Engagement Maximization

Benjamin Hébert (Stanford) and Weijie Zhong (Stanford)

October 12, 2022

Introduction

- Internet platforms sequentially present content to users
  - "news feeds," "timelines," and other recommendations
  - this content is accompanied by ads
- The platform profits in proportion to users' "engagement"
  - the total attention paid to the content
  - if user pays full attention, equivalent to time spent
- This is a principal-agent problem
  - The principal (platform) chooses the content
  - The agent (user) chooses whether or not to pay attention
- How can the platform maximize the user's engagement? What are the consequences of engagement maximization?

What We Do

- Model engagement maximization as principal-agent version of dynamic rational inattention (RI) problem
  - i.e. principal-agent version of Hébert and Woodford [2021], Zhong [2022]
- The agent chooses when to stop, and what action to take when stopping
  - Information is valued for instrumental purposes
- The agent receives signals, and updates her beliefs using Bayes' rule
- Unlike RI: the principal, not the agent, chooses the signals the agent receives
  - subject to the agent's information processing limits

What We Show

- **Welfare Minimization**: the agent is no better off with platform than without it
  - Even though the agent chooses when to stop
- **Belief Polarization**: when she stops and acts, the agent will hold extreme beliefs
  - relative to the RI benchmark in which the agent chooses the signals
- **Learning Paths**: all signals are either
  1. decisive (causing agent to stop and act), or
  2. suspensive (making the agent "more uncertain" in the appropriate sense)

Key point: paying for information with time instead of money distorts the information provider's incentives

Outline

1. Simplified example
    1.1 to illustrate main ideas
2. General model
3. Summary of extensions
4. Alternative interpretation: teaching test-motivated students

Example: States and Preferences

- Agent chooses action $a \in A \equiv \{l, r\}$
- State $x \in X \equiv \{L, R\}$, agent has uniform prior
- Utility if action a at time t in state x: $u_{a,x} - \kappa t$
- In this example: $\kappa = 2$,

$$u_{l|L} = u_{r|R} = 1,\ u_{l|R} = u_{r|L} = -1$$

  - note: agent gets zero expected utility from acting without info.
- Interpretation: κ is opportunity cost of time net of utility from using platform
  - tractable case: κ constant
  - extension: increasing opp. cost / decreasing utility of use

Example: Signals

- Example only: restricted signals, realizations $\{s_l, s_r\}$, Poisson arrival:

	L	R
s_l	$\alpha \cdot p$	$\alpha \cdot (1-p)$
s_r	$\alpha \cdot (1-p)$	$\alpha \cdot p$

- $\alpha \in \mathbb{R}_+$: arrival frequency $p \in [\frac{1}{2}, 1]$: precision
- symmetry: beliefs remain at prior absent Poisson arrival
- Interpretation: news feed presents content to user at a constant rate
  - α: likelihood content in a given instant is relevant for user
  - p: depth or informativeness of content, conditional on relevance

Example: Information Constraints

- User's information processing constraint:

$$\alpha \cdot \underbrace{\left(H^S\left(\frac{1}{2}\right) - H^S(p)\right)}_{\text{convex function of } p \text{, minimized at } p=\frac{1}{2}} \leq 1,$$

  where $H^S(p)$ is Shannon's entropy (specifics not important for example)
    - mutual information between signals and state from t to $t+\delta$ less than δ
- Interpretation: breadth vs. depth tradeoff
    - feed can cover broad range of topics, making it more likely one is relevant (high α)
    - or provide more info on a narrower range (high p)
    - but user can't process broad and in-depth coverage simultaneously

Example: The Agent-Optimal Benchmark

- Consider the signal with $p^B = \frac{e}{1+e}$ and $\alpha^B = (H^S\left(\frac{1}{2}\right) - H^S(\frac{e}{1+e}))^{-1}$.
- Suppose agent acts with action l (r) upon receiving s_l (s_r)
- Expected utility is

$$V^B = \underbrace{\frac{e-1}{1+e}}_{\text{Utility from Actions}} - \underbrace{2(\alpha^B)^{-1}}_{\text{Expected Cost of Delay}} \approx 0.24$$

- Strategy is optimal among all feasible (p, α)
  - and among all signals satisfying same mutual info. bound (Hébert and Woodford [2021])

Example: A Better Strategy for the Principal

- Engagement proportional to total attention, $H^S\left(\frac{1}{2}\right) - H^S(p)$
  - If constraint binds, proportional to α^{-1}, expected decision time
- Consider signal with $\tilde{p} = \frac{4e}{1+4e}$ and $\tilde{\alpha} = (H^S\left(\frac{1}{2}\right) - H^S(\frac{4e}{1+4e}))^{-1}$
  - More precise, less frequent than agent benchmark

$$\tilde{V} = \underbrace{\frac{4e-1}{1+4e}}_{\text{Utility from Actions}} - \underbrace{2\tilde{\alpha}^{-1}}_{\text{Expected Cost of Delay}} \approx 0.02 > 0$$

- The agent still benefits from the signal (and so will be willing to wait for it)
- The principal gets more engagement

Example: Takeaways

- Optimal principal: a slightly higher $\tilde{p}$, setting $\tilde{\alpha}$ to max and $\tilde{V}$ to zero:

1. Agent driven to reservation value (welfare minimization)
2. $\tilde{p} > p^B$: extreme beliefs when acting (extreme beliefs)
3. beliefs remain at prior until stopping (suspensive/decisive signals)

General Model

- Finite set of states X, actions A, arbitrary utility $u_{a,x}$, cost of delay $\kappa > 0$
- Notation: $\mathcal{P}(X)$ is simplex defined on X
- Indirect utility from acting with belief $q' \in \mathcal{P}(X)$:

$$\hat{u}(q') = \max_{a \in A} \sum_{x \in X} q'_x u_{a,x}.$$

- Sample space: $\Omega \equiv \mathbb{D}(\mathcal{P}(X))$ (set of càdlàg functions $\mathbb{R}_+ \to \mathcal{P}(X)$)
- Agent's belief process, $q : \Omega \times \mathbb{R}_+ \to \mathcal{P}(X)$, is canonical process on Ω,

$$q_t(\omega) = \omega(t)$$

- $\{\mathcal{F}_t\}$: natural filtration associated with q, $\mathcal{F} = \lim_{t \to \infty} \mathcal{F}_t$

The Agent's Problem

- $\mathcal{T} \subset \Omega \to \mathbb{R}_+$ is set of non-negative stopping times with respect to $\{\mathcal{F}_t\}$.
- Agent's problem: choose when to stop and which action to take
- Given a probability measure P defined on $(\Omega, \mathcal{F})$, solve

$$V(P) = \sup_{\tau \in \mathcal{T}} E^P[\hat{u}(q_\tau) - \kappa\tau | \mathcal{F}_0]$$

  - Given P, $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}, P)$ is a filtered probability space
  - q is the agent's belief process (an $\mathcal{F}_t$-martingale, by Bayes' rule)
  - $\hat{u}(q_\tau)$ is the utility from acting with the stopping belief q_τ
  - the measure P defines how the agent's beliefs will evolve given the principal's signals

The Information Processing Constraint

- The principal will choose P, subject to (for all $t > 0$, P-a.s.),

$$\frac{\mathrm{d}}{\mathrm{d}t} I_t = \limsup_{h \downarrow 0} \frac{1}{h} \mathbb{E}^P [H(q_{t+h}) - H(q_{t-}) | \mathcal{F}_{t-}] \leq \chi.$$

  - the c.t. version of a "UPS" rational inattention cost (Caplin et al. [2022])
  - H is a strongly convex, twice continuously-differentiable function
  - χ is a bound on the rate of information acquisition (example: $\chi = 1$)
  - I_t is the cumulative information acquisition, with $I_0 = 0$
    - $I_t = \chi t$ if bound is always tight
  - Example: H negative Shannon's entropy, I_t mutual information of state and signals received from zero to t
- Second constraint: agent's initial prior, $\bar{q}_0 \in \mathcal{P}(X)$, taken as given

The Principal's Problem

- Admissible strategies: $(P, I, \tau) \in \mathcal{A}(\bar{q}_0)$ if $I_0 = 0$,
  - P is measure on $(\Omega, \mathcal{F})$ such that $q_0 = \bar{q}_0$ with prob. one, and
  - Information processing constraint is satisfied
- Principal's problem: given $\bar{q}_0 \in \mathcal{P}(X)$, solve

$$J(\bar{q}_0) = \sup_{(P,I,\tau)\in\mathcal{A}(\bar{q}_0)} \underbrace{\mathbb{E}^P[I_\tau|\mathcal{F}_0]}_{\text{engagement}}$$

subject to agent being willing to follow stopping recommendation,

$$\tau \in \arg\max_{\tau'\in\mathcal{T}} \mathbb{E}^P[\hat{u}(q_{\tau'}) - \kappa\tau'|\mathcal{F}_0].$$

Assumptions

1. The principal acts first (chooses P before agent chooses τ)
   - interpretation: user takes recommendation algorithm as fixed
   - agent commitment irrelevant
2. The principal can implement any martingale belief process satisfying bounds
   - interpretation: large set of possible content to recommend
   - extension: principal can violate bounds, agent must selectively attend
3. The principal commits to P
   - extension: case without commitment differs only in minor respects
4. The agent follows principal recommendations when indifferent
   - a minor modification of P ensures strict incentives

No Full Learning Assumption

Assumption 1
We assume that

$$\hat{u}(\bar{q}_0) - \frac{\kappa}{\chi} H(\bar{q}_0) > \sum_{x \in X} \bar{q}_{0,x} (\hat{u}(e_x) - \frac{\kappa}{\chi} H(e_x)),$$

where $e_x \in \mathcal{P}(X)$ *has full support on state* $x \in X$*.*

- If agent can only learn everything or nothing, would choose nothing
- Could be automatic (if H infinite on boundary) or require κ large enough relative to χ (if not)
- Avoids case in which learning everything is agent-optimal policy
    - extreme beliefs result would not apply in this case

Solution Method

- Problem looks complicated (the set of martingales + stopping times is big)
- Strategy: solve a relaxed problem, then show solution is feasible in original problem
- Key idea: move from P (measure on $(\Omega, \mathcal{F})$) to π (measure on $q_\tau \in \mathcal{P}(X)$)

Lemma 1

$\forall (P, I, \tau) \in \mathcal{A}(q_0)$ *satisfying the agent's optimal stopping,*

$$\mathbb{E}^\pi[\widehat{u}(q) - \hat{u}(\bar{q}_0)] \geq \kappa \mathbb{E}^P[\tau | \mathcal{F}_0] \geq \frac{\kappa}{\chi} \mathbb{E}^\pi[H(q) - H(\bar{q}_0)]$$

Static RI

- Relaxed problem: consider only those constraints,

$$\bar{J}(q_0) = \sup_{\pi \in \mathcal{P}(\mathcal{P}(X)):\mathbb{E}^{\pi}[q]=q_0} \mathbb{E}^{\pi}[H(q) - H(\bar{q}_0)]$$

$$s.t. \ \frac{\kappa}{\chi}\mathbb{E}^{\pi}[H(q) - H(\bar{q}_0)] \leq \mathbb{E}^{\pi}[\widehat{\mu}(q) - \widehat{\mu}(\bar{q}_0)].$$

Proposition 1

Relaxed problem has a solution π^* with finite support. Either $\bar{J}(q_0) = 0$ or $\bar{J}(q_0) > 0$ and there exists $\lambda > \frac{\chi}{\kappa}$ s.t. all π^* satisfy

$$\pi^* \in \arg \max_{\pi \in \mathcal{P}(\mathcal{P}(X)):\mathbb{E}^{\pi}[q]=\bar{q}_0} \mathbb{E}^{\pi}\left[\hat{u}(q) - (\frac{\kappa}{\chi} - \frac{1}{\lambda})(H(q) - H(\bar{q}_0))\right].$$

- π^* is the solution to a standard (static) RI problem with cost $\theta = \frac{\kappa}{\chi} - \frac{1}{\lambda}$

Implementation

- Given $\pi^* \in \mathcal{P}(\mathcal{P}(X))$, define q_t as:

$$q_t = \bar{q}_0 + 1_{N_\alpha(t)\geq 1} \cdot (Q - \bar{q}_0),$$

  - $Q \in \mathcal{P}(X)$ is a random variable distributed according to π^*
  - $N_\alpha(t)$ is an independent Poisson counting process with parameter α.
  - We say q_t is an $\alpha-$*dilution* of π.

Theorem 1

$\forall \bar{q}_0 \in \mathcal{P}(X)$, there exists $\pi^* \in \mathcal{P}(\mathcal{P}(X))$ with finite support solving the relaxed problem. If $Supp(\pi^*) = \{\bar{q}_0\}$, the agent will immediately stop and any feasible policy is optimal. Otherwise, let $\alpha^* = \frac{\chi}{\mathbb{E}^{\pi^*}[H(q)-H(\bar{q}_0)]}$, and let (P^*, τ^*) be the law and jumping time of the α^*-dilution of π^*. Then, $(P^*, I_t^* = \chi t, \tau^*)$ is an optimal solution to the principal's problem.

Immediate Implications

- **Welfare Minimization**: under any optimal policy, $V(q_0) = \hat{u}(q_0)$
    - agent receives no surplus
- Commitment is unnecessary:
    - if both agents expect the policies of this solution with commitment to played going forward, they will be willing to adopt those policies at the current moment.
- Jumps are necessary:
    - will follow from extreme beliefs results
    - observation: agent will stop if beliefs reach agent-optimal stopping region

The Agent-Optimal Benchmark

- Compare to case in which agent chooses $(P^B, I^B, \tau^B) \in \mathcal{A}(\bar{q}_0)$
- By Hébert and Woodford [2021] and Zhong [2022]:

$$\mathbb{E}^{P^B}[\hat{u}(q_\tau) - \kappa\tau^B|\mathcal{F}_0] = V^B(\bar{q}_0)$$
$$= \max_{\pi\in\mathcal{P}(\mathcal{P}(X)):\ E^\pi[q]=\bar{q}_0} \mathbb{E}^\pi[\hat{u}(q) - \frac{\kappa}{\chi}(H(q) - H(\bar{q}_0))].$$

  - Same equivalence with static RI, cost $\theta^B = \frac{\kappa}{\chi}$
- Let π^B be an optimal policy in the agent-optimal benchmark

Extreme Beliefs

- $Q^i(\bar{q}_0)$ be the union of the support of all optimal π^*
  - in the benchmark ($i = B$) and principal-agent ($i = *$) models.
- Let $\mathrm{Conv}\, Q^B(\bar{q}_0)$ be the convex hull of $Q^B(\bar{q}_0)$

Proposition 2
$\forall q \in Q^*(\bar{q}_0), q \notin \mathrm{Conv}\, Q^B(\bar{q}_0)$.

- The principal ensures the agent stops with beliefs that lie outside the convex hull of what the agent would choose
- A formal definition of "extreme beliefs"
- Follows from static RI equivalence, $\theta^B > \theta$

Extreme Beliefs Illustrated

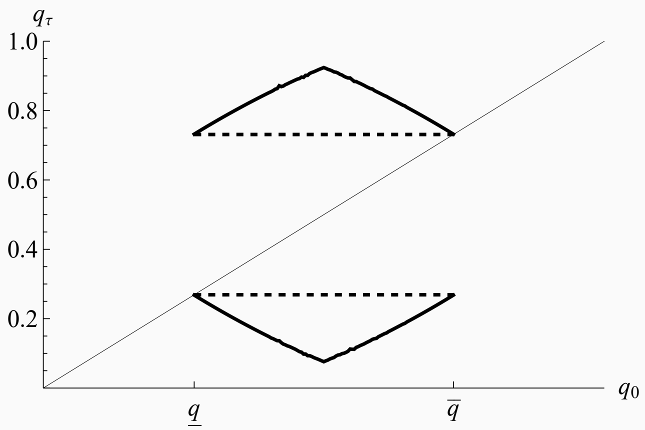

Figure 1: $\text{Supp}(q_\tau)$ as a correspondence of q_0

Engagement Maximization Illustrated

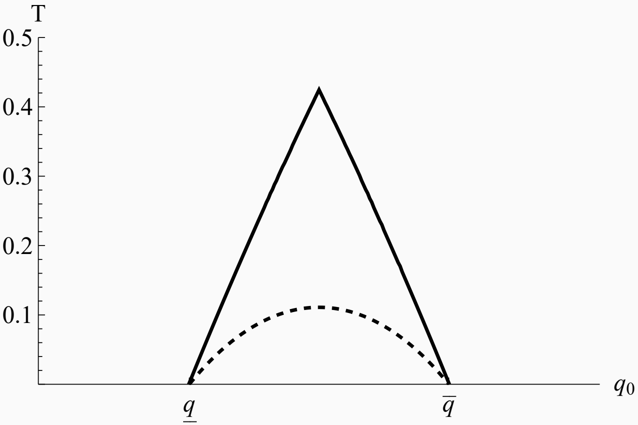

Figure 2: Engagement as a function of q_0

Decisive and Suspensive Beliefs

- Given an optimal policy π^*, $q_t \in \text{Conv}(\text{Supp}(\pi^*))$
- Define a "restricted" static RI problem,

$$V_R(q,\pi^*) = \max_{\pi' \in \mathcal{P}(\text{Supp}(\pi^*)) : E^{\pi'}[q'] = q} \mathbb{E}^{\pi'}\left[\widehat{\mu}(q') - \hat{u}(q) - \frac{\kappa}{\chi}(H(q') - H(q))\right]$$

- Interpretation: value of information to the agent, if the agent were in control but restricted to $q_t \in \text{Conv}(\text{Supp}(\pi^*))$

Definition 1
Given $\pi^* \in \mathcal{P}(\mathcal{P}(X))$, the belief $q \in \mathcal{P}(X)$ is *decisive* if $q \in \text{Supp}(\pi^*)$; the belief $q \in \mathcal{P}(X)$ is *suspensive* if $q \in \text{Conv}(\text{Supp}(\pi^*)) \setminus \text{Supp}(\pi^*)$ and satisfies $V_R(q,\pi^*) \geq 0$.

- Example: with prior $q_0 < \frac{1}{2}$ and stopping beliefs $\text{Supp}(\pi^*) = \{q_l, q_r\}$,
  - $\{q_l, q_r\}$ are the decisive beliefs, and $[q_0, 1-q_0]$ are the suspensive beliefs

Optimal Policies

- Let $\Delta(\pi) \subseteq P(X)$ denote the set of all suspensive beliefs given π.

Proposition 3
Given $\bar{q}_0 \in \mathcal{P}(X)$, let Π^* be the set of solutions to the relaxed problem. Then, in all solutions to the principal's problem, for all $t \in [0, \tau^*)$, beliefs are suspensive given some $\pi^* \in \Pi^*$,

$$Supp(q_t) \subseteq \bigcup_{\pi^* \in \Pi^*} \Delta(\pi^*),$$

and at $t = \tau^*$, beliefs are decisive given some $\pi^* \in \Pi^*$,

$$Supp(q_\tau) \subseteq \bigcup_{\pi^* \in \Pi^*} Supp(\pi^*).$$

Optimal Policies Illustrated

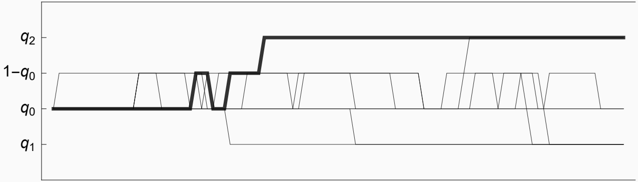

Figure 3: Sample paths of an optimal policy

Summary of Extensions

- If agent can selectively attend, and principal can send excess info:
  - nothing changes
- If $\kappa(t)$ increases over time:
  - problem not tractable without additional assumptions
  - under a lot of symmetry assumptions, main results continue to hold
- If size of jumps in beliefs bounded:
  - bounds not tight: all results hold
  - beliefs continuous + $|X| = 2$: agent-optimal solution
- Version without info. constraint, principal's goal is time-maximization:
  - similar results, but no agent-optimal benchmark

Engaging Test-Motivated Students

- Alternative application of our model
- States are $X = T \times Q$, T is truth, Q is test questions
- Actions are answers to questions
- Teacher: log scoring rule, wants student to learn state $T = \{R, L\} \times \{0, 1\}$
- Student wants to maximize number of correct answers
  - questions only depend on $\{R, L\}$ and not $\{0, 1\}$
- Student: mutual info. constraint
- Optimal policy: teach only test-relevant info, but provide more than student wants
  - result of congruence of mutual info. and log scoring rule

Teaching to the Test, Illustrated

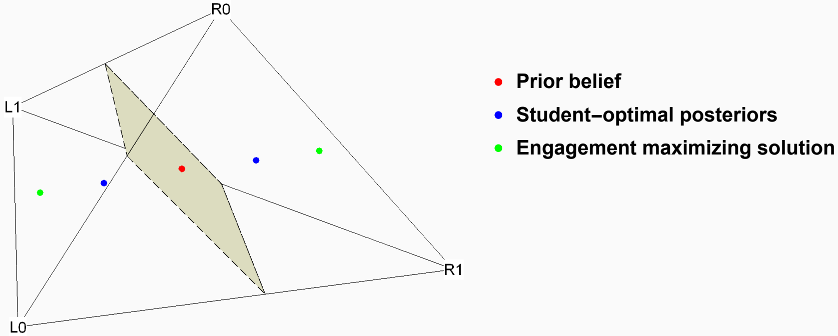

Figure 4: Teaching to the Test

Conclusion

- Maximizing engagement naturally leads to:
  - full surplus extraction by the principal (platform)
  - extreme beliefs held by the agent (user)
  - "sensational content" that is mostly disregarded, and occasionally leads the user down "rabbit holes"
- Despite the rationality of the user and the user's ability to stop at any time

References

Andrew Caplin, Mark Dean, and John Leahy. Rationally inattentive behavior: Characterizing and generalizing shannon entropy. *Journal of Political Economy*, 130 (6):1676–1715, 2022.

Benjamin M Hébert and Michael Woodford. Rational inattention when decisions take time. Technical report, National Bureau of Economic Research w.p. 26415, 2021.

Weijie Zhong. Optimal dynamic information acquisition. *Econometrica*, 90(4): 1537–1582, 2022.