Catering to the Bias

Maryam Saeedi Yikang Shen Ali Shourideh

Carnegie Mellon University
EIEF, July 2023

July 25, 2023

Personalized Recommender Systems

- Ubiquitous nowadays
  - eCommerce recommendations: Amazon, Google Shopping, NYT Wirecutter
  - Social Media: Facebook, TikTok, Instagram, Youtube, Twitter
  - News Aggregators: Feedly, Google News, Panda, Techmeme, Flipboard, Youtube, Twitter

Personalized Recommender Systems

- Ubiquitous nowadays
  - eCommerce recommendations: Amazon, Google Shopping, NYT Wirecutter
  - Social Media: Facebook, TikTok, Instagram, Youtube, Twitter
  - News Aggregators: Feedly, Google News, Panda, Techmeme, Flipboard, Youtube, Twitter
- 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

What we do

- Personalized aggregators: sometimes blamed for polarization in the media for amplifying *biases*
- This paper: what are its theoretical implications?
  - Hopefully later: what can/should we do about it?
- 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

- Principal and agent share the same prior beliefs,
  - The relative curvature of A's payoff w.r.t time to that of principal determines optimal provision of news
  - (A relative to P) Convex in time: Poisson revelation with an intensity determined by the agent's discount rate
  - (A relative to P) Concave in time: A period of no information followed by an immediate revelation
  - Intermediate cases: habit formation and boredom!

Overview of Results

- Principal and agent share the same prior beliefs,
  - The relative curvature of A's payoff w.r.t time to that of principal determines optimal provision of news
  - (A relative to P) Convex in time: Poisson revelation with an intensity determined by the agent's discount rate
  - (A relative to P) Concave in time: A period of no information followed by an immediate revelation
  - Intermediate cases: habit formation and boredom!
- When the agent has a biased prior
  - P always caters to A's biased prior
    - early in the game reveals the state where the agent thinks is more likely
  - Some form of gradual revelation is always necessary – no abrupt revelation

Simple Example

- (Simple and very limited) Game between an informed principal (committed) and an uninformed agent (uncommitted)
- 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 $e^{-\delta T^*} = 1/2$

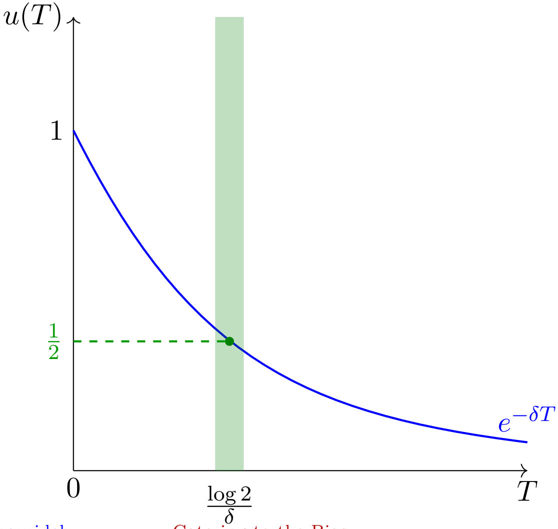

Simple Example

- Spread revelation time around T^*

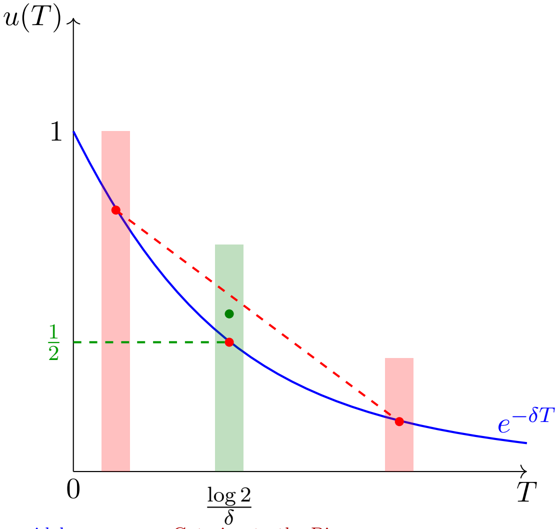

Simple Example

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

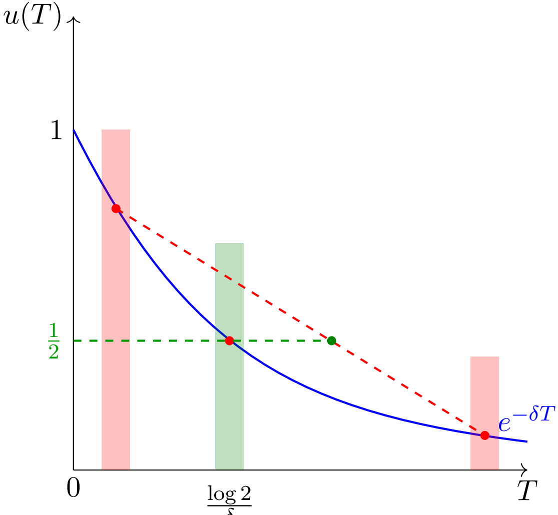

Simple Example

- Distribution: exponential at rate δ; Poisson revelation

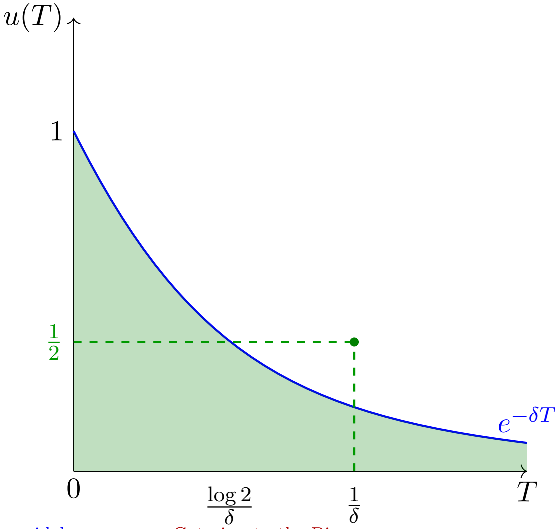

Simple Example

- Alternative: $u(T) = \left(1 - T^2/2\right) v(\text{info})$
- In this case, a mean preserving contraction of any distribution of T benefits A
  - $\Rightarrow$ its mean can be pushed up!
- Optimal revelation strategy is T^*

$$1 - (T^*)^2/2 = 1/2 \rightarrow T^* = 1$$

Simple Example

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

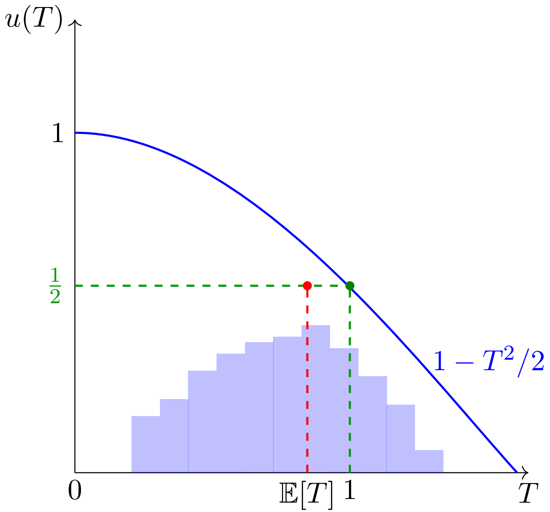

Summary of Example

- 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.

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

$$u_A(T, \omega, a) = D(T)\,\hat{u}(\omega, a)$$

- Underlying state: $\omega \in \Omega = \{0, 1\}$ – more would not make much of a difference
- Action: $a \in A$
- Time spent acquiring information: T
- $D(T)$ is strictly decreasing in T and $\hat{u}(\omega, a) \geq 0$
- Principal payoff : T
- Possibly uncommon priors $\mu_0^A = \mathbb{P}^A(\omega = 1), \mu_0^P = \mathbb{P}^P(\omega = 1) \in (0, 1)$. Common knowledge

Timing

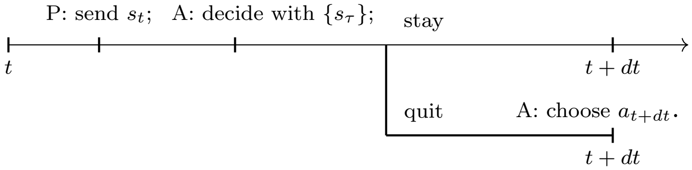

The Model

- 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

Some Examples

- Key assumption:

$$
\begin{aligned}
u^P &= T \\
u^A &= D\left(T\right)\hat{u}\left(\omega, a\right)
\end{aligned}
$$

- Useful to think about some examples to understand how to think about different shapes of $D\left(T\right)$.

Some Examples

Example 1. Standard Exponential Discounting: relative patience.

$$u^P = \int_0^{\hat{T}} e^{-\delta_P t} dt$$

$$\rightarrow T = u^P \Rightarrow D(T) = (1 - \delta_P T)^{\frac{\delta_A}{\delta_P}}$$

$$u^A = \delta_A e^{-\delta_A \hat{T}} \hat{u}(\cdot)$$

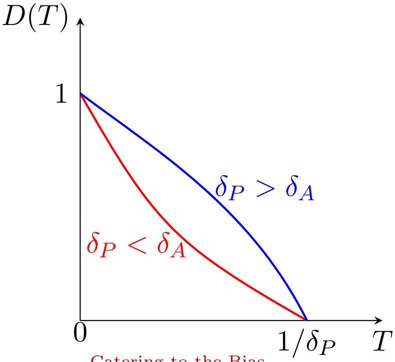

Some Examples

- **Example 2.** Gradually more patient agent $u^A = (1+\alpha T)^{-\beta}\,\hat{u}\,(\cdot)$; could be interpreted as habit formation
  - Set $T = u^P \Rightarrow D\,(T) = \left(1 - \frac{\alpha}{\delta_p}\log\left(1-\delta_p T\right)\right)^{-\beta}$

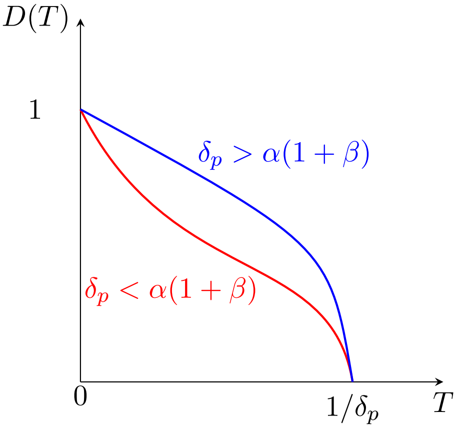

Some Examples

- **Example 3.** Gradually less patient agent $u^A = e^{\int_0^T g(\tau)d\tau} \hat{u}(\cdot), g' < 0, g'' > 0$; Boredom!

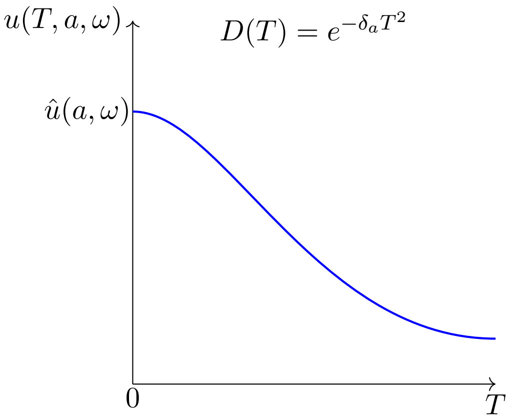

The Model – Characterization

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

$$\max_{G_0,G_1} \int_0^\infty \left(G_1(t) + \ell G_0(t)\right) dt$$

subject to

$$v(1)\,D(t)\,G(t) + v(1)\int_t^\infty G(s)\,D'(s)\,ds \geq G(t)\,D(t)\,v\left(\hat{\mu}^A(t)\right), \forall t$$

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

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

- $\ell = \frac{\mu_0^A}{1-\mu_0^A} / \frac{\mu_0^P}{1-\mu_0^P}$: 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

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. Concave Discounting. When $D(T)$ is concave, optimal solution is

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

The Agreement Case

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

Proposition. Concave Discounting. When $D(T)$ is concave, 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. Convex Discounting. When $D(T)$ is convex, optimal solution has two phases (if $\mu_0 > 1/2$)

$$t \leq t^* : G_1'(t) < 0, \hat{\mu}'(t) < 0, G_0(t) = 1 - \mu_0$$
$$t \geq t^* : \hat{\mu}(t) = 1/2, \frac{G_0'(t)}{G_0(t)} = \frac{G_1'(t)}{G_1(t)} = \frac{D'(t)}{D(t)}$$

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

The Agreement Case

Proposition. Convex Discounting. When $D(T)$ is convex, optimal solution has two phases (if $\mu_0 > 1/2$)

$$t \leq t^* : G_1'(t) < 0, \hat{\mu}'(t) < 0, G_0(t) = 1 - \mu_0$$
$$t \geq t^* : \hat{\mu}(t) = 1/2, \frac{G_0'(t)}{G_0(t)} = \frac{G_1'(t)}{G_1(t)} = \frac{D'(t)}{D(t)}$$

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: Convex Discounting

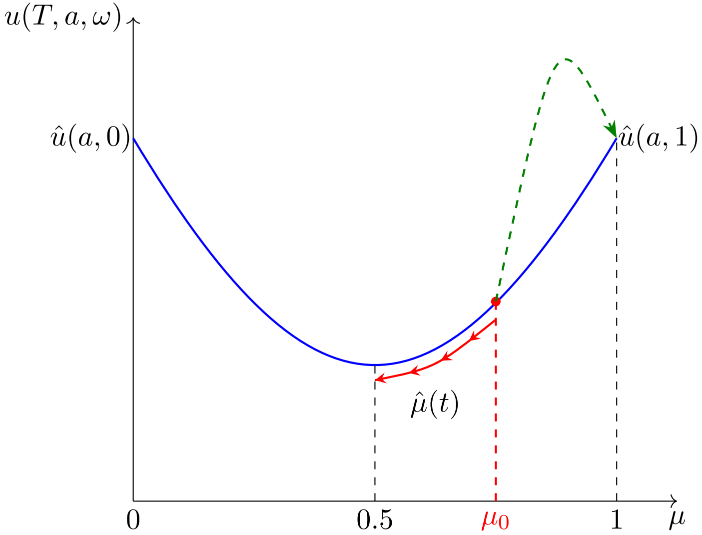

Agreement: Convex Discounting

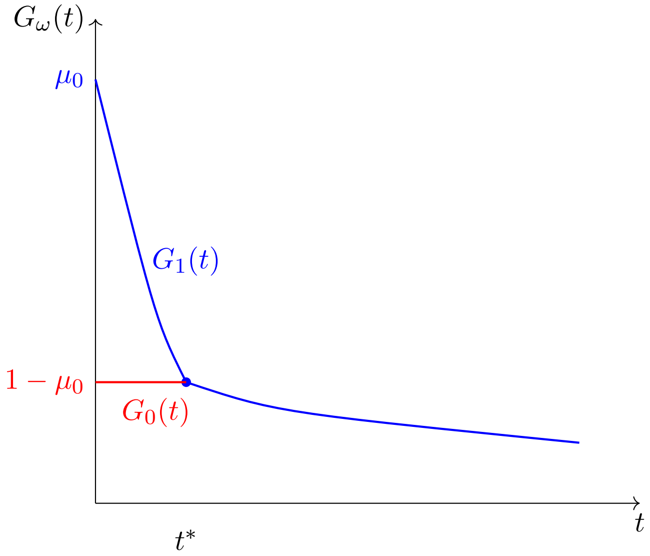

Agreement: Convex Discounting

- Two phases with time-varying Poisson revelation of information
  - Phase 1: Arrival of news about the more likely state at rate $> -\frac{D'(t)}{D(t)}$
  - Phase 2: Arrival of news about both state at rate $-\frac{D'(t)}{D(t)}$
- 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

Agreement: Convex-Concave

- Suppose there exists an inflection point T_i where $D(T)$ is convex below T_i and concave above T_i.
  - Possible with habit formation:

Proposition. Optimal information structure has (at most) three phases:

Phase1. More likely state is revealed according to poisson,

Phase2. Both states are revleaed at rate $-D'(t)/D(t)$,

Phase3. Silence followed by revelation of both states.

- Phase 3 often starts before T_i

Agreement: Convex-Concave

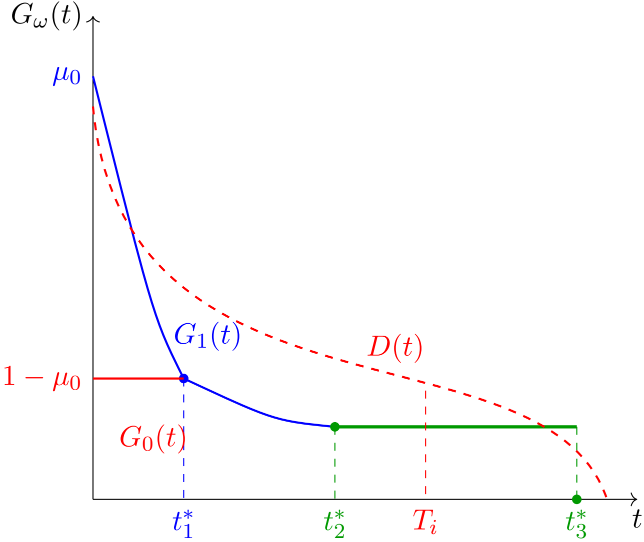

Disagreement

- Payoff of P
$$\int_0^\infty \left(\hat{\mu}^A(t) + \left(1-\hat{\mu}^A(t)\right)\ell\right)\left[G_0(t) + G_1(t)\right] dt$$
where $\ell = \frac{\mu_0^A}{1-\mu_0^A} / \frac{\mu_0^P}{1-\mu_0^P}$ is the relative likelihood ratios.
- We are writing everyone's payoff as a function of beliefs of the agent.
- 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: Convex Discounting

Proposition. Convex Discounting and Disagreement. Suppose $D(T) = e^{-\delta T}$ and $\mu_0^A < \mu_0^P$, then optimal solution two phase

$$t \leq t^* : G_0'(t) < 0, \hat{\mu}'(t) > 0, G_1(t) = \mu_0^A$$
$$t \geq t^* : \hat{\mu}(t) = \mu^*(t) > \mu_0^A, \frac{G_0'(t)}{G_0(t)} = \frac{G_1'(t)}{G_1(t)} = -\delta$$

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

Catering to the Bias

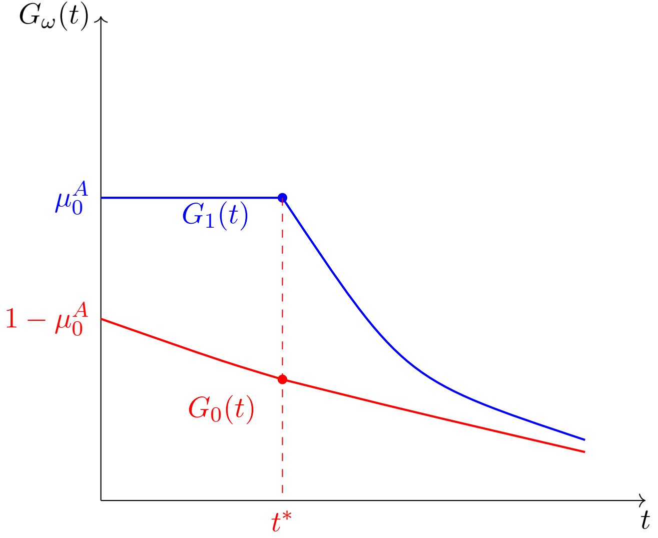

Disagreement: Concave Discounting

Proposition. Concave Discounting and Disagreement. Suppose $\mu_0^A < \mu_0^P$, then optimal solution is

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

- Abrupt full revelation is not optimal:
  - P gets different payoffs in each state; would rather reveal state 1 later

Disagreement: Concave Discounting

Proposition. Concave Discounting and Disagreement. Suppose $\mu_0^A < \mu_0^P$, then optimal solution is

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

- Abrupt full revelation is not optimal:
  - P gets different payoffs in each state; would rather reveal state 1 later
- Only one phase:
  - *Cater to the bias phase*: reveal the A-optimistic state until A is fully pushed to the pessimistic state

Catering to the Bias

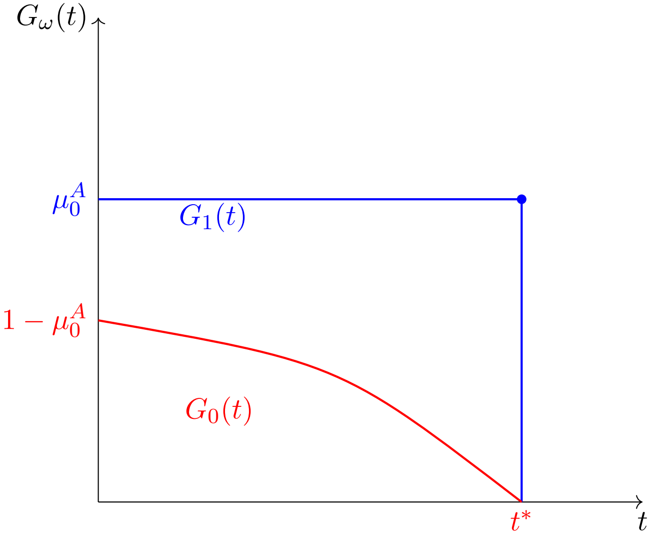

Conclusion

- 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 convex case
  - Nonlinear taxes might
- A lot more to be done:
  - Behavioral aspects: digital addiction, entertainment/suspense and surprise
  - Competition
  - Optimal regulation without violating first ammendment (in the U.S.)

Thank You