Getting the Agent to Wait

Maryam Saeedi Yikang Shen Ali Shourideh

Michigan State University
April 2024

June 3, 2024

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!!!
- The incentives of the recommender system (principal) and users (agents) are not aligned
  - Principal: Maximize engagement; in order to maximize ad revenue
  - Agent: Acquire information, time cost

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!!!
- The incentives of the recommender system (principal) and users (agents) are not aligned
  - Principal: Maximize engagement; in order to maximize ad revenue
  - Agent: Acquire information, time cost
- Why do we care?
  - Personalized news aggregators: sometimes blamed for polarization in the media for amplifying *biases*

What we do

- This paper:
  - Develope a framework:
    - What are the key determinants of information flows?
    - How is information flow catered to the agents?

What we do

- This paper:
  - Develope a framework:
    - What are the key determinants of information flows?
    - How is information flow catered to the agents?
- The Model:
  - Principal: wants to give information as late as possible
  - Agent: Wants to learn as soon as possible! Time cost (variety of cases)
  - A and P: Bayesian; possibly different prior
  - P can commit but A cannot

Overview of Results

- Key determinant of information flow:
  - ***Marginal Cost of Engagement (MCE)*** for A

$$\text{MCE} = \frac{-\partial u_A/\partial T}{\partial u_P/\partial T}, T : \text{end of engagement}$$

Overview of Results

- Key determinant of information flow:
  - ***Marginal Cost of Engagement (MCE)*** for A

$$\text{MCE} = \frac{-\partial u_A/\partial T}{\partial u_P/\partial T}, T : \text{end of engagement}$$

  - The degree of the bias in the prior

Overview of Results

- Evolution of MCE over time determines how information is revealed:
  - If MCE decreases over time (e.g., P more patient than A): gradual (Poisson) revelation
  - If MCE increases over time (e.g., P less patient than A): abrupt revelation

Overview of Results

- Evolution of MCE over time determines how information is revealed:
  - If MCE decreases over time (e.g., P more patient than A): gradual (Poisson) revelation
  - If MCE increases over time (e.g., P less patient than A): abrupt revelation
- Disagreement in prior:
  - Cater to the bias: reveal the state that the agent think is more likely first
  - Some form of gradual revelation is optimal

Overview of Results

- Evolution of MCE over time determines how information is revealed:
  - If MCE decreases over time (e.g., P more patient than A): gradual (Poisson) revelation
  - If MCE increases over time (e.g., P less patient than A): abrupt revelation
- Disagreement in prior:
  - Cater to the bias: reveal the state that the agent think is more likely first
  - Some form of gradual revelation is optimal
- Compare Personalized and non-Personalized News
  - Trade-off between quality of information and timing.

Related Literature

- Basics of information economics:
  - Kamenica and Gentzkow (2011) and many many many more!
  - Information design with incentives: Boleslavsky and Kim (2022), Onuchic and Ray (2022), Saeedi and Shourideh (2023), Best, Quigley, Saeedi, Shourideh (2023)
- 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
    - 3S: New insights on the change of optimal disclosure
- Small literature on recommender systems in economics: Calvano, Calzolari, Denicolo, and Pastorello (2023): focus on effect on competition
- Lots of commentary on the issue:
  - Example: Acemoglu and Robinson: tax online advertisement; Our model: not so straightforward

Full Model

- As before time is continuous
- Agent utility function

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

- Underlying state: $\omega \in \Omega = \{0,1\}$
- Action: $a \in A$
- Time spent acquiring information: T
- Principal's payoff :

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

- Possibly uncommon priors $\mu_0^A = \mathbb{P}^A(\omega = 1), \mu_0^P = \mathbb{P}^P(\omega = 1) \in (0,1)$.
  - Common knowledge now; later private information for the agent

Timing

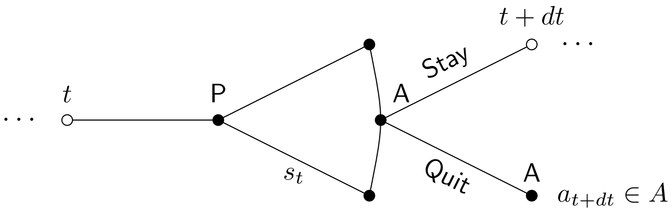

The Model

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

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

  - 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
  - $\mathcal{F}_t \subset \mathcal{F}_{t'} \subset \mathcal{F}, \forall t < t'$ is a filtration.

The Model

- A's information is similar except that it does not include Ω and

$$\mathbb{P}^A(S) = \mu_0^A \cdot \mathbb{P}^P(S \times \Omega | \omega = 1) + \left(1 - \mu_0^A\right) \cdot \mathbb{P}^P(S \times \Omega | \omega = 0)$$

  - $\mathcal{F}_t^A$ is similarly calculated
- Equilibrium is standard:
  - A cannot commit to exit strategies
  - P can commit to information structure

The Model – Characterization

Lemma. If A exits after history s_t, then $\mu_t^A = \mathbb{E}^A[\omega|s_t] = 0, 1$ a.e.

- Idea of proof: If not, then split the signal into two fully revealing signals each with probability μ_t^A and $1 - \mu_t^A$. Increases the value of staying at all histories. Allows P to reduce the probability of exit and increase his payoff.

The Model

Assumption. The Payoff function $v(\mu) = \max_{a \in A} \mathbb{E}_\mu [\hat{u}(a, \omega)]$ is strictly convex, differentiable and symmetric around $\mu = 1/2$.

- Allows us to take derivatives
- An example is $\hat{u}(a, \omega) = a(\omega - 1/2) - a^2/2$, $A = [-1, 1]$
- Does not include $|A| < \infty$, since $v(\mu)$ is piecewise linear
  - can approximate with smooth convex functions

The Model

- Can apply Caratheodory theorem
  - 3 signals in each period is sufficient: $\Omega \cup \{\text{No News}\}$
- Choice of information structure is equivalent to choice of two D.D.F functions (decumulutive distribution functions)

$$
\begin{aligned}
G_1(t) =& \mathbb{P}^A(\text{exit} \geq t, \omega = 1) \\
G_0(t) =& \mathbb{P}^A(\text{exit} \geq t, \omega = 0) \\
\hat{\mu}^A(t) =& \mathbb{P}^A(\omega | \text{stay until } t) \\
=& \frac{G_1(t)}{G_1(t) + G_0(t)} = \frac{G_1(t)}{G(t)}
\end{aligned}
$$

- D.D.F's are decreasing and $G_1(0) = \mu_0^A = 1 - G_0(0)$

Optimal Information Provision

$$\max_{G_0,G_1} \int_0^\infty e^{-\delta_P t} \left(G_{P,1}(t) + G_{P,0}(t)\right) dt$$

subject to

$$v(1)\, G_A(t) - v(1)\, \delta_A \int_t^\infty e^{-\delta_A (s-t)} G_A(s)\, ds \geq G(t)\, v\left(\hat{\mu}^A(t)\right), \forall t$$

$$G_\omega(t) : \text{ non-increasing}$$

$$G_1(0) = 1 - G_0(0) = \mu_0^A$$

- $\ell = \frac{\mu_0^A}{1-\mu_0^A} / \frac{\mu_0^P}{1-\mu_0^P}$: likelihood ratio; adjustment needed for difference in prior

Solution Method

- Objective is linear in $G_{\omega}(t)$
- Constraint set is convex and has a non-empty interior. We can use standard Lagrangian techniques
  - Guess a Lagrangian
  - Use first order condition
  - Use ironing when necessary
- Somewhat similar to Kleiner, Moldovanu, and Strack (2021) and Saeedi and Shourideh (2023)
  - key difference: it is not a linear program

Simple Example

- Restrict to extrme discosure policies: can only change the timing.
- Actions:
  - P: choose time $T \in \mathbb{R}_+ \cup \{0\}$ to reveal the state
  - A: chooses between quitting or staying at any time $t < T$ (no reason to stay after knowing the state)
- Payoffs:
  - P: T, i.e., he values engagement
  - A: $u(T) = e^{-\delta T} v(\text{Info})$, i.e., she values time not listening to the principal!!

$$v(\text{Info}) = \begin{cases} 1 & \text{Info} = \text{State} \\ 1/2 & \text{Info} = \text{Prior} \end{cases}$$

Simple Example

- Revelation strategy: reveal at $e^{-\delta T^*} = 1/2$

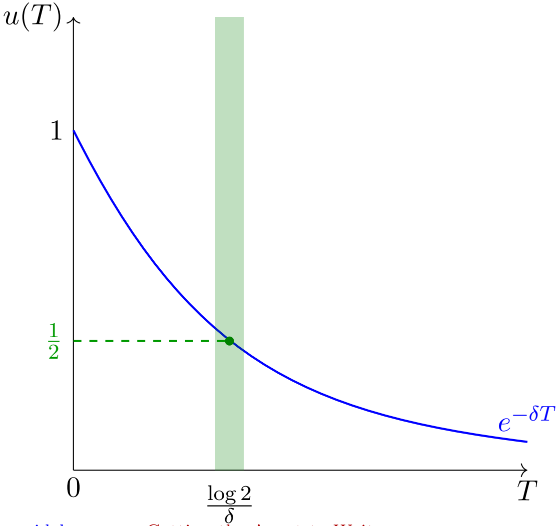

Simple Example

- Spread revelation time around T^*

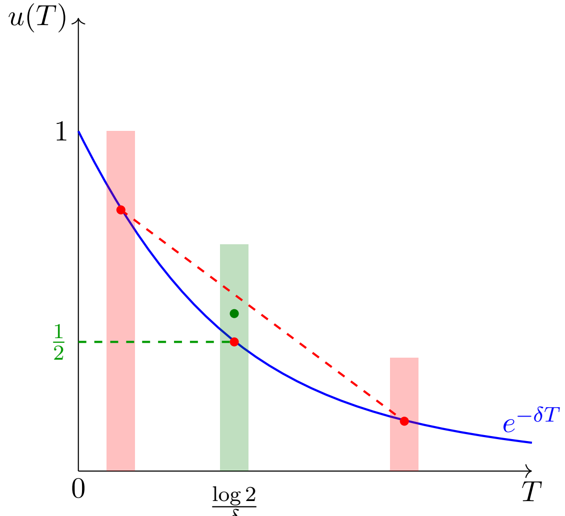

Simple Example

- Spread revelation time around T^* and increase its mean

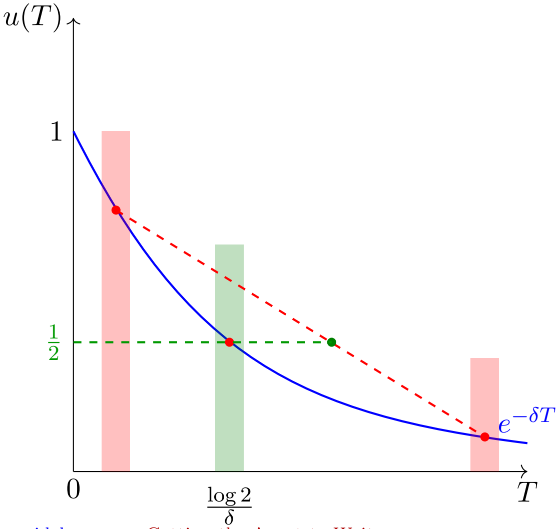

Simple Example

- Distribution: exponential at rate δ; Poisson revelation

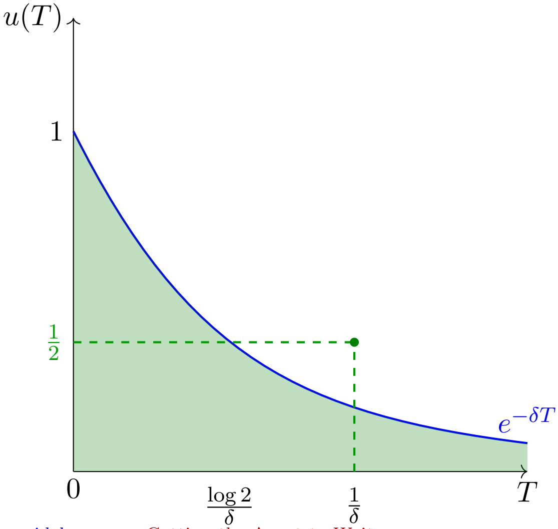

Simple Example

- Alternative: $u_P(T) = 1 - e^{-\delta_P T}$ with $\delta_P > \delta$; (A more patient)
- Rewrite:

$$u_A = (1 - u_P)^{\frac{\delta_P}{\delta}} \text{ : concave in } u_P$$

- In this case, a mean preserving contraction of any distribution of T (or u_P) benefits A
  - $\Rightarrow$ its mean can be pushed up!
- Optimal revelation strategy is T^*

$$e^{-\delta T^*} = 1/2 \to T^* = \frac{\log 2}{\delta}$$

Simple Example

- Concave payoff: Jensen's inequality: $\mathbb{E}[T] < 1$

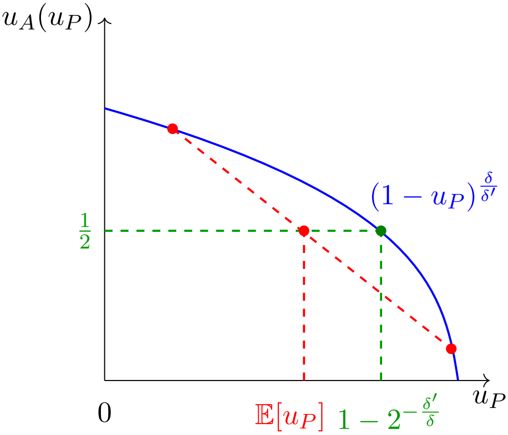

Summary of Example

- Relative concavity of the payoffs matter:
  - A convex relative to P: poisson revelation of information
  - A concave relative to P: abrupt revelation
- Example: quantity of information is fixed
  - Clearly can be varied by gradual slant, mixed messaging, etc.

The Agreement Case

- Suppose that $\mu_0^A = \mu_0^P \to \ell = 1$.
- First the easy one!

The Agreement Case

- Suppose that $\mu_0^A = \mu_0^P \to \ell = 1$.
- First the easy one!

Proposition. Impatient Principal. When $\delta_A < \delta_P$, optimal solution is

$$\begin{aligned} G_1(t) &= \mu_0 \mathbf{1}[t < t^*] \\ G_0(t) &= (1-\mu_0)\mathbf{1}[t < t^*] \\ v(1)D(t^*) &= v(\mu_0)D(0) \end{aligned}$$

The Agreement Case

- Suppose that $\mu_0^A = \mu_0^P \to \ell = 1$.
- First the easy one!

Proposition. Impatient Principal. When $\delta_A < \delta_P$, optimal solution is

$$G_1(t) = \mu_0 \mathbf{1}[t < t^*]$$
$$G_0(t) = (1 - \mu_0) \mathbf{1}[t < t^*]$$
$$v(1) D(t^*) = v(\mu_0) D(0)$$

- Silence until t^* is optimal!
- Agent is only indifferent at time 0 $\to$ Time inconsistency

The Agreement Case

Proposition. Patient Principal. When $\delta_P < \delta_A$, optimal solution has two phases (if $\mu_0 > 1/2$)

$$t \leq t^* : G_1'(t) < 0, \hat{\mu}'(t) < 0, G_0(t) = 1 - \mu_0$$

$$t \geq t^* : \hat{\mu}(t) = 1/2, \frac{G_0'(t)}{G_0(t)} = \frac{G_1'(t)}{G_1(t)} = \lambda^*$$

The case with $\mu_0 < 1/2$ is symmetric.

The Agreement Case

Proposition. Patient Principal. When $\delta_P < \delta_A$, optimal solution has two phases (if $\mu_0 > 1/2$)

$$t \leq t^* : G_1'(t) < 0, \hat{\mu}'(t) < 0, G_0(t) = 1 - \mu_0$$

$$t \geq t^* : \hat{\mu}(t) = 1/2, \frac{G_0'(t)}{G_0(t)} = \frac{G_1'(t)}{G_1(t)} = \lambda^*$$

The case with $\mu_0 < 1/2$ is symmetric.

- *Belief-Smoothing*
  - A's value function $v(\mu)$, i.e., cost of delay, is strictly convex
- Agent is always indifferent → Time consistency

Agreement: Patient Principal

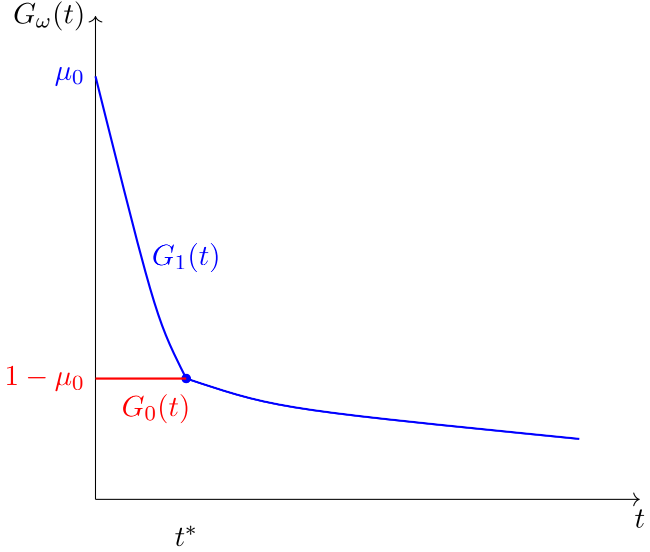

Agreement: Patient Principal

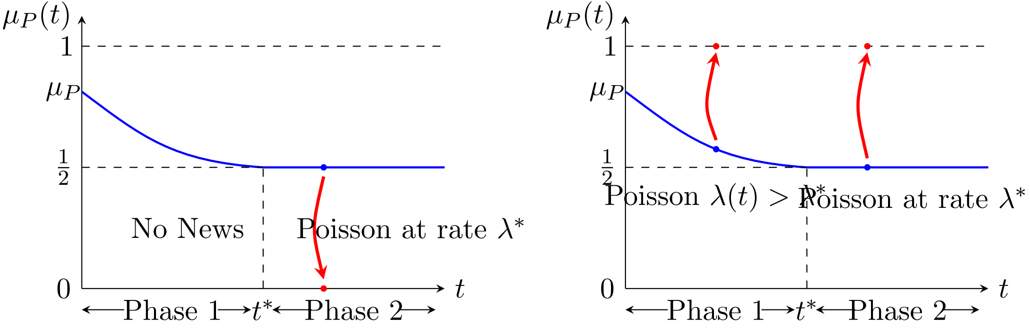

(a) Beliefs when $\omega = 0$ (b) Beliefs when $\omega = 1$

Figure: Evolution of beliefs in each state when the agent is more impatient and $\mu_P > 1/2$

Agreement: Patient Principal

- Two phases with time-varying Poisson revelation of information
  - Phase 1: Arrival of news about the more likely state at high rate
  - Phase 2: Arrival of news about both state at constant rate
- Phase 1 depends on the curvature of $v(\mu)$
  - The more convex it is, the longer is Phase 1
  - Belief-smoothing: Agent values smoothness of beliefs

Disagreement

- We are writing everyone's payoff as a function of beliefs of the principal.
- WLOG, let's say $\ell < 1$ so A is more optimistic about $\omega = 0$.
- Given that P prefers μ closer to 1, wants A to spend the most time strictly above $\hat{\mu} = 1/2$.

Disagreement: Impatient Principal

Proposition. Impatient Principal and Disagreement. Suppose $\delta_A < \delta_P$ and $\mu_A < \mu_P$, then optimal solution has two phases In phase 1, $t \in [0, t_1^*)$, no information is revealed.

1. At t_1^*, $\omega = 0$ is revealed with a positive probability.
2. In phase 2, $t \in [t_1^*, t_2^*]$, $\omega = 0$ is revealed gradually and according to a Poisson process at a rate so that the agent's beliefs satisfies the following ODE

$$\delta_A = \frac{\mu_A'(t)}{\mu_A(t)} \frac{v(1) - v(\mu_A(t)) + v'(\mu_A(t))\,\mu_A(t)}{v(\mu_A(t))}$$

3. At t_2^*, $\omega = 1$ is revealed so that $\mu_A(t_2^*) = 1$.

- Again two phases:
  - *Cater to the bias phase*: reveal the A-optimistic state
  - Settle on higher belief

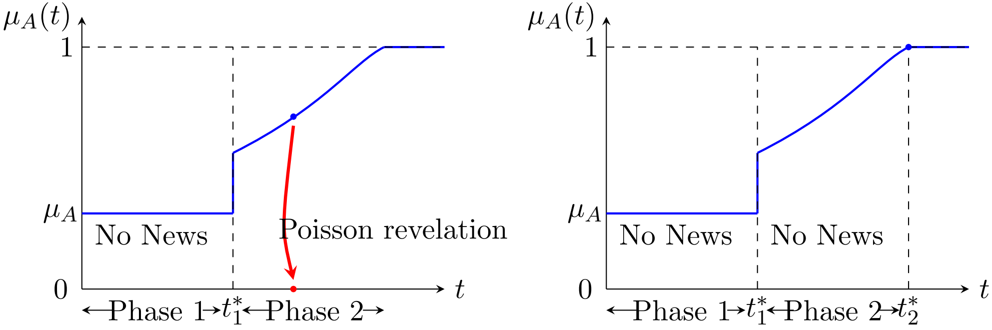

(a) Beliefs when $\omega = 0$

(b) Beliefs when $\omega = 1$

Figure: Catering to the bias with an Impatient Principal

Proposition. Patient Principal and Disagreement. Suppose that $\delta_A > \delta_P$. Then there exists a threshold $\mu_P^*(\mu_A)$ such that optimal communication consists of two phases:

1. If $\mu_P < \mu_P^*(\mu_A)$, in phase 1 only state $\omega = 0$ is gradually revealed so that the agent's beliefs satisfy

$$\delta_A = \frac{\mu_A'(t)}{\mu_A(t)} \frac{v(1) - v(\mu_A(t)) + v'(\mu_A(t))\,\mu_A(t)}{v(\mu_A(t))}$$

2. If $\mu_P^*(\mu_A) < \mu_P$, in phase 1 only state $\omega = 1$ is gradually revealed so that the agent's beliefs satisfy

$$-\delta_A = \frac{\mu_A'(t)}{1 - \mu_A(t)} \frac{v(1) - v(\mu_A(t)) - v'(\mu_A(t))\,(1 - \mu_A(t))}{v(\mu_A(t))}$$

3. In phase 2, when $\mu_P^*(\mu_A) = \mu_P$, both states are gradually revealed according

Patient Principal

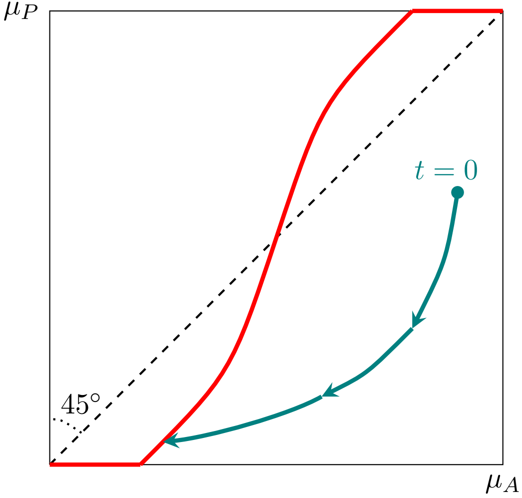

Application to Personalized News

- Personalized news is blamed for creating eco-chambers
- Possibly leading to polarization
- One way to think about it is via comparison of non-personalized and personalized benchmarks
- Agent's prior belief is private information: $\mu_A^H > \mu_A^L$
- Agent exits at rate ρ

Conclusion

- Developed a dynamic model of information provision when the principal wants to maximize engagement
- Relative curvature of principal and agent's payoffs determines revelation
- With biased beliefs: principal always initially caters to the bias
- Implications:
  - flat tax an advertisement might just not work
    - wont work in the patient case
  - Nonlinear taxes might
- A lot more to be done:
  - Time Inconsistency: digital addiction (Already showed that results dont change!)
  - Competition
  - Optimal regulation without violating first ammendment (in the U.S.)

Thank You for Staying Engaged