# Information Design with Large Language Models


Tao Lin

Microsoft Research → CUHK-Shenzhen


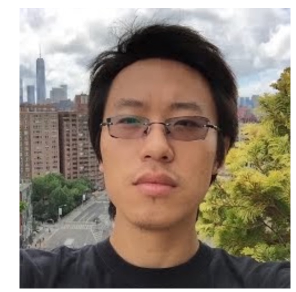


Safwan Hossain
Harvard

Paul Dütting
Google

Renato Paes Leme
Google

Sai S Ravindranath
Harvard

Haifeng Xu
UChicago

Song Zuo
Google


EC 2026

# Information Design

An economic scenario where *one player ("Sender") strategically reveals information to influence the decision of another player ("Receiver").*


*Example: advertising*

# Information Design

An economic scenario where *one player ("Sender") strategically reveals information to influence the decision of another player ("Receiver").*

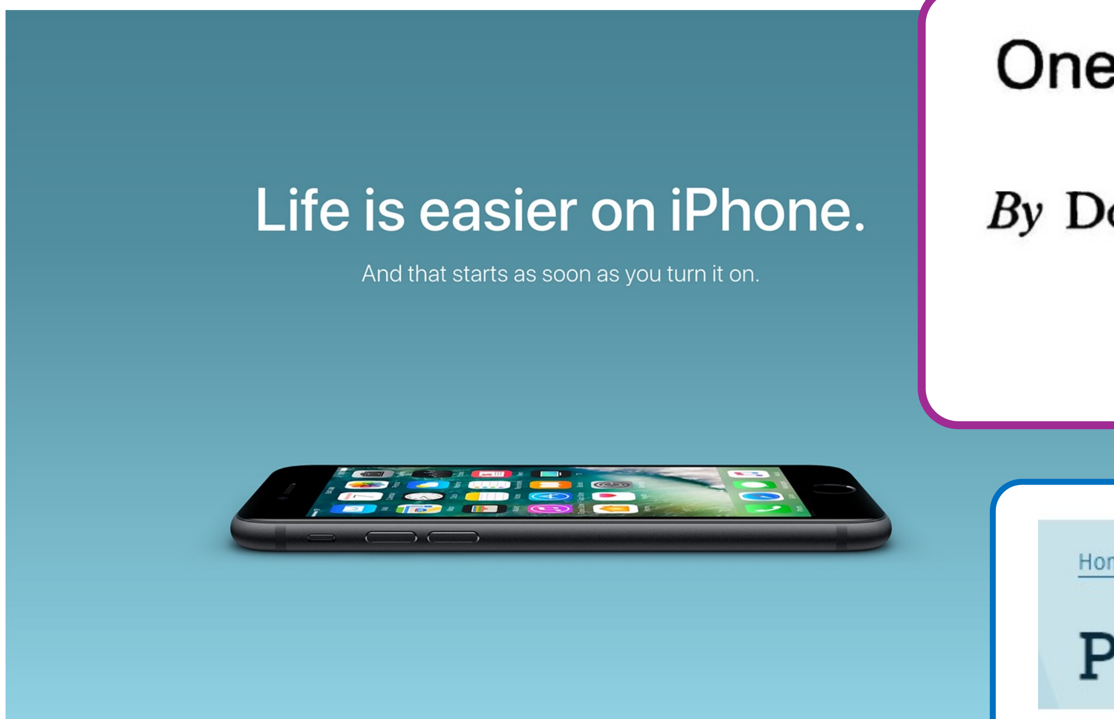


*Example: advertising*

**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

# Information Design

An economic scenario where *one player ("Sender") strategically reveals information to influence the decision of another player ("Receiver").*

Many economic models about information design and communication games:

- "Information Disclosure Games" (Grossman, 1981; Milgrom 1981)
- "Cheap Talk" (Crawford & Sobel, 1982)
- "Bayesian Persuasion" (Kamenica & Gentzkow, 2011)
- …

# A Common Assumption in Classical Models

**Bayesian Assumption**

- The information transmitted from Sender to Receiver is modeled by a random variable $s$ correlated with the state of the world $\omega$

$$\omega \sim \mu_0, \qquad s \sim \pi(\cdot|\omega)$$

- Receiver does Bayes update after receiving $s$:

$$\mu(\omega|s)$$

and takes a best-responding action.

- Importantly, *how the signal $s$ is communicated ("framed")* doesn't matter.
- But real-world decision-makers are affected by the framing of information.

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

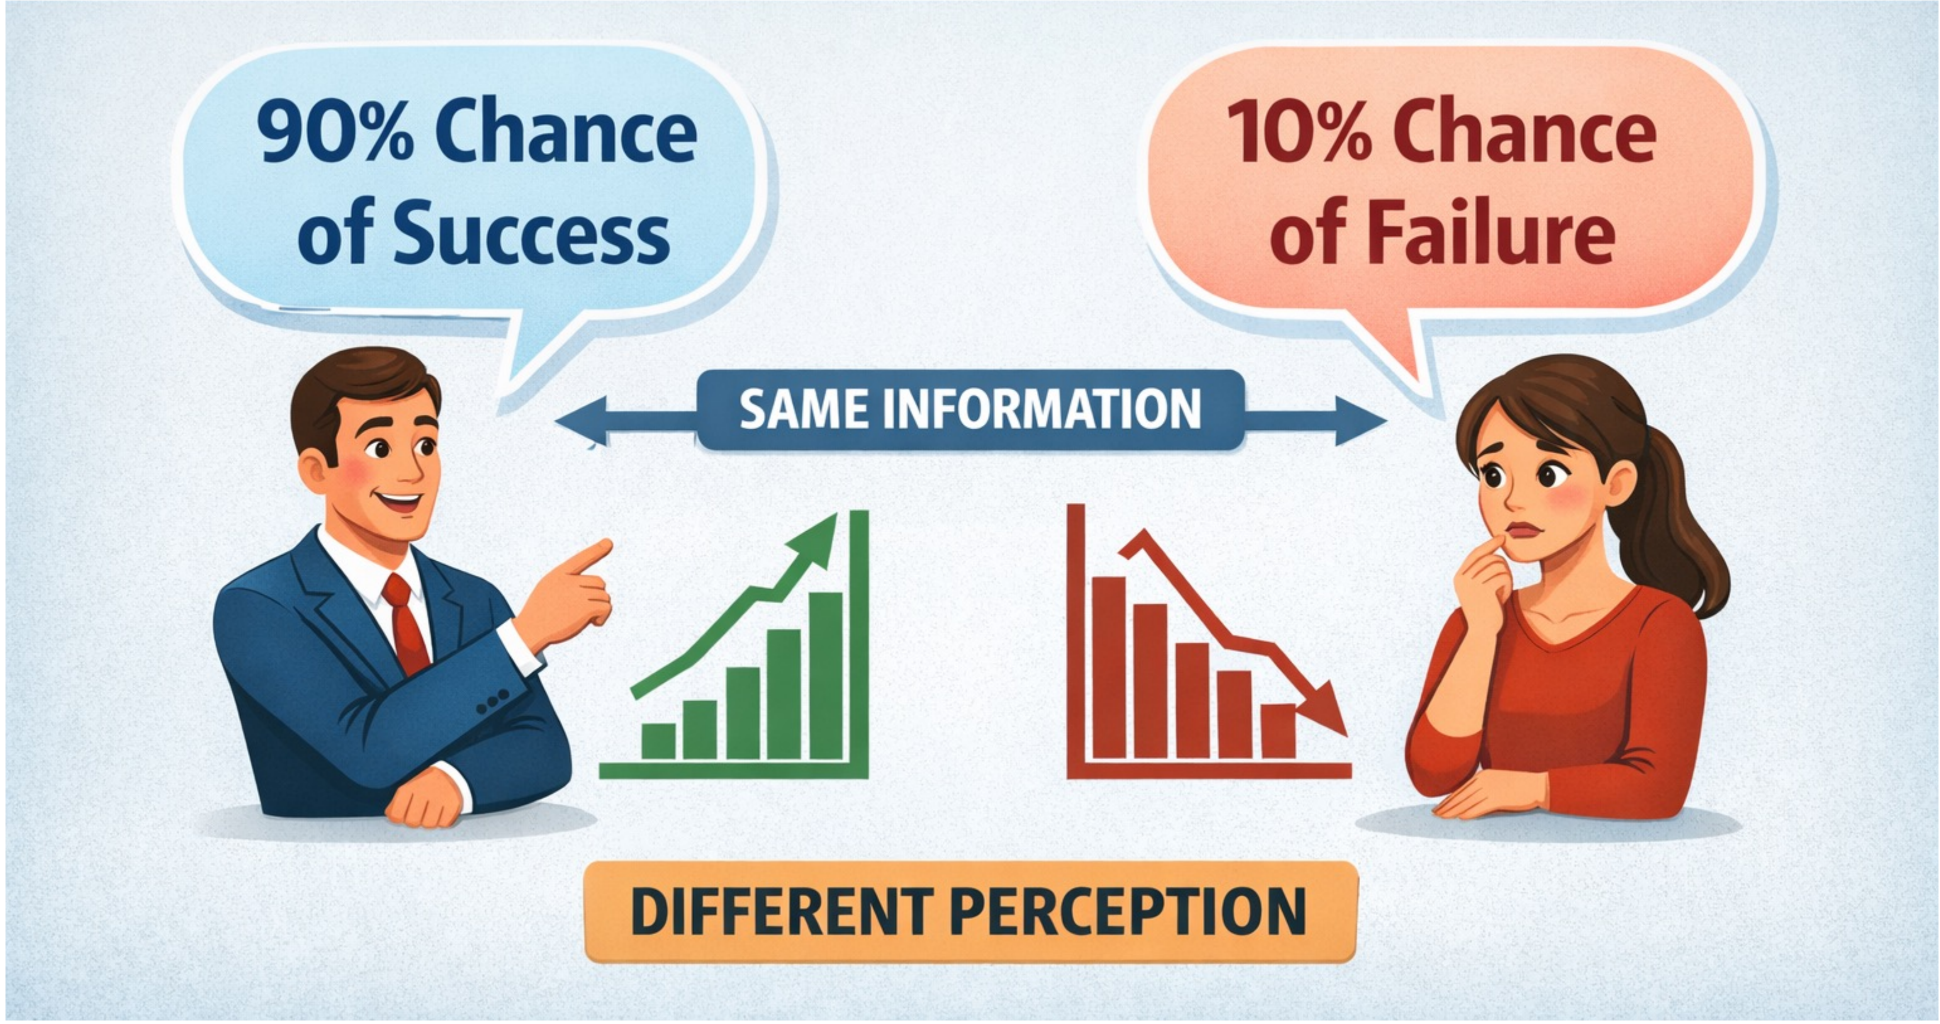

# **Example 2:** Advertising – Slogan/Logo Design

Our work aims to incorporate (non-Bayesian) framing effect

into classical (Bayesian) information design models.

# Our Contributions

1) **Propose a theoretical model** that incorporates framing effect into the classical Bayesian persuasion model:
   - Propose framing optimization problems;
   - Characterize their complexity.

2) **Demonstrate the ability of large language models (LLMs)** to
   - Simulate real-world framing effect;
   - Optimize framing.

# Theoretical Model for "Persuasion with Framing Effect"

- Two players: Sender 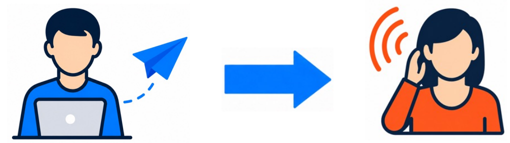 Receiver
- A state of the world $\omega \sim \mu_0 \in \Delta(\Omega)$. Sender knows the prior distribution $\mu_0$.

# Theoretical Model for "Persuasion with Framing Effect"

- Two players: Sender 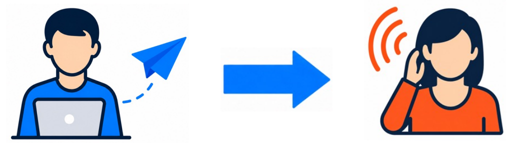 Receiver
- A state of the world $\omega \sim \mu_0 \in \Delta(\Omega)$. Sender knows the prior distribution $\mu_0$.
- Sender chooses a framing $c$ from a set of framings $C$:
  - The choice of $c$ is independent of $\omega$
  - $c$ shapes the Receiver's *subjective prior belief* to be $\mu_c = \ell(c)$
  - $\ell: C \rightarrow \Delta(\Omega)$ is a *framing-to-belief mapping*

# Theoretical Model for "Persuasion with Framing Effect"

- Two players: Sender 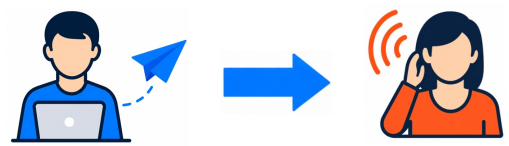 Receiver
- A state of the world $\omega \sim \mu_0 \in \Delta(\Omega)$. Sender knows the prior distribution $\mu_0$.
- Sender chooses a framing $c$ from a set of framings $C$:
  - The choice of $c$ is independent of $\omega$
  - $c$ shapes the Receiver's *subjective prior belief* to be $\mu_c = \ell(c)$
  - $\ell: C \to \Delta(\Omega)$ is a *framing-to-belief mapping*

- With priors $(\mu_0, \ \mu_c)$, (standard) Bayesian Persuasion happens:
  - Sender designs a signaling scheme $\pi: \Omega \to \Delta(S)$, and sends signal $s \sim \pi(\cdot \mid \omega)$
  - Given $s$, Receiver obtains posterior belief $\mu_c(\cdot \mid s, \pi)$ via Bayes update from $\mu_c$, 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 $u\left(a^*_{s,\pi}(\mu_c), \omega\right)$

# The Advertising Example

**Sender:** a clothing brand

**4 states of products:**
{Fashionable, Functional} × {Durable, Not Durable}

**Receiver:** consumer

**3 actions:** { buy on sale,
buy at regular price,
not buy }

# The Advertising Example

**Sender:** a clothing brand

**4 states of products:**
{Fashionable, Functional} × {Durable, Not Durable}

**Receiver:** consumer

**3 actions:** { buy on sale,
buy at regular price,
not buy }

**Framing** $c$ (product-independent):

**Slogan:** *"Build the best product."*

**Description of product line:** *"Whether you're seeking boulders and alpine views or stalking rainbow trout, this low-profile, waterproof jacket will keep you dry and provide rain protection from the fork in the trail to the places less traveled, and back again. This 3-layer shell meets our H2No® Performance Standard for exceptional waterproof/breathable protection, and the fabric, membrane and durable water repellent (DWR) finish are made without added PFAS."*

# The Advertising Example

**Sender:** a clothing brand

**Receiver:** consumer

**4 states of products:**
{Fashionable, Functional} × {Durable, Not Durable}

**3 actions:** { buy on sale,
buy at regular price,
not buy }

**Framing** $c$ (product-independent):

**Slogan:** *"Build the best product."*

**Description of product line:** *"Whether you're seeking boulders and alpine views or stalking rainbow trout, this low-profile, waterproof jacket will keep you dry and provide rain protection from the fork in the trail to the places less traveled, and back again. This 3-layer shell meets our H2No® Performance Standard for exceptional waterproof/breathable protection, and the fabric, membrane and durable water repellent (DWR) finish are made without added PFAS."*

**Signaling scheme** $\pi$:

$$\text{product state} \rightarrow \begin{cases} \text{Ad with Discount} \\ \text{Ad without Discount} \\ \text{No Ad} \end{cases}$$

# Two Optimization Problems

- Two players: Sender  Receiver
- A state of the world $\omega \sim \mu_0 \in \Delta(\Omega)$. Sender knows the prior distribution $\mu_0$.
- Sender chooses a framing $c$ from a set of framings $C$:
  - The choice of $c$ is independent of $\omega$
  - $c$ shapes the Receiver's *subjective prior belief* to be $\mu_c = \ell(c)$
  - $\ell: C \to \Delta(\Omega)$ is a *framing-to-belief mapping*
- With priors $(\mu_0, \ \mu_c)$, (standard) Bayesian Persuasion happens:
  - Sender designs a signaling scheme $\pi: \Omega \to \Delta(S)$, and sends signal $s \sim \pi(\cdot \mid \omega)$
  - Given $s$, Receiver obtains posterior belief $\mu_c(\cdot \mid s, \pi)$ via Bayes update from $\mu_c$, 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 $u\big(a^*_{s,\pi}(\mu_c), \omega\big)$

# Two Optimization Problems

- Two players: Sender  Receiver
- A state of the world $\omega \sim \mu_0 \in \Delta(\Omega)$. Sender knows the prior distribution $\mu_0$.
- Sender chooses a framing $c$ from a set of framings $C$:
  - The choice of $c$ is independent of $\omega$
  - $c$ shapes the Receiver's *subjective prior belief* to be $\mu_c = \ell(c)$
  - $\ell: C \to \Delta(\Omega)$ is a *framing-to-belief mapping*
- With priors $(\mu_0,\ \mu_c)$, (standard) Bayesian Persuasion happens:
  - Sender designs a signaling scheme $\pi: \Omega \to \Delta(S)$, and sends signal $s \sim \pi(\cdot \mid \omega)$
  - Given $s$, Receiver obtains posterior belief $\mu_c(\cdot \mid s, \pi)$ via Bayes update from $\mu_c$, and

**Problem 1: Framing-Only Optimization**

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

**Problem 2: Joint Optimization**

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

# Theoretical Simplification: Optimizing $\mu_c$ instead of $c$

Assume:

- The *framing-to-belief mapping* $\ell: c \mapsto \mu_c$ can be perfectly simulated (by, e.g., LLM)
- Then, optimizing framing $c$ is equivalent to optimizing the induced prior belief

$$\mu \in B = \{\ell(c): c \in C\}$$

- Further assume that $B$ is given.

**Problem 1: Framing-Only Optimization**

Fix $\pi$, 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**

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

# Theoretical Simplification: Optimizing $\mu_c$ instead of $c$

Assume:

- The *framing-to-belief mapping* $\ell: c \mapsto \mu_c$ can be perfectly simulated (by, e.g., LLM)
- Then, optimizing framing $c$ is equivalent to optimizing the induced prior belief

$$\mu \in B = \{\ell(c): c \in C\}$$

- Further assume that $B$ is given.

**Problem 1: Framing-Only Optimization**

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

**Problem 2: Joint Optimization**

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

# Theoretical Result 1: Computational Complexity

Both problems are **NP-hard**!

**Theorem 1:** For some fixed $\pi$, computing the optimal framing (prior) $\mu_\pi^*$ is NP-hard, even when $B = \Delta(\Omega)$

**Theorem 2:** Computing the optimal $(\mu^*, \pi^*)$ pair is NP-hard, for some simple $B \subset \Delta(\Omega)$

**Problem 1: Framing-Only Optimization**

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

**Problem 2: Joint Optimization**

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

# Theoretical Result 1: Computational Complexity

Both problems are **NP-hard**!

**Theorem 1:** For some fixed $\pi$, computing the optimal framing (prior) $\mu_\pi^*$ is NP-hard, even when $B = \Delta(\Omega)$

Proved by a reduction from Bayesian Stackelberg Games [Conitzer & Sandholm, EC 2006]:

$\mu \leftrightarrow$ leader's strategy, $s \leftrightarrow$ follower's type

**Theorem 2:** Computing the optimal $(\mu^*, \pi^*)$ pair is NP-hard, for some simple $B \subset \Delta(\Omega)$

Proved by a reduction from Independent Set

**Problem 1: Framing-Only Optimization**

Fix $\pi$, find $\max_{\mu \in B} \mathbb{E}_{\omega \sim \mu_0,\, s \sim \pi(\cdot|\omega)} \left[ u\left(a_{s,\pi}^*(\mu), \omega\right) \right]$

**Problem 2: Joint Optimization**

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

# Theoretical Result 2: Continuity

The sender's objective function is

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

**Problem 1: Framing-Only Optimization**

Fix $\pi$, find $\max_{\mu \in B} U(\mu, \pi)$

**Problem 2: Joint Optimization**

$$\max_{\mu \in B,\, \pi : \Omega \to \Delta(S)} U(\mu, \pi) = \max_{\mu \in B} U^*(\mu)$$

# Theoretical Result 2: Continuity

The sender's objective function is

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

- **FO is not continuous:** Fixing $\pi$, $U(\mu, \pi)$ is **discontinuous** in $\mu$
  - *Small change in $c$ (small change in $\mu_c$) $\rightarrow$ Small change in posterior belief $\rightarrow$ Sudden change in receiver's action $\rightarrow$ Large change in sender's utility*

**Problem 1: Framing-Only Optimization**

Fix $\pi$, find $\max_{\mu \in B} U(\mu, \pi)$

**Problem 2: Joint Optimization**

$$\max_{\mu \in B,\, \pi:\Omega \to \Delta(S)} U(\mu, \pi) = \max_{\mu \in B} U^*(\mu)$$

# Theoretical Result 2: Continuity

The sender's objective function is

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

- **FO is not continuous:** Fixing $\pi$, $U(\mu, \pi)$ is **discontinuous** in $\mu$
  - *Small change in $c$ (small change in $\mu_c$) → Small change in posterior belief → Sudden change in receiver's action → Large change in sender's utility*
- **JO is continuous:** With optimized $\pi$, $U^*(\mu) = \max_{\pi:\Omega\to\Delta(S)} U(\mu, \pi)$ is **continuous** in $\mu$

**Problem 1: Framing-Only Optimization**

Fix $\pi$, find $\max_{\mu\in B} U(\mu, \pi)$

**Problem 2: Joint Optimization**

$$\max_{\mu\in B,\, \pi:\Omega\to\Delta(S)} U(\mu, \pi) = \max_{\mu\in B} U^*(\mu)$$

# Theoretical Result 2: Continuity

The sender's objective function is

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

- **FO is not continuous:** Fixing $\pi$, $U(\mu, \pi)$ is **discontinuous** in $\mu$
  - *Small change in $c$ (small change in $\mu_c$) $\rightarrow$ Small change in posterior belief $\rightarrow$ Sudden change in receiver's action $\rightarrow$ Large change in sender's utility*
- **JO is continuous:** With optimized $\pi$, $U^*(\mu) = \max_{\pi:\Omega\rightarrow\Delta(S)} U(\mu, \pi)$ is **continuous** in $\mu$

***Joint Optimization** might be more tractable than **Framing-Only Optimization***

**Problem 1: Framing-Only Optimization**

Fix $\pi$, find $\max_{\mu \in B} U(\mu, \pi)$

**Problem 2: Joint Optimization**

$$\max_{\mu \in B,\, \pi:\Omega\rightarrow\Delta(S)} U(\mu, \pi) = \max_{\mu \in B} U^*(\mu)$$

# Our Contributions

1) **Propose a theoretical model** that incorporates framing effect into the classical Bayesian persuasion model:
   - Propose framing optimization problems;
   - Characterize their complexity.

2) **Demonstrate the ability of large language models (LLMs)** to
   - Simulate real-world framing effect;
   - Optimize framing.

# **Experiment Pipeline** (Framing-Signaling Joint Optimization)

We use LLM to:

- *Simulate* the framing-to-belief mapping $\ell: c \mapsto \mu_c$
- *Optimize* framing $c$

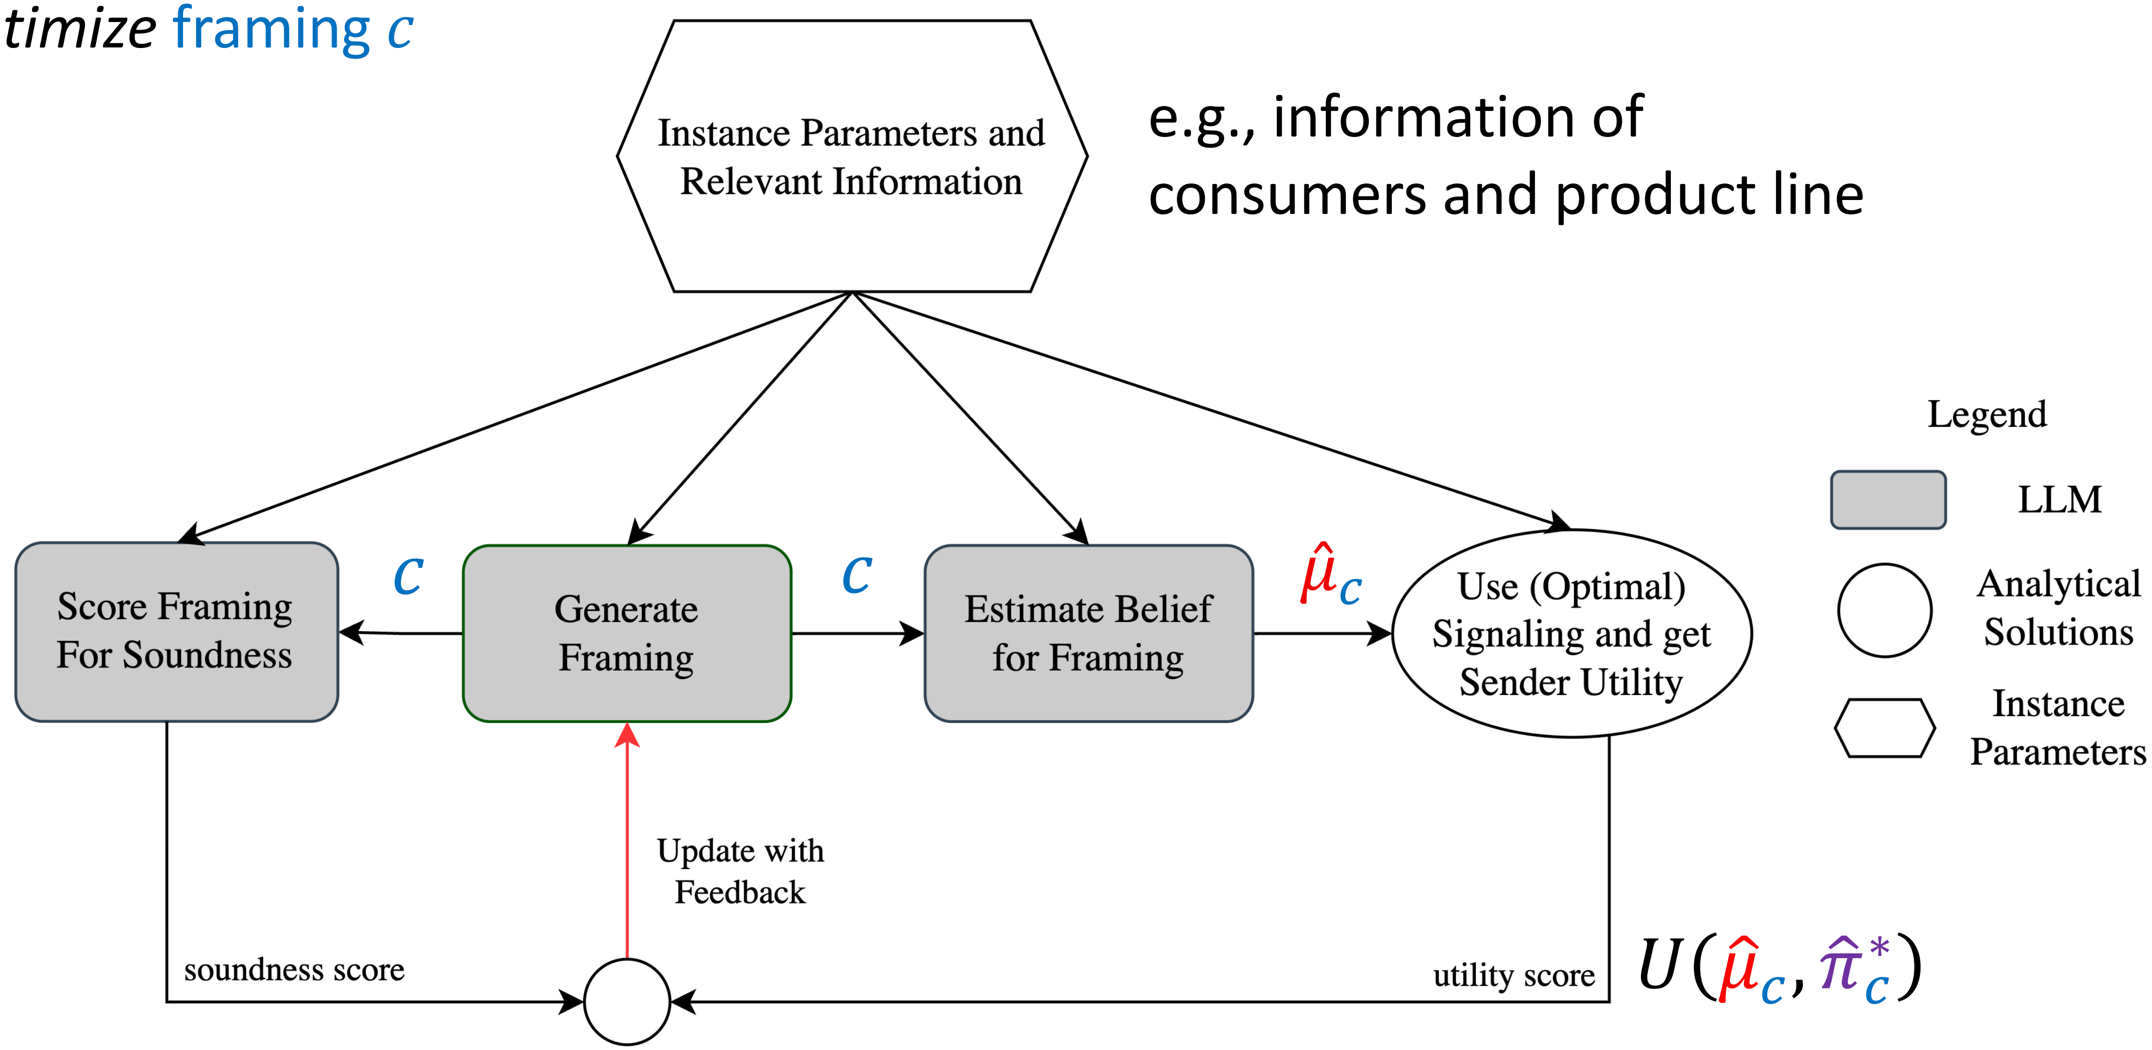

# Case Study: Advertising

**Sender:** a clothing brand

**4 states of products:**
{Fashionable, Functional} × {Durable, Not Durable}

**Receiver:**
a target group of consumers

**3 actions:** { buy on sale,
buy at regular price,
not buy }

**Framing $c$:**

**Slogan:** *"Build the best product."*
**Description of product line:** *"Whether you're seeking boulders and alpine views or stalking rainbow trout, this low-profile, waterproof jacket will keep you dry and provide rain protection from the fork in the trail to the places less traveled, and back again. This 3-layer shell meets our H2No® Performance Standard for exceptional waterproof/breathable protection, and the fabric, membrane and durable water repellent (DWR) finish are made without added PFAS."*

**Consumer Information:**

"They are fashion-aware consumers. They like the idea of performance wear (e.g., hiking and skiing) but are not deeply familiar with technical characteristics. Functionality and durability is a nice bonus, but aesthetic appeal primarily drives their interest. Lululemon, Nike are main competitors in this market segment."

**3 signals:** {Ad with Discount, Ad without Discount, No Ad}

# Experiment 1: Use LLM to simulate $\ell: c \mapsto \mu_c$

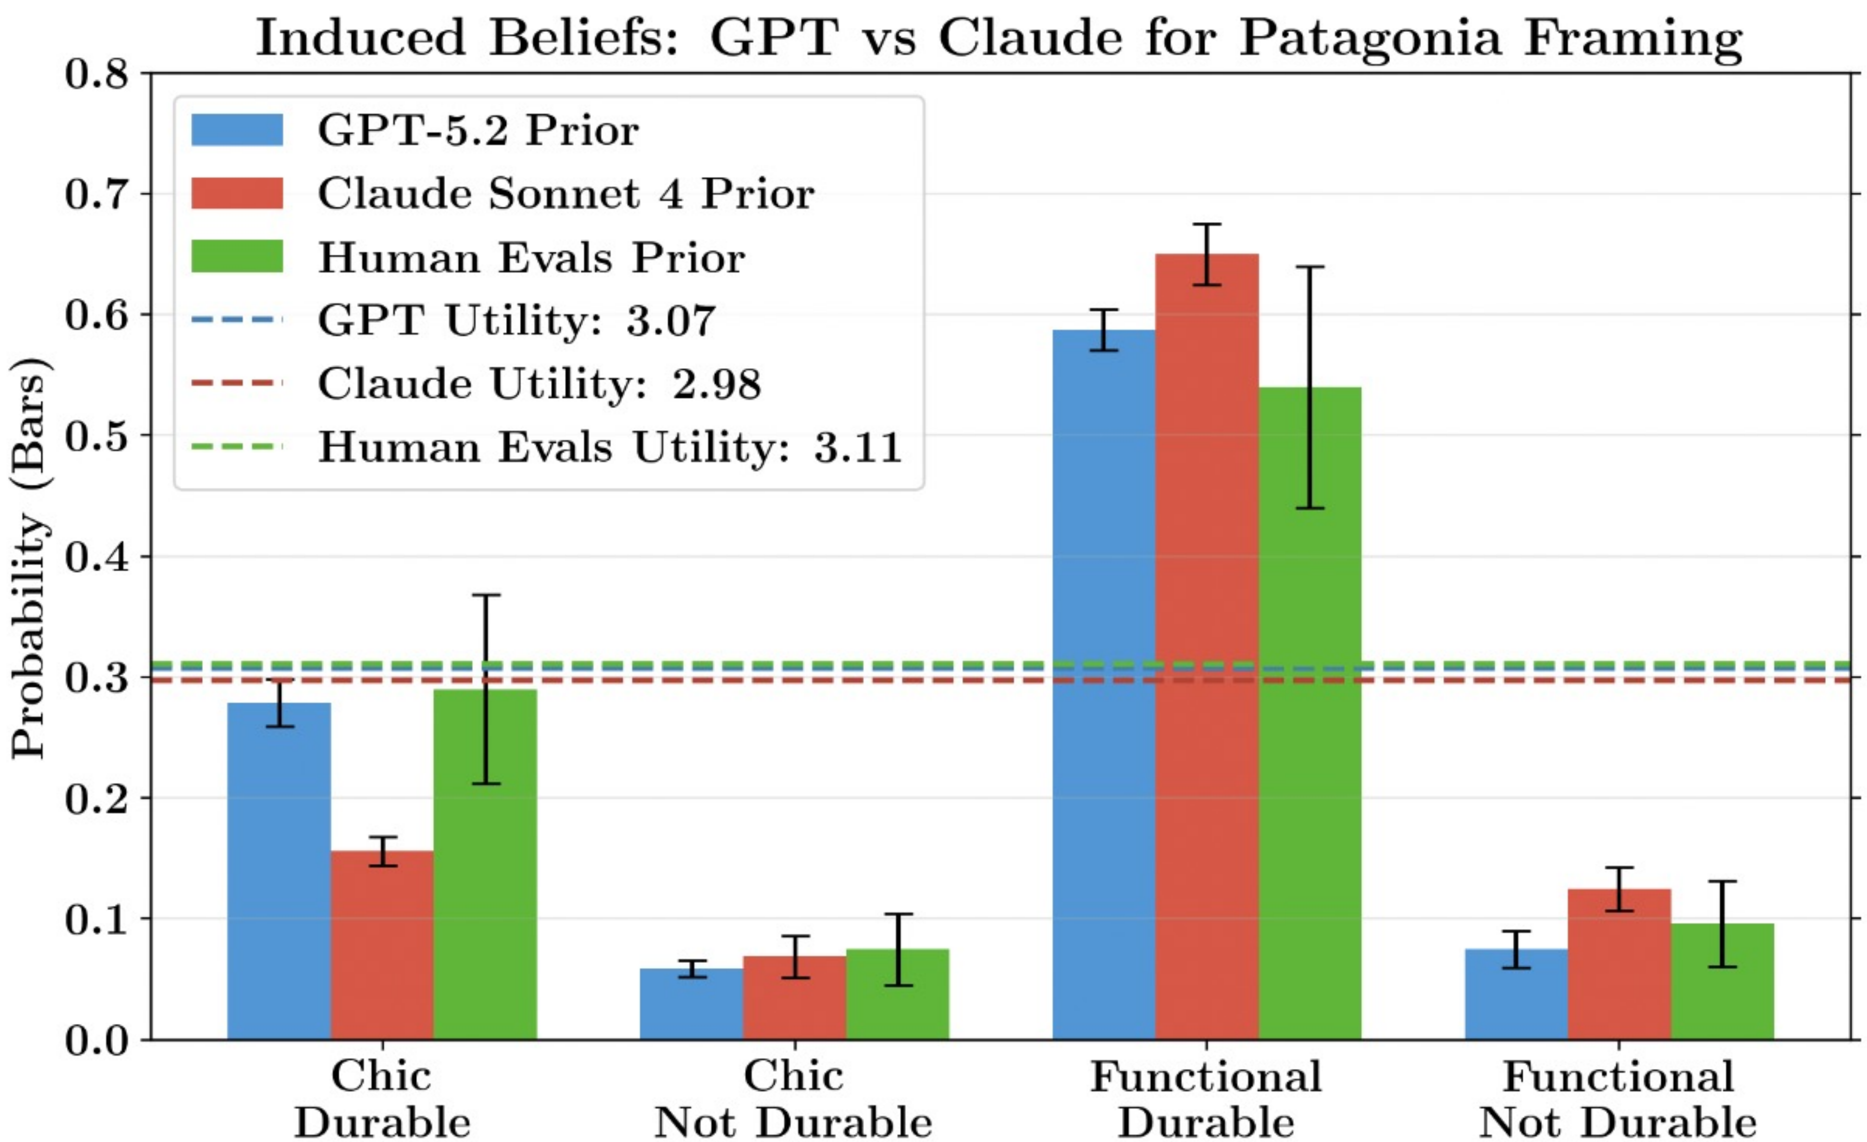


*(Human responses are from Prolific)*

# Experiment 2:
# Use LLM to optimize framing $c$

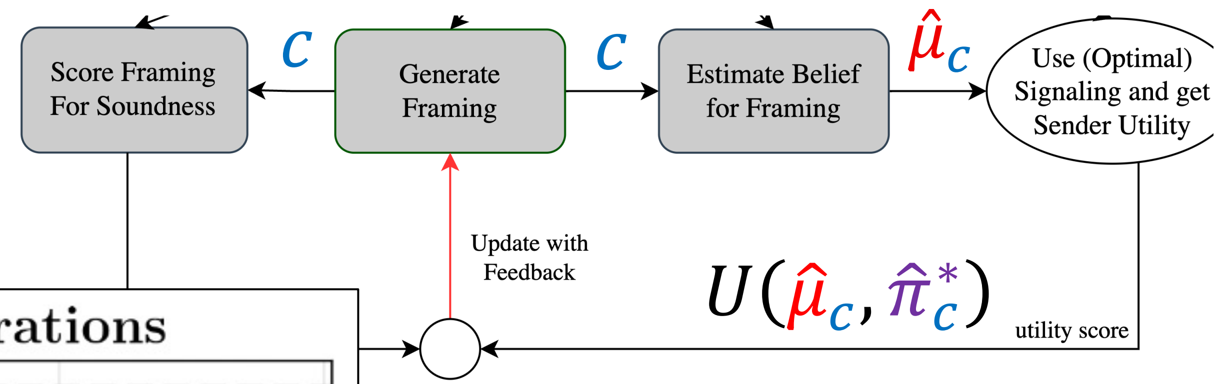


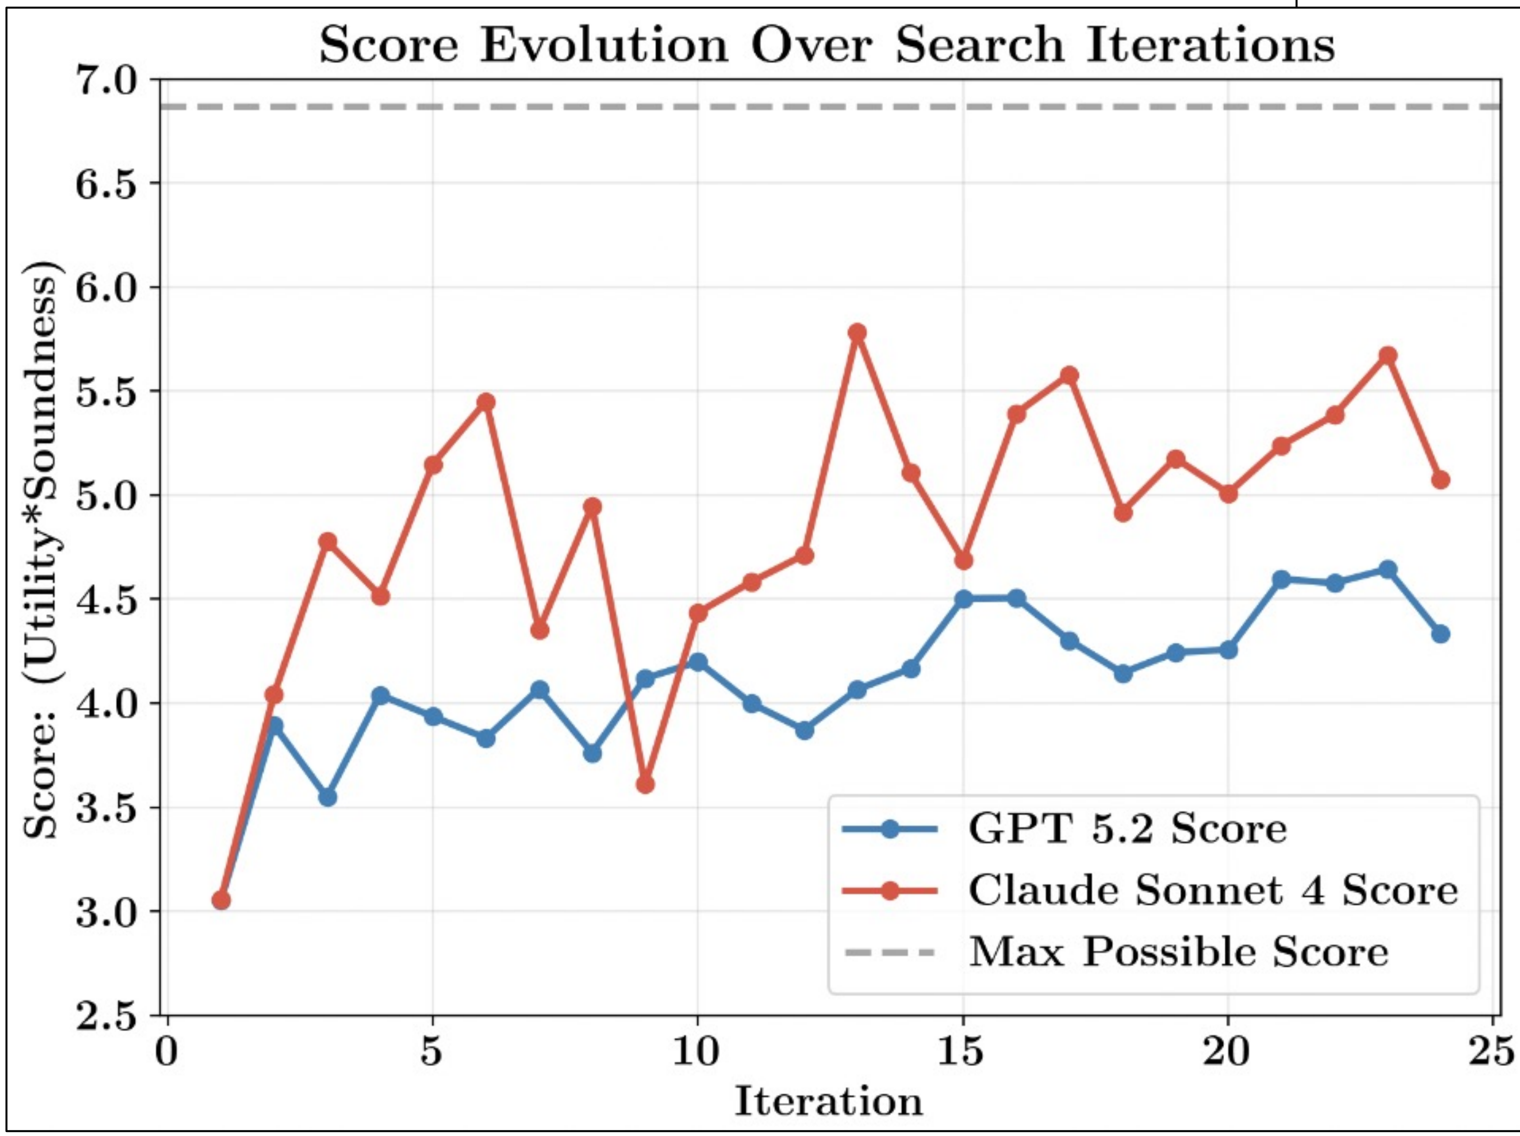

# Example Framings Found by LLM

**Original:**

**Slogan:** *Build the best product.*
**Product Line Description:**
*Whether you're seeking boulders and alpine views or stalking rainbow trout, this low-profile, waterproof jacket will keep you dry and provide rain protection from the fork in the trail to the places less traveled, and back again. This 3-layer shell meets our H2No® Performance Standard for exceptional waterproof/breathable protection, and the fabric, membrane and durable water repellent (DWR) finish are made without added PFAS.*

**GPT Result:**

**Slogan:** *Where style meets substance.*
**Product Line Description:**
*Himalaya's new outerwear brings modern, everyday style to real weather: streamlined parkas, ski jackets and pants, windbreakers, and thermal layers that look clean on campus, in cafés, and on the go. The fits are sleek and easy to layer—made to pair with denim, joggers, and sneakers. Underneath the minimal finish: 100% recycled nylon ripstop, a durable water-repellent finish, and high performance waterproof/breathable comfort for rainy days and weekend trips.*

# Example Framings Found by LLM

**Original:**

**Slogan:** *Build the best product.*
**Product Line Description:**
*Whether you're seeking boulders and alpine views or stalking rainbow trout, this low-profile, waterproof jacket will keep you dry and provide rain protection from the fork in the trail to the places less traveled, and back again. This 3-layer shell meets our H2No® Performance Standard for exceptional waterproof/breathable protection, and the fabric, membrane and durable water repellent (DWR) finish are made without added PFAS.*

**Claude Result:**

**Slogan:** *City Sharp. Mountain Ready.*
**Product Line Description:** *Our newest collection redefines modern outerwear with clean lines, contemporary cuts, and versatile designs that transition effortlessly from city to countryside. Each piece features premium sustainable fabrics and thoughtful construction details that ensure lasting quality without compromising on style. Perfect for the modern lifestyle – whether you're commuting to work, meeting friends for brunch, or enjoying weekend adventures. These aren't just clothes, they're wardrobe essentials that reflect your values and elevate your everyday look with confidence and sophistication.*

# Summary



1) We propose a theoretical model that incorporates framing effects into the classical Bayesian persuasion model:
   - Both framing-only and joint optimization problems are hard, but joint optimization might be easier.
2) We demonstrate LLM's ability to simulate real-world framing effect and optimize framing.

# Summary



1) We propose a theoretical model that incorporates framing effects into the classical Bayesian persuasion model:
   - Both framing-only and joint optimization problems are hard, but joint optimization might be easier.
2) We demonstrate LLM's ability to simulate real-world framing effect and optimize framing.

- Our work is just one initial attempt to bridge classical information design theory with real-world persuasion, using recent LLM technologies.
- Other works: "Verbalized Bayesian Persuasion" (Li et al, ICML 2026)
  "AI Realtor" (Wu et al, CAIS 2026)
  ……

***Questions?***