Soft Information, Hard Decisions: AI Advising

Jing Huang
Texas A&M University

Shumiao Ouyang
University of Oxford

Joint OMI & SBS Conference on FinTech and Blockchain Economics

June 2026

Advice, Inverted

- Classic friction in expert advice: conflict of interest (Crawford-Sobel)

Advice, Inverted

- Classic friction in expert advice: conflict of interest (Crawford-Sobel)
- AI advisors **invert** this friction:
  - No conflict of interest — no commission to chase, no product to push
  - But may misunderstand what the user wants when it depends on her *soft traits*
  - . . . and the user may not fully know herself

Advice, Inverted

- Classic friction in expert advice: conflict of interest (Crawford-Sobel)
- AI advisors **invert** this friction:
  - No conflict of interest — no commission to chase, no product to push
  - But may misunderstand what the user wants when it depends on her *soft traits*
  - . . . and the user may not fully know herself
- **Research Question**: How do prompt quality, memory, and the AI's prior determine the value of AI advice?

Advice, Inverted

- Classic friction in expert advice: conflict of interest (Crawford-Sobel)
- AI advisors **invert** this friction:
  - No conflict of interest — no commission to chase, no product to push
  - But may misunderstand what the user wants when it depends on her *soft traits*
  - . . . and the user may not fully know herself
- **Research Question**: How do prompt quality, memory, and the AI's prior determine the value of AI advice?

Key friction: communicating soft information to a machine

AI doesn't know user's preference $\Rightarrow$ User must debias AI's recommendation

Literature

This Paper

This Paper

- **Theory**: Preference uncertainty + cheap talk framework
  - AI cannot observe user's preference type (soft information)
  - Key friction: AI recommends based on its belief; user must debias

This Paper

- **Theory**: Preference uncertainty + cheap talk framework
  - AI cannot observe user's preference type (soft information)
  - Key friction: AI recommends based on its belief; user must debias
- **Key Results**:
  - Prompt quality can substitute for memory
  - Memory benefits depend on prior alignment
  - Opinionated AI can be optimal

This Paper

- **Theory**: Preference uncertainty + cheap talk framework
  - AI cannot observe user's preference type (soft information)
  - Key friction: AI recommends based on its belief; user must debias
- **Key Results**:
  - Prompt quality can substitute for memory
  - Memory benefits depend on prior alignment
  - Opinionated AI can be optimal
- **Empirics**: LLM simulations with 500 investor profiles
  - Two-sided simulation: both advisor and investor are LLM-generated
  - Validate theoretical predictions

Model

Setting: Two Layers of Uncertainty

- **Layer 1: Preference Uncertainty** (Novel)
  - Investor's preference type $\omega \in \{0, 1\}$
  - $\omega = 1$: target is $\tilde{\theta}_1$ (equity); $\omega = 0$: target is $\tilde{\theta}_0$ (fixed income)
  - AI cannot directly observe ω

Setting: Two Layers of Uncertainty

- **Layer 1: Preference Uncertainty** (Novel)
  - Investor's preference type $\omega \in \{0, 1\}$
  - $\omega = 1$: target is $\tilde{\theta}_1$ (equity); $\omega = 0$: target is $\tilde{\theta}_0$ (fixed income)
  - AI cannot directly observe ω
- **Layer 2: Fundamental Uncertainty**
  - $\tilde{\theta}_1 \sim N(\mu_1, 1)$, $\tilde{\theta}_0 \sim N(\mu_0, 1)$

Setting: Two Layers of Uncertainty

- **Layer 1: Preference Uncertainty** (Novel)
  - Investor's preference type $\omega \in \{0, 1\}$
  - $\omega = 1$: target is $\tilde{\theta}_1$ (equity); $\omega = 0$: target is $\tilde{\theta}_0$ (fixed income)
  - AI cannot directly observe ω
- **Layer 2: Fundamental Uncertainty**
  - $\tilde{\theta}_1 \sim N(\mu_1, 1)$, $\tilde{\theta}_0 \sim N(\mu_0, 1)$
- **Beliefs**: $p = \mathbb{P}^{\text{investor}}(\omega = 1)$, $\hat{p} = \mathbb{P}^{\text{AI}}(\omega = 1)$

Setting: Two Layers of Uncertainty

- **Layer 1: Preference Uncertainty** (Novel)
  - Investor's preference type $\omega \in \{0, 1\}$
  - $\omega = 1$: target is $\tilde{\theta}_1$ (equity); $\omega = 0$: target is $\tilde{\theta}_0$ (fixed income)
  - AI cannot directly observe ω
- **Layer 2: Fundamental Uncertainty**
  - $\tilde{\theta}_1 \sim N(\mu_1, 1)$, $\tilde{\theta}_0 \sim N(\mu_0, 1)$
- **Beliefs**: $p = \mathbb{P}^{\text{investor}}(\omega = 1)$, $\hat{p} = \mathbb{P}^{\text{AI}}(\omega = 1)$
- **Investor's Utility**: Quadratic loss

$$U = -p(a - \theta_1)^2 - (1 - p)(a - \theta_0)^2$$

Full Notation

The AI Advisor and the Debiasing Problem

- AI has no conflict of interest but doesn't know user's preference type ω

The AI Advisor and the Debiasing Problem

- AI has no conflict of interest but doesn't know user's preference type ω
- AI observes (θ_1, θ_0) and recommends:

$$\boxed{m = \hat{p}\theta_1 + (1 - \hat{p})\theta_0}$$

The AI Advisor and the Debiasing Problem

- AI has no conflict of interest but doesn't know user's preference type ω
- AI observes (θ_1, θ_0) and recommends:

$$\boxed{m = \hat{p}\theta_1 + (1 - \hat{p})\theta_0}$$

- **The Core Friction**:
  - If $\hat{p} \neq p$: recommendation is "biased" from user's perspective
  - User must debias: extract signal about θ_ω from m

The AI Advisor and the Debiasing Problem

- AI has no conflict of interest but doesn't know user's preference type ω
- AI observes (θ_1, θ_0) and recommends:

$$\boxed{m = \hat{p}\theta_1 + (1 - \hat{p})\theta_0}$$

- **The Core Friction**:
  - If $\hat{p} \neq p$: recommendation is "biased" from user's perspective
  - User must debias: extract signal about θ_ω from m
- **Trade-off**:
  - More extreme $\hat{p} \Rightarrow$ more informative about one fundamental
  - But larger debiasing cost if $\hat{p} \neq p$

Communication Details Memory Details

The Stopping Value

- Payoff upon stopping:

$$g(p,\hat{p}) = -\underbrace{\frac{(p-\hat{p})^2}{\hat{p}^2+(1-\hat{p})^2}}_{\text{Debiasing Cost}} - \underbrace{p(1-p)(\Delta_\mu^2+2)}_{\text{Pref. Uncertainty}}$$

The Stopping Value

- Payoff upon stopping:

$$g(p, \hat{p}) = -\underbrace{\frac{(p-\hat{p})^2}{\hat{p}^2 + (1-\hat{p})^2}}_{\text{Debiasing Cost}} - \underbrace{p(1-p)(\Delta_\mu^2 + 2)}_{\text{Pref. Uncertainty}}$$

- **Debiasing Cost**:
  - Zero when $p = \hat{p}$
  - Extreme $\hat{p} \Rightarrow$ more informative

The Stopping Value

- Payoff upon stopping:

$$g(p,\hat{p}) = -\underbrace{\frac{(p-\hat{p})^2}{\hat{p}^2+(1-\hat{p})^2}}_{\text{Debiasing Cost}} - \underbrace{p(1-p)(\Delta_\mu^2+2)}_{\text{Pref. Uncertainty}}$$

- **Debiasing Cost**:
  - Zero when $p = \hat{p}$
  - Extreme $\hat{p} \Rightarrow$ more informative
- **Pref. Uncertainty**:
  - Highest at $p = 0.5$

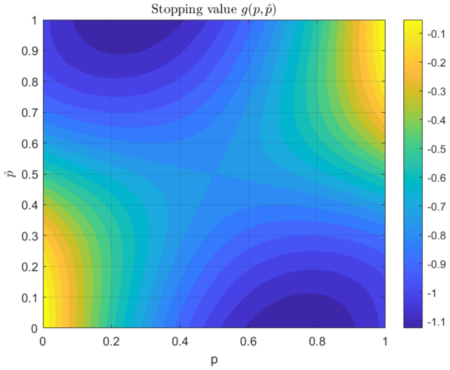

Equilibrium: Threshold Policy

- When $\kappa = 0$ (no memory):

Equilibrium: Threshold Policy

- When $\kappa = 0$ (no memory):
- **Threshold Structure**:
  - Continue if $p \in (\underline{p}, \overline{p})$
  - Stop otherwise

Equilibrium: Threshold Policy

- When $\kappa = 0$ (no memory):
- **Threshold Structure**:
  - Continue if $p \in (\underline{p}, \overline{p})$
  - Stop otherwise
- **Symmetric**: $\underline{p} + \overline{p} = 1$

Equilibrium: Threshold Policy

- When $\kappa = 0$ (no memory):
- **Threshold Structure**:
  - Continue if $p \in (\underline{p}, \overline{p})$
  - Stop otherwise
- **Symmetric**: $\underline{p} + \overline{p} = 1$
- **Intuition**:
  - Quadratic $g(p)$, symmetric volatility
  - States $\omega \in \{0, 1\}$ interchangeable

Equilibrium: Threshold Policy

- When $\kappa = 0$ (no memory):
- **Threshold Structure**:
  - Continue if $p \in (\underline{p}, \overline{p})$
  - Stop otherwise
- **Symmetric**: $\underline{p} + \overline{p} = 1$
- **Intuition**:
  - Quadratic $g(p)$, symmetric volatility
  - States $\omega \in \{0, 1\}$ interchangeable
- **Closed-form**: Up to one unknown ($\overline{p}$)

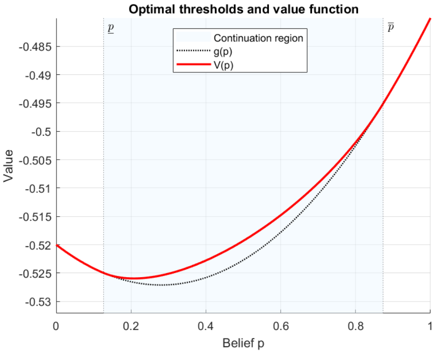

HJB Equation

Design Levers

Three Key Results

Three Key Results

1. **Prompt Quality $\leftrightarrow$ Memory**
   - High prompt quality $\Rightarrow$ belief converges to certainty
   - At certainty, even minimal memory suffices
   - Good prompts reduce need for long context windows

Three Key Results

1. **Prompt Quality $\leftrightarrow$ Memory**
   - High prompt quality $\Rightarrow$ belief converges to certainty
   - At certainty, even minimal memory suffices
   - Good prompts reduce need for long context windows
2. **Memory Benefits $\leftrightarrow$ Prior Alignment**
   - Aligned priors: more memory always helps
   - Misaligned priors: more memory can *hurt* (overshooting)

Three Key Results

1. **Prompt Quality $\leftrightarrow$ Memory**
   - High prompt quality $\Rightarrow$ belief converges to certainty
   - At certainty, even minimal memory suffices
   - Good prompts reduce need for long context windows

2. **Memory Benefits $\leftrightarrow$ Prior Alignment**
   - Aligned priors: more memory always helps
   - Misaligned priors: more memory can *hurt* (overshooting)

3. **Opinionated AI Can Be Optimal**
   - Lean aggressive? Prefer aggressive AI
   - Informativeness effect dominates alignment effect

Optimal AI Figure

LLM Simulation Evidence

Simulation Design

- **Data**: 500 investor profiles from SCF 2022
  - Ground truth: Vanguard Investor Questionnaire

Simulation Design

- **Data**: 500 investor profiles from SCF 2022
  - Ground truth: Vanguard Investor Questionnaire
- **LLM**: OpenAI GPT-5, two-sided simulation

Simulation Design

- **Data**: 500 investor profiles from SCF 2022
  - Ground truth: Vanguard Investor Questionnaire
- **LLM**: OpenAI GPT-5, two-sided simulation
- **Three Advisor Conditions**:
  - No Memory / Full Memory / Full Information

Simulation Design

- **Data**: 500 investor profiles from SCF 2022
    - Ground truth: Vanguard Investor Questionnaire
- **LLM**: OpenAI GPT-5, two-sided simulation
- **Three Advisor Conditions**:
    - No Memory / Full Memory / Full Information
- **Scale**: 2,500 conversations (500 $\times$ 5 iterations)

Full Design Details

H1: Self-Learning Dominates

Stage	Accuracy
Prior (before any interaction)	69.0
Final Before Rec (after Q&A, before rec)	83.4
Final (after seeing recommendation)	86.6
Improvement from interaction alone	+14.4pp
Additional improvement from recommendation	+3.2pp

Finding: Primary value is investor self-learning, not personalized recommendations

Each additional Q&A round adds $\sim$0.77pp to accuracy

Regression Results

H2: Impatience Hurts

	Advisor Rec. Accuracy
Exogenous Termination	-0.99***
	(0.39)
Profile FE	YES

Finding: Impatience-driven (exogenous) termination cuts conversations short, reducing recommendation accuracy by $\sim$1pp

$\Rightarrow$ Validates model prediction that impatience hurts

H3: Memory Helps

Condition	Accuracy
No Memory	83.5
Full Memory	88.3
Full Information	89.8

- Memory: +4.8pp
- Full Info gap: +1.5pp

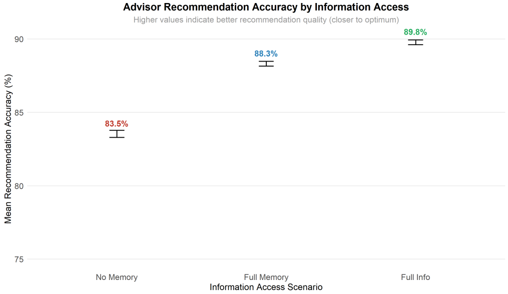

H4: Opinionated AI Optimal

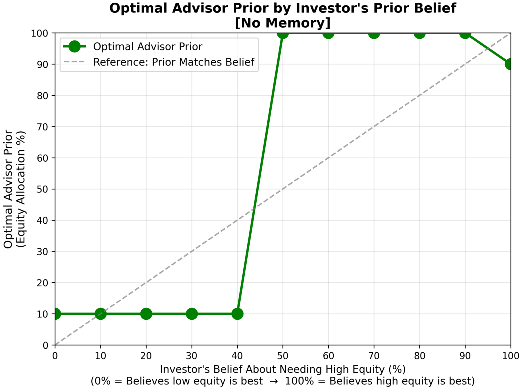

Key finding: Optimal $\hat{p}_0^*$ curve matches theoretical prediction

Investors benefit from advisors more extreme than their own prior beliefs

Methodology

Summary

- **Question**: AI advisors have no conflict of interest — what determines the value of their advice?

Summary

- **Question**: AI advisors have no conflict of interest — what determines the value of their advice?
- **Answer**: Debiasing friction — AI doesn't know user's preference

Summary

- **Question**: AI advisors have no conflict of interest — what determines the value of their advice?
- **Answer**: Debiasing friction — AI doesn't know user's preference
- **Key Insights**:
  1. Prompt quality can substitute for memory
  2. Memory benefits depend on prior alignment
  3. Opinionated AI can be optimal

Summary

- **Question**: AI advisors have no conflict of interest — what determines the value of their advice?
- **Answer**: Debiasing friction — AI doesn't know user's preference
- **Key Insights**:
  1. Prompt quality can substitute for memory
  2. Memory benefits depend on prior alignment
  3. Opinionated AI can be optimal
- **Empirical Validation**:
  - Self-learning dominates (+14.4pp)
  - Memory helps (+4.8pp)
  - Opinionated AI pattern validated

Thank You!

Questions and Comments Welcome

Jing Huang jing.huang@tamu.edu
Shumiao Ouyang shumiao.ouyang@sbs.ox.ac.uk

Related Literature

- **AI in Economics and Finance**
  - Ide & Naganuma (2025); Chen et al. (2026); Ferreira (2026)
- **Cheap Talk and Financial Advising**
  - Crawford & Sobel (1982); Bolton, Freixas & Shapiro (2007)
  - Gennaioli, Shleifer & Vishny (2015)
- **Soft Information and FinTech**
  - Liberti & Petersen (2019); He (2024)
- **LLM Simulations**
  - Horton (2023); Ouyang, Yun & Zheng (2024)

Back

Key Notation

Symbol	Meaning
$\omega \in \{0, 1\}$	True (latent) preference type
$p = \mathbb{P}^i(\omega = 1)$	Investor's belief about her preference
$\hat{p} = \mathbb{P}^L(\omega = 1)$	AI's belief about investor's preference
$p_0, \hat{p}_0$	Prior beliefs (before communication)
σ	Noise in signals (prompt quality $= 1/\sigma$)
$\kappa \in [0, 1]$	AI memory parameter
λ	Poisson shock intensity (impatience)
c	Flow cost of communication
$\Delta_\mu \equiv \mu_1 - \mu_0$	Preference uncertainty magnitude

Back

Communication Process: Brownian Learning

- Information about ω revealed through Brownian diffusion:

$$ds_t = \omega \, dt + \sigma \, dB_t$$

- **Investor's Belief Update** (via Bayes + Itô):

$$\boxed{dp_t = \frac{p_t(1 - p_t)}{\sigma} dB_t}$$

- Key properties:
  - p_t is a **martingale**
  - Absorbing at $p = 0$ or $p = 1$
  - **Prompt Quality** $= 1/\sigma$

Back

AI Memory (κ Model)

- AI observes signal z_t but only retains fraction κ:

$$\boxed{d\hat{z}_t = \kappa\, dz_t}$$

AI Memory (κ Model)

- AI observes signal z_t but only retains fraction κ:

$$\boxed{d\hat{z}_t = \kappa\, dz_t}$$

- **No Memory** ($\kappa = 0$):
  - AI ignores all signals: $\hat{p}_t = \hat{p}_0$

AI Memory (κ Model)

- AI observes signal z_t but only retains fraction κ:

$$\boxed{d\hat{z}_t = \kappa\, dz_t}$$

- **No Memory** ($\kappa = 0$):
  - AI ignores all signals: $\hat{p}_t = \hat{p}_0$
- **Partial Memory** ($\kappa \in (0, 1)$):
  - AI partially updates belief

AI Memory (κ Model)

- AI observes signal z_t but only retains fraction κ:

$$\boxed{d\hat{z}_t = \kappa \, dz_t}$$

- **No Memory** $(\kappa = 0)$:
  - AI ignores all signals: $\hat{p}_t = \hat{p}_0$
- **Partial Memory** $(\kappa \in (0, 1))$:
  - AI partially updates belief
- **Full Memory** $(\kappa = 1)$ + aligned prior $(\hat{p}_0 = p_0)$:
  - $\hat{p}_t = p_t$ (AI tracks investor perfectly)

Back

The Optimal Stopping Problem

- Investor solves:

$$\sup_{\tau \geq 0} \mathbb{E}^{\omega} \left\{ \int_0^{\tau} e^{-\lambda t} [-c + \lambda g(p_t, \hat{p}_t)] dt + e^{-\lambda \tau} g(p_\tau, \hat{p}_\tau) \right\}$$

- **Interpretation**:
  - Choose when to stop communicating
  - Flow cost c; Poisson shock λ may force early stop
- **HJB Equation**:

$$\boxed{-c + \lambda(g(p) - V(p)) + \frac{p^2(1-p)^2}{2\sigma^2} V''(p) = 0}$$

Back

Optimal AI Training

- Optimal AI prior $\hat{p}^*$:

$$\hat{p}^*(p_0) \begin{cases} > p_0 & \text{if } p_0 > 0.5 \\ < p_0 & \text{if } p_0 < 0.5 \\ = p_0 & \text{if } p_0 = 0.5 \end{cases}$$

- **Two Effects**:
  - Alignment: favors $\hat{p} = p$
  - Informativeness: favors extreme $\hat{p}$

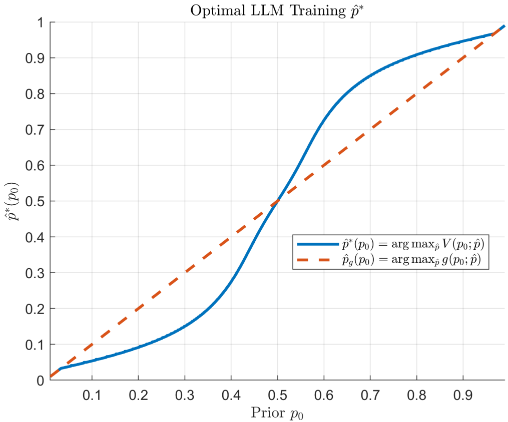

Back

Simulation Design (Full)

- **Data Source**: Survey of Consumer Finances (SCF) 2022
  - $n = 500$ hypothetical investor profiles
  - Ground truth: Vanguard 11-question Investor Questionnaire
- **LLM Implementation**: OpenAI GPT-5
  - Two-sided simulation: Both advisor and investor are LLM-generated
  - Multi-turn, role-structured conversations
- **Three Advisor Conditions**:
  - No Memory: Only most recent Q&A pair
  - Full Memory: Complete chat history
  - Full Information: Given complete profile text
- **Termination**: Poisson shock ($p = 0.10$ per round), max 11 rounds

Back

H1: Regression Results

	Accuracy$_i$			
	(1)	(2)	(3)	(4)
Prior (before any interaction)	69.0			
Final Before Rec	83.4			
Final (after rec)	86.6			
Number of Rounds (interim)		0.73***		
		(0.08)		
Number of Rounds (final)			0.77***	
			(0.12)	
Total Words (final)				0.012***
				(0.002)

Back

H4: Methodology

- **Weighted Pool Analysis**:
  - Partition: $\omega = 0$ (Vanguard $\leq 25\%$) vs $\omega = 1$ (Vanguard $\geq 75\%$)
  - Subsample: 50 profiles, 30,250 observations
- **Expected Payoff**:

$$\mathbb{E}[\text{Payoff}|p_0, p'] = p_0 \cdot \mathbb{E}[\text{Payoff}|\omega = 1, p'] + (1 - p_0) \cdot \mathbb{E}[\text{Payoff}|\omega = 0, p']$$

- **Result**: Empirical $\hat{p}^*(p_0)$ matches theory

Back

Conversation Dynamics

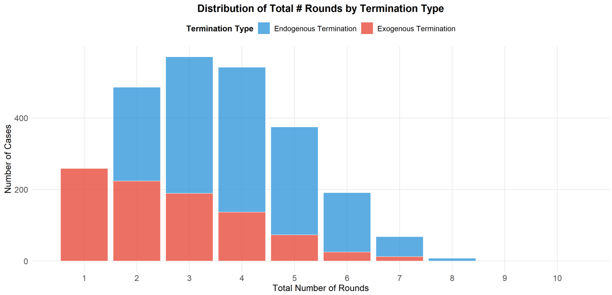

Distribution of conversation lengths and termination reasons

Back to H1

Comparative Statics: Communication Cost

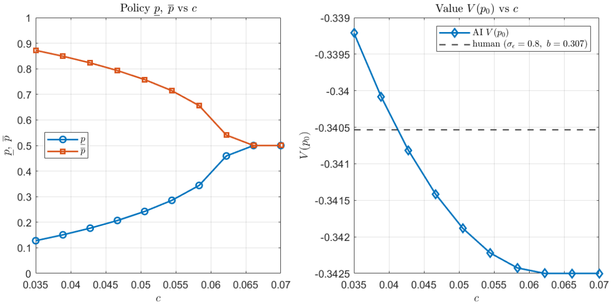

Higher $c \Rightarrow$ narrower continuation region

Comparative Statics: Preference Uncertainty

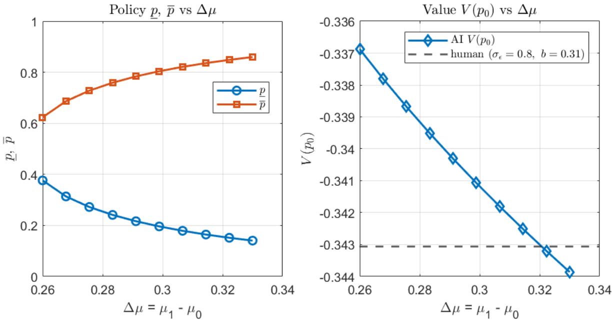

Higher $\Delta_\mu \Rightarrow$ wider continuation region

Comparative Statics: LLM Prior

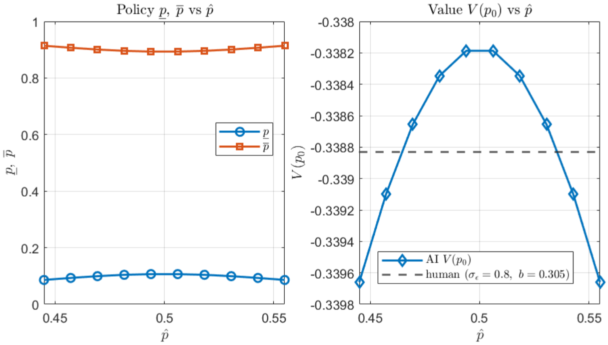

More extreme $\hat{p}_0 \Rightarrow$ more informative

Comparative Statics: Fundamental Uncertainty

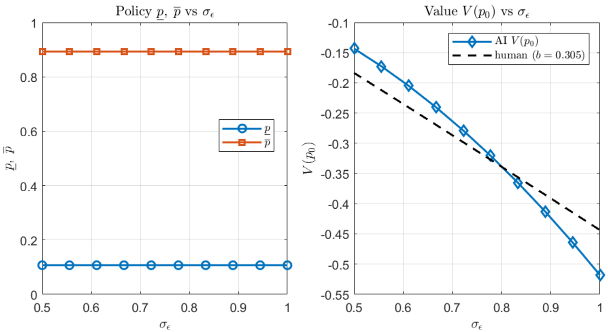

Thresholds invariant to σ_ϵ^2

LLM Timeline

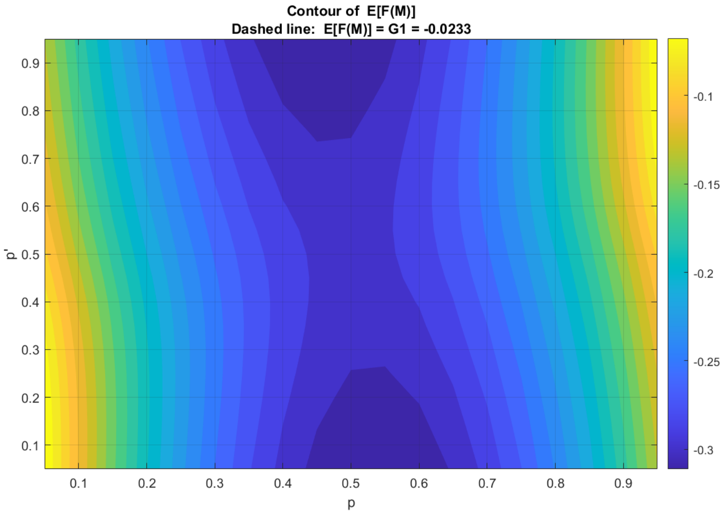

Policy Extension