Getting the Agent to Wait

Maryam Saeedi Yikang Shen Ali Shourideh

SITE 2024: Dynamic Games, Contracts, and Markets

August 6, 2024

Maryam Saeedi, Yikang Shen, Ali Shourideh

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
 - $\circ~$ 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
 - $\circ\,$ Principal: Maximize engagement
 - 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
 - $\circ\,$ Principal: Maximize engagement
 - Agent: Acquire information, time cost
- Why do we care? Filter Bubbles!
 - $\circ~$ Personalized news aggregators: sometimes blamed for polarization for amplifying \underline{biases}

• Relative intertemporal preferences of the parties is important

- Relative intertemporal preferences of the parties is important
- Disagreement in prior a form of bias:
 - Cater to the bias: reveal the state that the agent think is more likely first
 - Some form of gradual revelation is optimal

- Relative intertemporal preferences of the parties is important
- Disagreement in prior a form of bias:
 - 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 speed.

Related Literature _____

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

Related Literature ____

- Information economics and Bayesian Persuasion:
 - Rayo and Segal (2010), Kamenica and Gentzkow (2011) and 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): Let's wait for that!

Related Literature _____

- Information economics and Bayesian Persuasion:
 - Rayo and Segal (2010), Kamenica and Gentzkow (2011) and 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): Let's wait for that!
 - Complimentary to our results: our focus is on heterogeneity in priors with and without private information

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

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

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

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

Possibility of (belief) Disagreement

- Priors $\mu_0^A = \mathbb{P}^A \left(\omega = 1 \right), \mu_0^P = \mathbb{P}^P \left(\omega = 1 \right) \in (0, 1)$
- Learning is Bayesian (details later)

Possibility of (belief) Disagreement

- Priors $\mu_0^A = \mathbb{P}^A \left(\omega = 1 \right), \mu_0^P = \mathbb{P}^P \left(\omega = 1 \right) \in (0, 1)$
- Learning is Bayesian (details later)
- Two interpretations;
 - Subjective beliefs (Savage(1972)) and agree to disagree a la Aumann (1976)
 - Different source of prior knowledge/information
- A form of mis-specified learning with "dogmatic" prior as the source of mis-specification:
 - Berk (1966), Fudenberg, Lanzani, Strack (2021, 2022, etc.), Bohren, Hauser (2021)
- Start from μ_0^A and μ_0^P being public information:
 - Later allow for μ_0^A being private.

Timing



Maryam Saeedi, Yikang Shen, Ali Shourideh

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

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

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

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

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

• \mathcal{F}_t^A is similarly calculated

- Bayes rule: relative likelihood ratio stays constant
- Equilibrium is standard:
 - A cannot commit to exit strategies
 - P can commit to information structure

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

- Idea of proof: If not, then split the signal into two fully revealing signals each with probability μ_t^P and $1 \mu_t^P$. 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

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
- Does not include $|A| < \infty$, since $v(\mu)$ is piecewise linear
 - $\,\circ\,$ Handwavy argument: can approximate with smooth convex functions

- Can apply Caratheodory theorem
 - $\circ~3$ signal realizations in each period is sufficient: $\Omega \cup \{ \mathrm{No} \ \mathrm{News} \}$
- Choice of information structure is equivalent to choice of two D.D.F functions (decumulative distribution functions)

$$G_{P,\omega}(t) = \mathbb{P}^{P} (\text{exit} \ge t, \omega)$$
$$\mu^{P}(t) = \mathbb{P}^{P} (\omega | \text{stay until } t)$$
$$= \frac{G_{P,1}(t)}{G_{P,1}(t) + G_{P,0}(t)}$$

• D.D.F's are decreasing and $G_{P,1}(0) = \mu_0^P = 1 - G_{P,0}(0)$

Optimal Information Provision

$$\max_{G_{0},G_{1}} \int_{0}^{\infty} e^{-\delta_{P}t} \left(G_{P,1} \left(t \right) + G_{P,0} \left(t \right) \right) dt$$

subject to

$$v(1) G_A(t) - v(1) \delta_A \int_t^\infty e^{-\delta_A(s-t)} G_A(s) ds \ge G(t) v(\mu^A(t)), \forall t$$
$$G_\omega(t) : \text{ non-increasing}$$
$$G_1(0) = 1 - G_0(0) = \mu_0^A$$

•
$$\ell = \frac{\mu_0^A}{1-\mu_0^A} / \frac{\mu_0^P}{1-\mu_0^P}$$
: relative likelihood ratio

Maryam Saeedi, Yikang Shen, Ali Shourideh

- Objective is linear in $G_{\omega}(t)$; constraint set is convex \rightarrow use results from Luenberger (1969)
- A technical issue: cannot readily use strong duality
 - By $T = \infty$, all information should be revealed $(e^{-\delta_A T} \to 0)$
 - IC is going to be binding
 - Constraint set has an empty interior
- Work around:
 - $\circ~$ Assume time is finite: $T \leq \hat{T} < \infty \rightarrow$ strong duality holds
 - $\circ\,$ Bound the multiplier by a function of \hat{T} and take limits. Use Berge's maximum theorem

SIMPLE EXAMPLE

Maryam Saeedi, Yikang Shen, Ali Shourideh

- Restrict to extreme revelation policies: Have to reveal the state
- P: chooses time $T \in \mathbb{R}_+ \cup \{0\}$ to reveal the state
- Time preferences: $\delta_P = 0 < \delta_A$, i.e., P more patient
- Priors: $\mu_A = \mu_P = 1/2$

$$v(\mu) = \begin{cases} 1 & \mu = 0, 1\\ 1/2 & \mu = 1/2 \end{cases}$$

Maryam Saeedi, Yikang Shen, Ali Shourideh

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



• Spread revelation time around T^*



Maryam Saeedi, Yikang Shen, Ali Shourideh

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



• Distribution: exponential at rate δ_A ; Poisson revelation



Maryam Saeedi, Yikang Shen, Ali Shourideh

Simple Example _____

- Alternative: $\delta_P > \delta_A$, i.e., A more patient
- Rewrite:

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

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

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

Maryam Saeedi, Yikang Shen, Ali Shourideh

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



Maryam Saeedi, Yikang Shen, Ali Shourideh

Summary of Example _____

• Key object: Marginal cost of waiting for the agent relative to marginal benefit of engagement:

$$RMCE\left(T\right) = -\frac{\partial u_{A}/\partial T}{\partial u_{P}/\partial T}$$

Summary of Example ____

• Key object: Marginal cost of waiting for the agent relative to marginal benefit of engagement:

$$RMCE(T) = -\frac{\partial u_A/\partial T}{\partial u_P/\partial T} = \frac{\delta_A}{\delta_P} e^{-(\delta_A - \delta_P)T}$$

• How RMCE changes over time determines communication

- With decreasing RMCE, random revelation earlier allows reaching time with lower RMCE and thus extend further.
- Quantity of information is fixed
 - Clearly can be varied by gradual slant, mixed messaging, etc.

PROPERTIES OF OPTIMAL COMMUNICATION

Maryam Saeedi, Yikang Shen, Ali Shourideh

More Impatient Principal

Proposition. Suppose that $\delta_P > \delta_A$ and $\ell \leq 1$. Then optimal communication consists of two phases and two instantaneous revelations times:

- In phase 1, $t \in [0, t_1^*)$, no information is revealed.
- At t_1^* , $\omega = 0$ is revealed with a positive probability.
- In phase 2, $t \in [t_1^*, t_2^*]$, $\omega = 0$ is revealed gradually according to a time-varying Poisson process so that IC binds.
- At t_2^* , $\omega = 1$ is revealed such that $\mu_A(t_2^*) = 1$.

Length of phase 2 is positive, i.e., $t_1^* < t_2^*$, if and only if $\ell < 1$.

More Impatient Principal

Proposition. Suppose that $\delta_P > \delta_A$ and $\ell \leq 1$. Then optimal communication consists of two phases and two instantaneous revelations times:

- In phase 1, $t \in [0, t_1^*)$, no information is revealed.
- At t_1^* , $\omega = 0$ is revealed with a positive probability.
- In phase 2, $t \in [t_1^*, t_2^*]$, $\omega = 0$ is revealed gradually according to a time-varying Poisson process so that IC binds.
- At t_2^* , $\omega = 1$ is revealed such that $\mu_A(t_2^*) = 1$.

Length of phase 2 is positive, i.e., $t_1^* < t_2^*$, if and only if $\ell < 1$.

- $\ell > 1$: mirror case
- Catering to the bias: if $Pr^{A}(\omega) > Pr^{P}(\omega)$, reveal ω first.

Catering to the Bias



Maryam Saeedi, Yikang Shen, Ali Shourideh

More Patient Principal

• What if $0 \leq \delta_P < \delta_A$?

More Patient Principal ____

- What if $0 \le \delta_P < \delta_A$?
- Stationary Engagement: values of $(\mu_A, \mu_P^*(\mu_A))$ so that both states are revealed gradually at the same rate
- Why it exists?
 - Start from symmetric revelation to keep A indifferent

$$\frac{\lambda}{\lambda + \delta_A} v\left(1\right) = v\left(\mu_A\right)$$

- Deviate to reveal only $\omega = 1$, over (t, t + dt) at rate q and switch back to symmetric poisson revelation after t + dt.
- Costs and benefits:



• Stationary Engagement:
$$\Delta = 0$$
.

Maryam Saeedi, Yikang Shen, Ali Shourideh

More Patient Principal



Maryam Saeedi, Yikang Shen, Ali Shourideh

Catering to the Bias



Maryam Saeedi, Yikang Shen, Ali Shourideh

Proposition. 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 IC binds.
- 2. If $\mu_P^*(\mu_A) < \mu_P$, in phase 1 only state $\omega = 1$ is gradually revealed so that the agent's IC binds.
- 3. In phase 2, when $\mu_P^*(\mu_A) = \mu_P$, both states are gradually revealed according to a Poisson process with intensity λ which satisfies $\frac{\lambda}{\lambda + \delta_A} v(1) = v(\mu_A)$.

Extreme Catering to the Bias



Maryam Saeedi, Yikang Shen, Ali Shourideh

NON-PERSONALIZED COMMUNICATION

Maryam Saeedi, Yikang Shen, Ali Shourideh

Non-Personalized Communication

- How does communication change when μ_A is not observed?
- A's belief is either μ_A^L or μ_A^H , with $\mu_A^L < \mu_A^H$, while the principal's belief is μ_P , $\alpha^j = \Pr\left(\mu_A^j\right)$.
- Focus on $\delta_A > \delta_P$.
- 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

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

where $\mu_A^j(s^t)$ is the agent of type j'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.



Speed vs. Quality ____

- If type L is closer to certainty, P might want to reveal <u>some</u> information so that L exits but H does not.
- As a result, L is exiting with less than perfect information
- Since L needs to be incentivized beforehand, speed has to be higher than the personalized case
- Trade-off of Personalized vs. Non-personalized:
 - Higher Quality vs. Higher Speed

Stationary Engagement ____

- Beliefs in phase 1 remain constant
- Stationary strategy for the principal:
 - Arrival rate of the transition signal λ^* from phase 1, and a distribution of posteriors, $\{p_{\sigma}^*, \mu_{\sigma}^*\}_{\sigma \in \Sigma}$.
 - Beliefs are martingale

$$\mu_P^* = \sum_{\sigma \in \Sigma} p_\sigma^* \mu_\sigma^*$$

- $\circ\,$ For each realization of $\sigma,$ recommend the type with higher personalized payoff to stay.
- Use duality to cast stationary engagement as a constrained concavification

Stationary Engagement _

Proposition. The steady state level of belief for the principal μ_P^* is either 0 or 1 and is achieved in finite time, or $\mu_P^* \in (0, 1)$ and $\lambda^*, \Lambda_L^*, \Lambda_H^* \ge 0$ exists that satisfy:

- 1. The Belief Smoothing (Euler-Lagrange) equation holds,
- 2. The phase 1 optimality of symmetric transition holds,
- 3. The following incentive compatibility and complementary slackness conditions are satisfied:

$$\frac{\lambda^*}{\lambda^* + \delta_A} \sum_{\sigma} p_{\sigma}^* v_j \left(\mu_{\sigma}^*\right) \ge v_j \left(\mu_P^*\right), \text{ with equality if } \Lambda_j^* > 0.$$

Maryam Saeedi, Yikang Shen, Ali Shourideh

Computational Example .



Maryam Saeedi, Yikang Shen, Ali Shourideh

RANDOM EXIT AND BELIEF POLARIZATION

Maryam Saeedi, Yikang Shen, Ali Shourideh

Distribution of Exit Times

• Allow for random exit so that even in the personalized model A does not have perfect information upon exit



Maryam Saeedi, Yikang Shen, Ali Shourideh

Belief Distributions



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 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
 - Competition
 - Optimal regulation without violating first ammendment (in the U.S.)

THANK YOU FOR STAYING ENGAGED!

Maryam Saeedi, Yikang Shen, Ali Shourideh