Getting the Agent to Wait

Maryam Saeedi Yikang Shen Ali Shourideh

Carnegie Mellon University

Women In Economic Theory, LSE

Engagement as Objective

- Several environments:
  - Expert advice: legal and consulting services
  - Social Media: main source of revenue is advertising
  - Recommender Systems: TikTok, YouTube, Google News

Engagement as Objective

- Several environments:
  - Expert advice: legal and consulting services
  - Social Media: main source of revenue is advertising
  - Recommender Systems: TikTok, YouTube, Google News
  - This talk!!!

Engagement as Objective

- Several environments:
  - Expert advice: legal and consulting services
  - Social Media: main source of revenue is advertising
  - Recommender Systems: TikTok, YouTube, Google News
  - This talk!!!
- Incentives of the platform/expert (Principal) and the user (Agent) are *misaligned*:
  - Principal: maximize engagement (time on platform / billable)
  - Agent: acquire accurate information *as quickly as possible*

Engagement as Objective

- Several environments:
  - Expert advice: legal and consulting services
  - Social Media: main source of revenue is advertising
  - Recommender Systems: TikTok, YouTube, Google News
  - This talk!!!
- Incentives of the platform/expert (Principal) and the user (Agent) are *misaligned*:
  - Principal: maximize engagement (time on platform / billable)
  - Agent: acquire accurate information *as quickly as possible*
- **Central question:** how does the principal optimally release information over time, and what does that imply for what users end up believing?

This Paper

- A dynamic information-design framework where the principal chooses *when* and *what* to reveal, allowing:
  - **Belief disagreement** (agree-to-disagree, à la Aumann 1976)
  - **Private agent beliefs** – non-personalized vs. personalized communication
- Two design questions:
  1. How should an expert release information to a user whose prior differs from his own?
  2. Does personalization (one feed per user) improve welfare relative to mass communication (one feed for everyone)?
- **Why we should care:** personalized news feeds are routinely blamed for polarization and *filter bubbles*.

Preview of Results

- **Catering to the bias** (disagreeing priors):
  - Principal first reveals the state the agent **prefers**
  - Principal first reveals the state the agent already thinks is **likely**
  - Engagement converges to a *steady-state belief frontier*

Preview of Results

- **Catering to the bias** (disagreeing priors):
  - Principal first reveals the state the agent **prefers**
  - Principal first reveals the state the agent already thinks is **likely**
  - Engagement converges to a *steady-state belief frontier*
- **Speed vs. quality** (private beliefs):
- If types differ too much:
  - Non-personalized: information arrives *faster* but agents exit *less informed*
  - Personalized: slower exit, but beliefs are closer to the truth

The Model

The Model

- Time is continuous
- Agent trying to take the most accurate action $a \in A$ as soon as possible

$$u_A(\omega, a, T) = \delta_A e^{-\delta_A T} \hat{u}(\omega, a)$$

- Underlying state: $\omega \in \Omega = \{0, ..., N\}$
  - In most of the presentation: $N = 2$
- Time spent acquiring information: T
- Principal's payoff :
-

$$\int_0^T e^{-\delta_P t} dt = \frac{1 - e^{-\delta_P T}}{\delta_P}$$

Possibility of (belief) Disagreement

- Priors

$$\mu_A = \mathbb{P}_A(\omega = 1), \mu_P = \mathbb{P}_P(\omega = 1) \in (0,1)$$

- Learning is Bayesian
- Start from $\mu_{A,0}$ and $\mu_{P,0}$ being public information:
  - Later allow for $\mu_{A,0}$ being private.

The Model: Timing

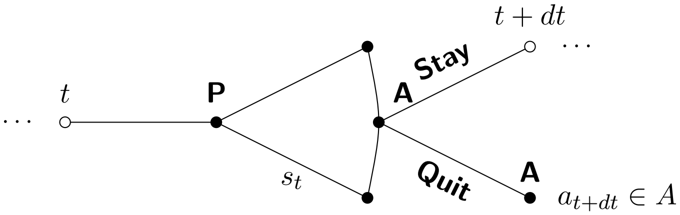

- P: commitment type
- A: no commitment

The Model

- P chooses an information structure.
- A mapping from the space of history realizations to probability distributions over signals at t.

$$\left(S_\infty \times \Omega, \mathcal{F}, \mathbb{P}_P, \left\{\mathcal{F}_{t,P}\right\}_{t\in\mathbb{R}_+}\right)$$

  - S_∞: the set of history of signal realizations,
  - Each member is of the form s^∞, $\mathcal{F}$ is a σ-algebra over $S_\infty \times \Omega$,
  - $\mathbb{P}_P$: probability measure from the principal's perspective
  - $\mathcal{F}_{P,t} \subset \mathcal{F}_{P,t'} \subset \mathcal{F}, \forall t < t'$ is a filtration.
  - A's information is similar:

$$\mathbb{P}_A(S) = \sum_{j\in\Omega} \mu_{A,j} \cdot \mathbb{P}_P(S \times \Omega | \omega = j)$$

The Model

- We define the payoff function of the agent upon exiting:

$$v(\mu) = \max_{a \in A} \mathbb{E}_{\mu}\left[\hat{u}(a, \omega)\right]$$

Assumption. The Payoff function is strictly convex and differentiable.

- allows us to take derivatives

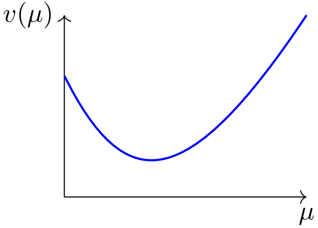

Leave with Perfect Info

Lemma. If A exits after history s^t, then $\mu_A(s^t) = \mathbb{E}_A\left[\omega = 1 | s^t\right] \in \{0, 1\}$ a.e.

- Idea of proof: If not, then split the signal into 2 fully revealing signals each with probability $\mu_A(s^t)$.
- Mean preserving spread: Increases the value of staying at all histories.
- Allows P to reduce the probability of exit and increase his payoff.
- Crucial Assumption: common knowledge about priors

Recommendation Mechanism

- By Carathéodory's theorem: 3 signals suffice per period
  - Combined with the Lemma (leave with perfect info): signals are $\Omega \cup \{\text{No News}\}$
  - Each period: agent either exits with **full information** about ω, or stays and updates beliefs from the absence of revelation

Recommendation Mechanism

- By Carathéodory's theorem: 3 signals suffice per period
  - Combined with the Lemma (leave with perfect info): signals are $\Omega \cup \{\text{No News}\}$
  - Each period: agent either exits with **full information** about ω, or stays and updates beliefs from the absence of revelation
- Recommendation mechanism: Myerson ('82, '86)
  - P recommends exit accompanied by simultaneous revelation of ω
  - **Conditional** D.D.F's (decumulative distribution functions):

$$G_\omega(t) = \Pr(\text{exit} \geq t \mid \omega), \quad \forall \omega \in \Omega, \qquad G_\omega(0) = 1$$

  - $G_\omega(t)$: non-increasing; same for both P and A conditional on ω

Beliefs and Likelihood Ratio

- Principal's unconditional survival and posterior:

$$G_P(t) = \mu_{P,0} G_1(t) + \left(1 - \mu_{P,0}\right) G_0(t), \qquad \mu_P(t) = \frac{\mu_{P,0} G_1(t)}{G_P(t)}$$

- Agent's unconditional survival:

$$G_A(t) = \mu_{A,0} G_1(t) + \left(1 - \mu_{A,0}\right) G_0(t), \qquad \mu_A(t) = \frac{\mu_{A,0} G_1(t)}{G_A(t)}$$

- Relative likelihood ratio (constant over time):

$$\ell = \frac{\mu_{A,0}}{1 - \mu_{A,0}} \frac{1 - \mu_{P,0}}{\mu_{P,0}}, \qquad \mu_A(t) = \frac{\ell \mu_P(t)}{\ell \mu_P(t) + 1 - \mu_P(t)}$$

Optimal Information Provision

$$\max_{G_0, G_1} \int_0^\infty \frac{e^{-\delta_P t} - 1}{\delta_P} \, dG_P(t)$$

subject to

$$-\int_t^\infty e^{-\delta_A (s-t)} \left[\mu_{A,0} \, v(1) \, dG_1(s) + \left(1 - \mu_{A,0}\right) v(0) \, dG_0(s) \right] \geq G_A(t) \, v(\mu_A(t)), \ \forall t$$

$$G_\omega(t) \ : \ \text{non-increasing}, \quad G_\omega(0) = 1, \quad \omega \in \{0, 1\}$$

Optimal Information Provision

$$\max_{G_0, G_1} \int_0^\infty \frac{e^{-\delta_P t} - 1}{\delta_P} \, dG_P(t)$$

subject to

$$-\int_t^\infty e^{-\delta_A (s-t)} \left[\mu_{A,0} \, v(1) \, dG_1(s) + \left(1 - \mu_{A,0}\right) v(0) \, dG_0(s)\right]$$
$$\geq G_A(t) \, v(\mu_A(t)), \; \forall t$$

$$G_\omega(t) \; : \; \text{non-increasing}, \quad G_\omega(0) = 1, \quad \omega \in \{0, 1\}$$

- Everything in red is written under the *agent's* prior
- Constant likelihood ratio $\ell \Rightarrow$ a change of variable absorbs the disagreement

Reformulation: Disagreement → Planner's Prior

- Define the *disagreement-adjusted* agent payoff:

$$\tilde{v}(\mu_P) \equiv [\ell\mu_P + 1 - \mu_P]\, v\left(\frac{\ell\mu_P}{\ell\mu_P + 1 - \mu_P}\right), \qquad \tilde{v}(1) = \ell\, v(1), \;\; \tilde{v}(0) = v(0)$$

$$\max_{G_0, G_1} \int_0^\infty \frac{e^{-\delta_P t} - 1}{\delta_P}\, dG_P(t)$$

subject to

$$-\int_t^\infty e^{-\delta_A (s-t)} \left[\mu_{P,0}\, \tilde{v}(1)\, dG_1(s) + \left(1 - \mu_{P,0}\right) \tilde{v}(0)\, dG_0(s)\right] \geq G_P(t)\, \tilde{v}(\mu_P(t)), \; \forall t$$

- Same structure as the common-prior problem – but with $\tilde{v}$ in place of v
- Disagreement is absorbed into the curvature of $\tilde{v}$ via ℓ

Properties of Optimal Communication

More Patient Principal

Assumption. The principal is more patient than the agent, i.e., $0 \leq \delta_P < \delta_A$.

- Key implication: information is revealed at a slow rate via a state-dependent jump process
  - Exit is rare: $\Pr(\text{exit in } (t, t+dt)) \to 0$ as $dt \to 0$
  - A is kept always **indifferent** between staying and exiting

More Patient Principal

Assumption. The principal is more patient than the agent, i.e., $0 \leq \delta_P < \delta_A$.

- Key implication: information is revealed at a slow rate via a state-dependent jump process
  - Exit is rare: $\Pr(\text{exit in } (t, t+dt)) \to 0$ as $dt \to 0$
  - A is kept always **indifferent** between staying and exiting
- Why? Agent's payoff $e^{-\delta_A T}$ as a function of exit time T is **more convex** than the principal's $\frac{1-e^{-\delta_P T}}{\delta_P}$

Revelation Rates and Belief Dynamics

- Belief upon staying, $\mu(t)$, evolves differentially over time
- Let $\frac{G'_\omega(t)}{G_\omega(t)} = -g_\omega(t) \leq 0$ be the rate at which state ω is revealed. Then:

$$\mu'(t) = (g_0(t) - g_1(t))\,\mu(t)\,(1 - \mu(t))$$

Revelation Rates and Belief Dynamics

- Belief upon staying, $\mu(t)$, evolves differentially over time
- Let $\frac{G'_\omega(t)}{G_\omega(t)} = -g_\omega(t) \leq 0$ be the rate at which state ω is revealed. Then:

$$\mu'(t) = (g_0(t) - g_1(t))\,\mu(t)\,(1 - \mu(t))$$

- **No news is news**: the *absence* of revelation is itself informative
  - If $g_0 > g_1$: state 0 is revealed faster
  - $\Rightarrow$ surviving without news $\Rightarrow$ more confident state is 1 $\Rightarrow$ μ drifts up

Binding IC: Flow Equation

- Since A is kept indifferent at all times, staying and exiting yield the same payoff
- Differentiating the IC ⇒ the **flow value** of staying must equal the flow value of the exit option:

$$\delta_A v(\mu(t)) = \underbrace{\mu(t) g_1(t)[v(1) - v(\mu(t))] + (1 - \mu(t)) g_0(t)[v(0) - v(\mu(t))]}_{\text{change in payoff when state is revealed}}$$
$$+ \underbrace{\mu(t)(1 - \mu(t))(g_0(t) - g_1(t)) v'(\mu(t))}_{\text{change from smooth belief drift (no news)}}$$

- Interpretation:
  - LHS: flow cost of staying = value of exit $v(\mu(t))$
  - RHS: expected change in continuation value from revelation + drift
    - Jump terms: state revealed ⇒ payoff jumps from $v(\mu)$ to $v(1)$ or $v(0)$
    - Drift term: no news ⇒ beliefs shift smoothly ⇒ exit option changes by $v'(\mu)\mu'$

HJB Equation

-

$$\underbrace{\delta_P V(\mu)}_{\text{required return}} = \max_{g_0, g_1 \geq 0} \underbrace{1}_{\text{flow}} - \underbrace{\left[\mu g_1 + (1-\mu) g_0\right] V(\mu)}_{\text{loss from exit}} + \underbrace{\mu (1-\mu)(g_0 - g_1) V'(\mu)}_{\text{belief drift effect}}$$

subject to the agent's binding IC.

  - Three forces acting on V:
    - Flow income: earns 1 per unit time while A stays
    - Termination risk: if state revealed, A exits, $V \to 0$
    - Belief drift: no news shifts μ smoothly, changes V

One State at a Time

- HJB is **linear** in $(g_0, g_1) \Rightarrow$ optimum is at a **corner**
  - Either reveal only $\omega = 1$, or only $\omega = 0$ (not both at once)

One State at a Time

- HJB is **linear** in $(g_0, g_1) \Rightarrow$ optimum is at a **corner**
  - Either reveal only $\omega = 1$, or only $\omega = 0$ (not both at once)
- P compares the **cost-per-unit-of-IC-benefit** of each state:
  - Reveal whichever state is **cheaper** to reveal

One State at a Time

- HJB is **linear** in $(g_0, g_1) \Rightarrow$ optimum is at a **corner**
  - Either reveal only $\omega = 1$, or only $\omega = 0$ (not both at once)
- P compares the **cost-per-unit-of-IC-benefit** of each state:
  - Reveal whichever state is **cheaper** to reveal
- "**Cater to the bias**": cheaper state is typically the one A is already leaning toward
  - A's exit option $v(\mu)$ is already close to $v(\mathbf{1}_\omega)$
  - Confirming what A already believes costs less continuation value

Steady State

- Belief $\mu(t)$ converges to a **steady state** μ^*:
  - At μ^*, both states revealed at the **same rate** $\Rightarrow$ belief drift vanishes
  - P switches between revealing $\omega = 0$ and $\omega = 1$ exactly at μ^*

Steady State

- Belief $\mu(t)$ converges to a **steady state** μ^*:
  - At μ^*, both states revealed at the **same rate** $\Rightarrow$ belief drift vanishes
  - P switches between revealing $\omega = 0$ and $\omega = 1$ exactly at μ^*
- Envelope theorem applied to HJB gives the **geometric** condition:

$$v'(\mu^*) = -\frac{\delta_P [v(1) - v(0)]}{\delta_A - \delta_P}$$

  - Tangent to agent's value at μ^* has slope equal to this critical value

Steady State: Catering to the Bias

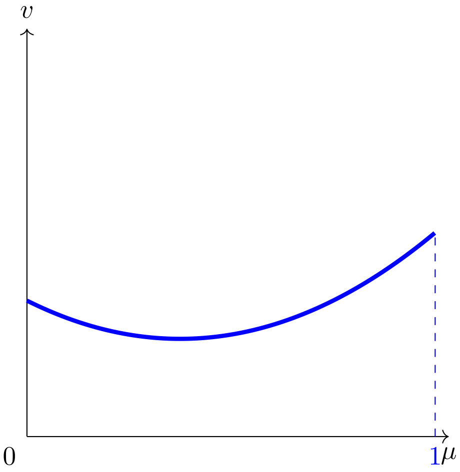

Steady State: Catering to the Bias

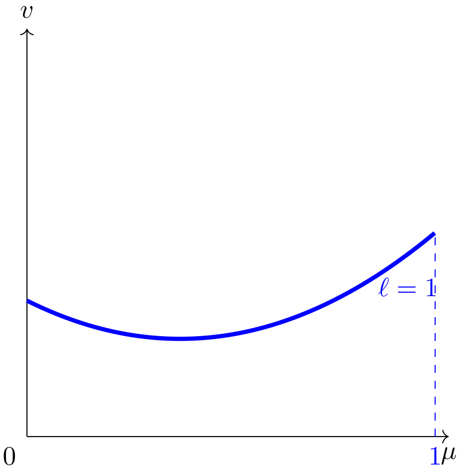

Steady State: Catering to the Bias

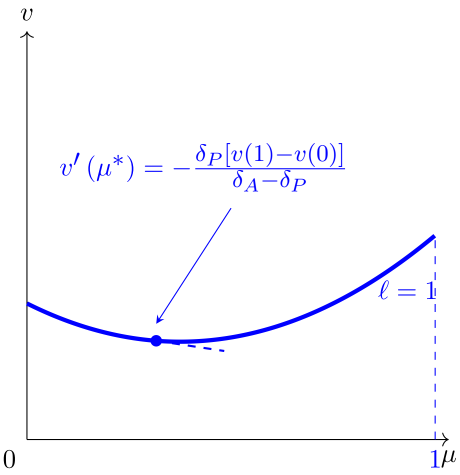

Steady State: Catering to the Bias

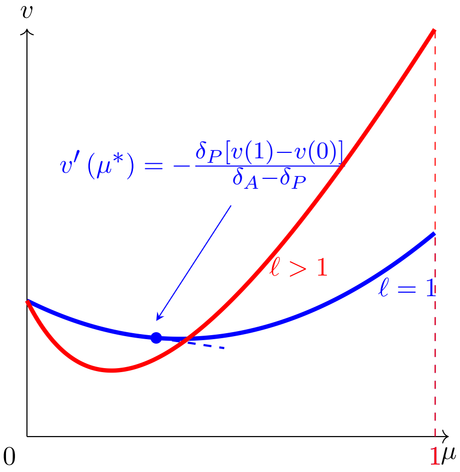

Steady State: Catering to the Bias

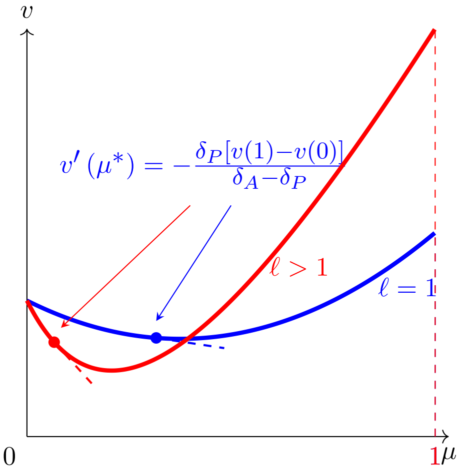

More Than Two States

Finite Number of States

- Finite number of states: $\Omega = \{0, 1, \cdots, N-1\}$
- HJB:

$$\delta_P V(\mu) = \max_{g_j \geq 0} 1 - \sum_j g_j \mu_j \left[V(\mu) + \frac{\partial V}{\partial \mu_j} - \nabla V \cdot \mu \right]$$

subject to

$$\delta_A v(\mu) = \sum_j \mu_j g_j \left[v\left(\mathbf{1}_j\right) - v(\mu) - \frac{\partial v}{\partial \mu_j} + \mu \cdot \nabla v \right]$$

- Implications:
  - Local linearity: almost always one state is revealed
  - Hierarchical revelation: See example with $|\Omega| = 3$.

Example with 3 states

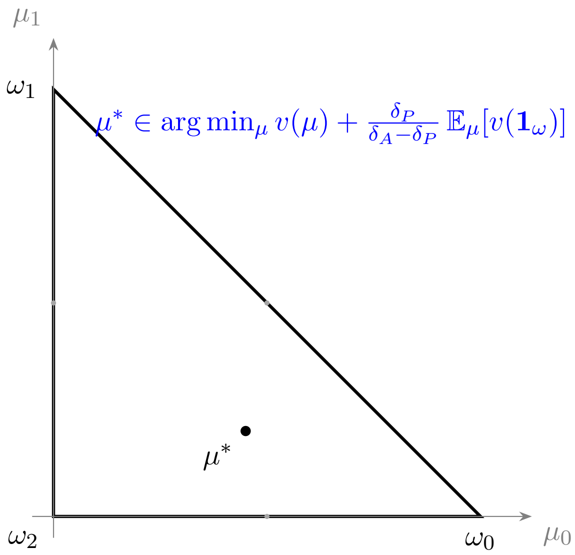

Example with 3 states

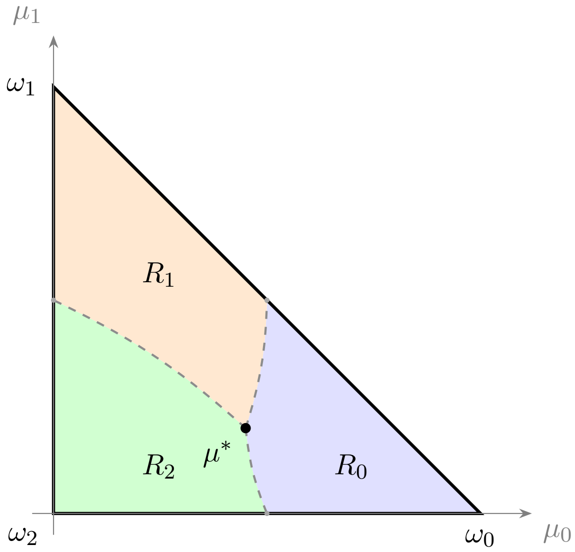

Example with 3 states

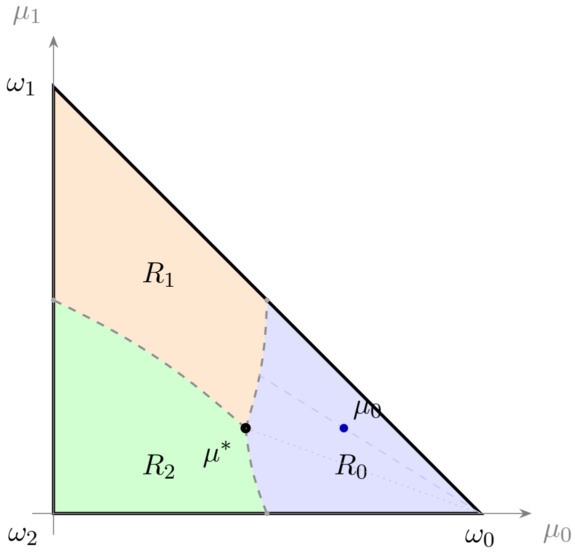

Example with 3 states

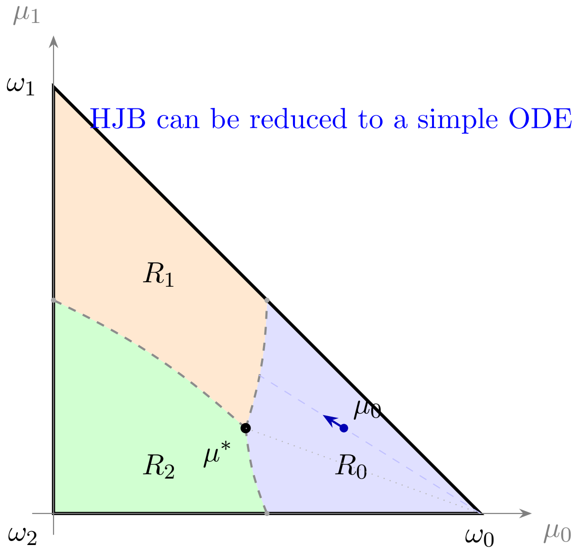

Example with 3 states

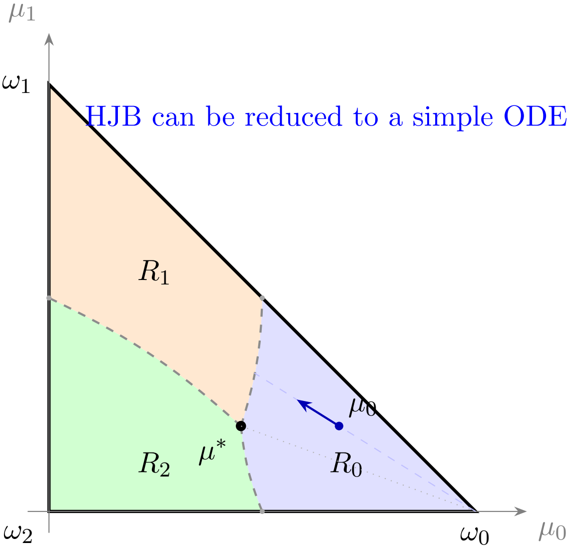

Example with 3 states

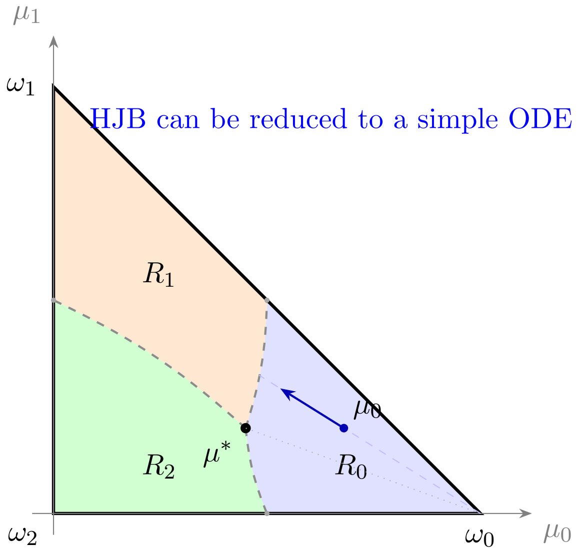

Example with 3 states

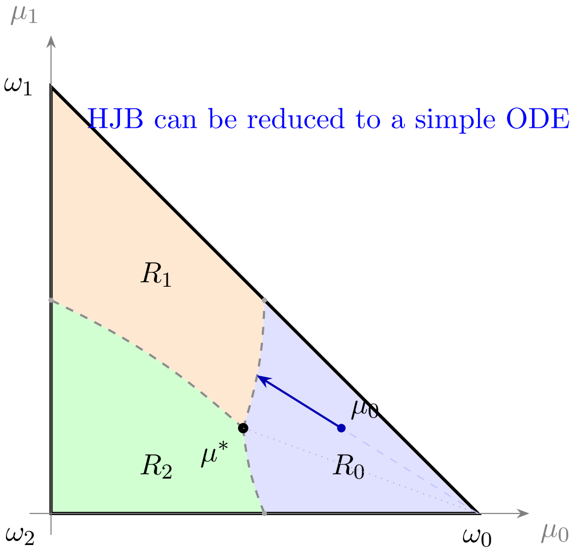

Example with 3 states

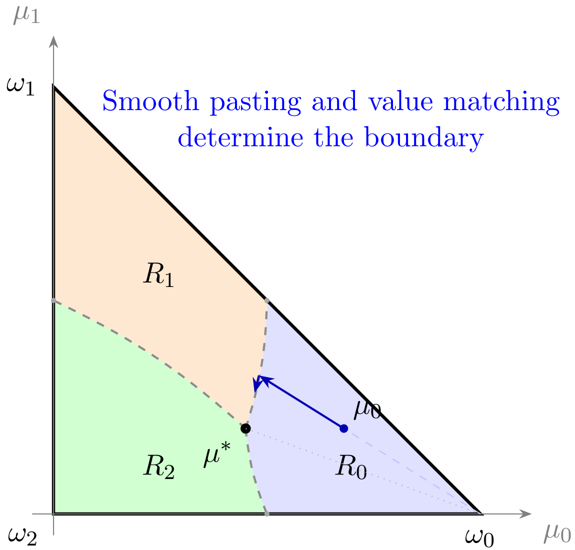

Example with 3 states

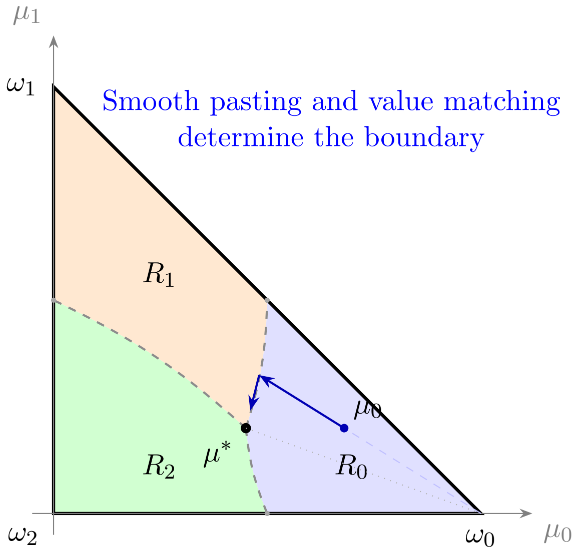

Example with 3 states

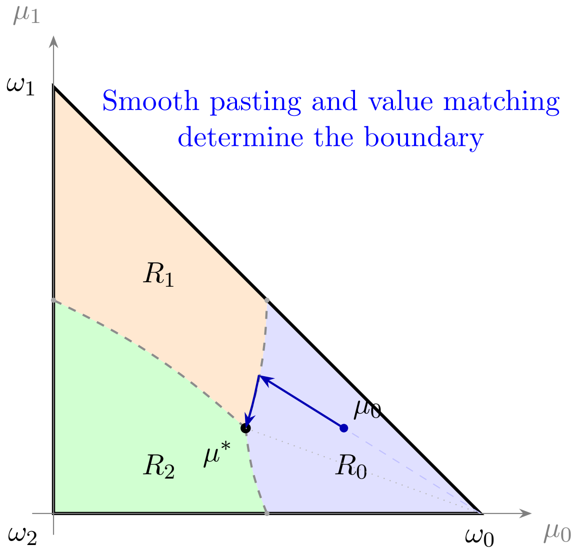

Non-Personalized Communication

Non-Personalized Communication

- How does communication change when μ_A is not observed?
- A's belief is either $\mu^L_{A,0}$ or $\mu^H_{A,0}$, with $\mu^L_{A,0} < \mu^H_{A,0}$, while the principal's belief is $\mu_{P,0}$, $\alpha_\theta = \Pr\left(\mu^\theta_{A,0}\right)$.
- P chooses $\left(S_\infty, \mathcal{F}, \mathbb{P}_P, \left\{\mathcal{F}_{t,P}\right\}_{t\geq 0}\right)$ but cannot control who listens and who exits

Non-Personalized Communication

Lemma. The best equilibrium of the game from the principal's perspective can be described by a communication policy together with a recommendation strategy for each type such that:

1. If type θ is recommended to quit following signal history s^t, the value of staying engaged for θ is not higher than $v\left(\mu_A^\theta\left(s^t\right)\right)$,
2. If type θ is recommended to stay following signal history s^t, the value of staying engaged for θ is not lower than $v\left(\mu_A^\theta\left(s^t\right)\right)$,

where $\mu_A^\theta\left(s^t\right)$ is the agent of type θ's belief induced by the communication policy.

Two Phases of Communication

1. *Full Engagement Phase* (Phase 1): Both types are engaged until a transition signal arrives.
2. *Partial Engagement Phase* (Phase 2): Transition to phase 2 happens when it is recommended that only one type stays. With one type engaged, we revert to the personalized case.

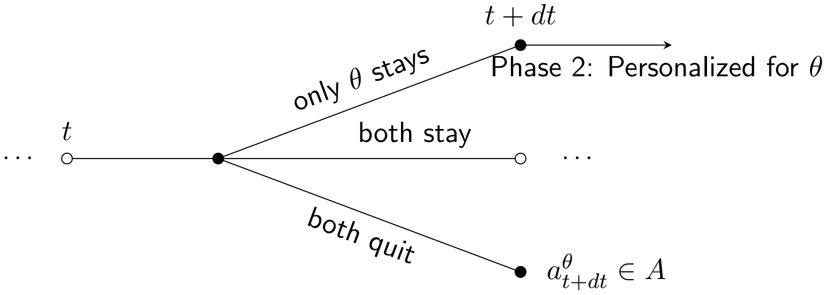

Resulting HJB

HJB

$$\delta_P V(\mu) = \max_{g_\omega, w_\theta, \mu_E} 1 + \mathbb{E}_\mu g_\omega \times (W - V(\mu)) + (g_0 - g_1)\,\mu\,(1-\mu)\,V'(\mu)$$

subject to

$$\mu_E = \frac{\mu g_1}{\mu g_1 + (1-\mu)\,g_0}, \text{IC}$$

Resulting HJB

HJB

$$\delta_P V(\mu) = \max_{g_\omega, w_\theta, \mu_E} 1 + \mathbb{E}_\mu g_\omega \times (W - V(\mu)) + (g_0 - g_1)\mu(1-\mu)V'(\mu)$$

subject to

$$\mu_E = \frac{\mu g_1}{\mu g_1 + (1-\mu) g_0}, \text{IC}$$

Exit Problem

$$W(w_H, w_L, \mu_E) = \max_{\tau \in \Delta\Delta\Omega} \sum_{\theta=H,L} \tau_\theta \alpha^\theta V_\theta(\mu_\theta)$$

subject to

$$\sum_{\theta=H,L} \tau_\theta \mu_\theta + \tau_1 \times 1 + \tau_0 \times 0 = \mu_E$$

$$\mathbb{E}_\tau v_\theta(\mu) = w_\theta$$

No Catering to the Bias

Proposition. If $\mu^L_{A,0} < \mu_{P,0} < \mu^H_{A,0}$, then there exists $\underline{\Delta} > 0$ such that if $\mu^H_{A,0} - \mu^L_{A,0} < \underline{\Delta}$, then at the stationary μ^*_P, $\tau^*_1 + \tau^*_0 = 1$. Revelation rate λ^* is higher than the personalized revelation rates:

$$\lambda^* > \lambda^*_H, \lambda^*_L.$$

No Catering to the Bias

Proposition. If $\mu^L_{A,0} < \mu_{P,0} < \mu^H_{A,0}$, then there exists $\underline{\Delta} > 0$ such that if $\mu^H_{A,0} - \mu^L_{A,0} < \underline{\Delta}$, then at the stationary μ^*_P, $\tau^*_1 + \tau^*_0 = 1$. Revelation rate λ^* is higher than the personalized revelation rates:

$$\lambda^* > \lambda^*_H, \lambda^*_L.$$

- No personalization stage
- Personalization leads to less uncertainty when uninformed.
- Faster learning: personalization leads to higher belief polarization

Quality–Speed Trade-off

Proposition. If $\mu_{A,0}^L < \mu_{P,0} < \mu_{A,0}^H$, then there exists $\overline{\Delta} > 0$ such that if $\mu_{A,0}^H - \mu_{A,0}^L > \overline{\Delta}$, then at the stationary μ_P^*, $\tau_1^* + \tau_0^* = 0$.

Quality–Speed Trade-off

Proposition. If $\mu^L_{A,0} < \mu_{P,0} < \mu^H_{A,0}$, then there exists $\overline{\Delta} > 0$ such that if $\mu^H_{A,0} - \mu^L_{A,0} > \overline{\Delta}$, then at the stationary μ^*_P, $\tau^*_1 + \tau^*_0 = 0$.

- Quality–Speed trade-off:
  - Personalization leads to higher speed of delivery of information
  - But information is not fully revealing in the first stage

Related Literature

- Information economics and Bayesian Persuasion:
  - Rayo and Segal (2010), Kamenica and Gentzkow (2011) and many many many more!

Related Literature

- Information economics and Bayesian Persuasion:
  - Rayo and Segal (2010), Kamenica and Gentzkow (2011) and many many many more!
- Models of Dynamic Communication
  - Ely and Szydlowski (2020), Orlov, Skrzypacz, Zryumov (2020), Che, Kim and Meierendorf (2022), Hebert and Zhong (2022): difference in payoffs and information revelation policies
  - Koh and Sanguanmoo (2024), Koh, Sanguanmoo and Zhong (2024)

Conclusion

- Developed a dynamic model of information provision when the principal wants to maximize engagement
- Principal caters to the bias
  - State that the agent likes
  - State that the agent thinks is more likely
- More than two types, two phases
  - Keep both types
  - Cater to one type
- A lot more to be done:
  - Time Inconsistency: digital addiction
  - Competition
  - Optimal regulation without violating first amendment (in the U.S.)

Thank You for Staying Engaged!