a family-oriented house</u>. He is looking for houses less than $500,000.

Experiment 1: Use LLM to simulate the belief oracle ℓ: $c \mapsto \mu_c$

Why?

- Previous works showed that prompted LLM can simulate specific group of people.
- When people delegate decisions to AI agents, we will persuade AI agents.

How?

- Provide the realtor description c to LLM (without recommendation s)
- Ask LLM to output the buyer's prior belief about the state of the house.

Chat-GPT 4o

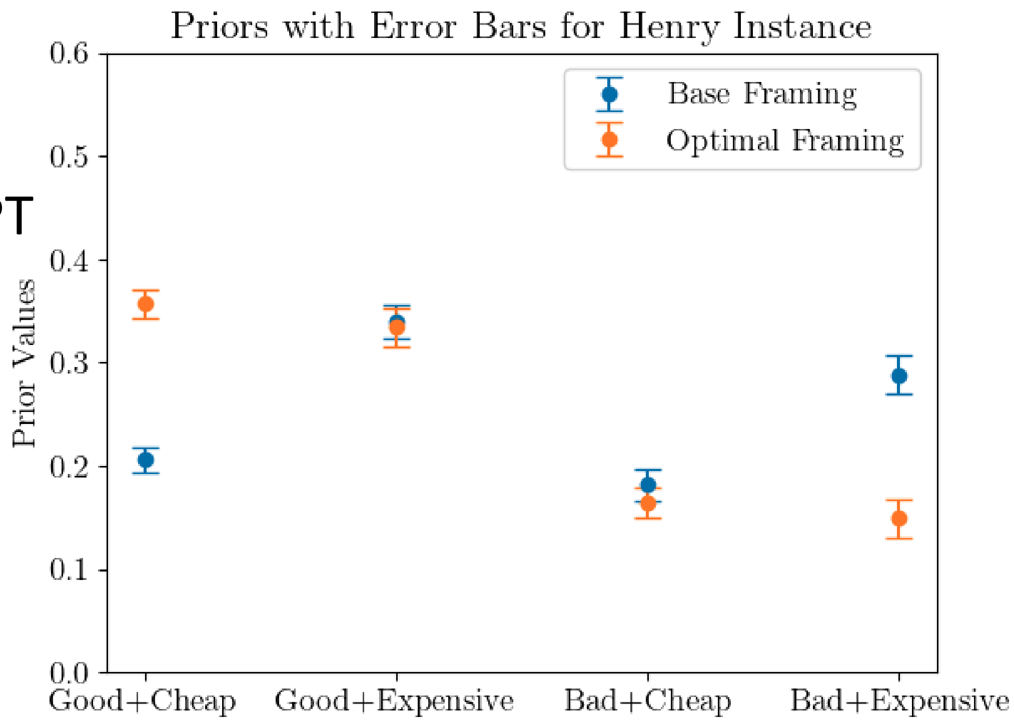

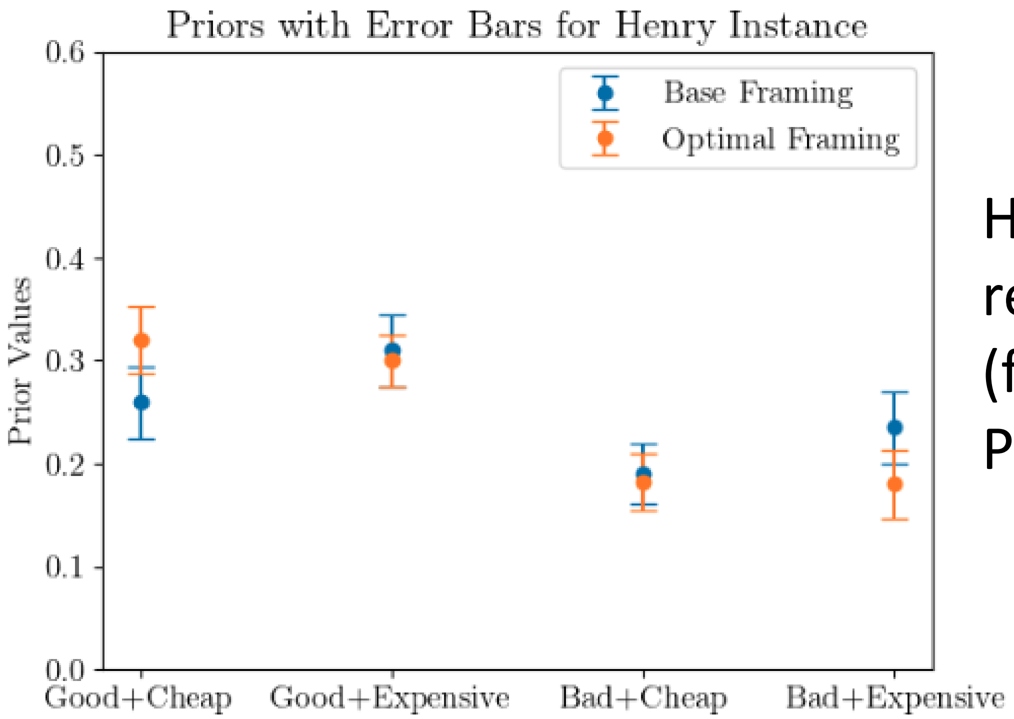

Human responses (from Prolific)

Experiment 2.1: Use LLM to optimize framing c

Framing for the "Henry" Instance	Utility
No framing - receiver prior equal to sender prior	0.28
Realtor Jeremy Profile from the Instance Description	0.30
Best LLM Framing: *Meet Jeremy Hammond, a dedicated realtor with over 8 years of experience, specializing in finding the perfect homes for outdoor enthusiasts like you. Living in Downtown Boston, Jeremy understands the balance between city life and access to nature. With a background as a contractor, he ensures that every property meets your low-maintenance needs. When he's not helping clients, you can find him hiking local trails or enjoying his backyard garden. Trust Jeremy to help you discover a home that complements your active lifestyle while staying within your budget.*	0.40
Analytical Upper Bound (Optimal Joint Strategy when $B = \Delta(\Omega)$)	0.41

LLM generates sentences not in the given realtor profile, tailored to Henry

Experiment 2.2: Use LLM to optimize framing c

Framing for the "Lilly" Instance	**Utility**
No framing - receiver prior equal to sender prior	0.33
Realtor Jeremy Profile from the Instance Description	0.33
Best LLM Framing: *Introducing Jeremy Hammond, a seasoned realtor with 8 years dedicated to helping families find their dream homes in Boston's suburbs. With a rich background as a contractor, Jeremy excels in identifying spacious, family-friendly properties with excellent school districts—just what you need for your kids. As a fellow dog owner, he knows the importance of a great yard and a welcoming neighborhood. Trust Jeremy to leverage his local expertise and commitment to family values as he guides you to affordable yet quality homes that fit your family's lifestyle.*	0.42
Analytical Upper Bound (Optimal Joint Strategy when $B = \Delta(\Omega)$)	0.46

LLM generates a different realtor description for another house-buyer

Outline

- Background on a Classical Information Design Model:

  “Bayesian Persuasion” [Kamenica & Gentzkow, 2011]

- Information Design with a Learning Receiver

- Information Design with Large Language Models

- Summary and one more thing

Summary: Information Design + Learning & LLM

- **Commitment**: *Learning outcomes might differ from classic outcomes*
  - Sender can commit to a randomized mapping ("signaling scheme") $\pi: \Omega \to \Delta(S)$ before state realization.
- **Bayesian receiver:**
  - Knowing the prior μ and signaling scheme π, the receiver does Bayes update after receiving signal s (and best responds)

- **Abstract signal space:**
  - Language doesn't matter – only the correlation between signal and state matters. *Capture framing effect by theory and LLM*

Many research opportunities!

My Research Interests

"Learning-Based Incentive Design":

- *Information Design*
- *Mechanism Design*
- *Algorithmic Game Theory*
- *Multi-Agent Learning*
- ……