

Getting the Agent to Wait*

Maryam Saeedi
Carnegie Mellon University
msaeedi@andrew.cmu.edu

Yikang Shen
Carnegie Mellon University
yikangs@andrew.cmu.edu

Ali Shourideh
Carnegie Mellon University
ashourid@andrew.cmu.edu

July 11, 2024

Abstract

We examine the strategic interaction between an expert (principal) maximizing engagement and an agent seeking swift information. Our analysis reveals: When priors align, relative patience determines optimal disclosure—impatient agents induce gradual revelation, while impatient principals cause delayed, abrupt revelation. When priors disagree, *catering to the bias* often emerges, with the principal initially providing signals aligned with the agent’s bias. With private agent beliefs, we observe two phases: one engaging both agents, followed by catering to one type. Comparing personalized and non-personalized strategies, we find faster information revelation in the non-personalized case, but higher quality information in the personalized case.

1 Introduction

Maximizing engagement is a central objective across many economic settings, from traditional expert services such as management consulting and legal advice to modern digital platforms with content provision and recommender systems. In these contexts, an expert (principal) aims to prolong user engagement to maximize value extraction, while the user (agent) incurs costs

*We owe the title of this paper to Orlov, Skrzypacz and Zryumov’s work! We thank several seminar and conference participants as well as Nageeb Ali, Arjadha Bardhi, James Best, Lukas Bolte, Joyee Deb, Alexey Kushnir, George Loewenstein, Jesse Shapiro, Vasiliki Skreta, and Ariel Zetlin-Jones.

from extended information acquisition periods. In this paper, we develop a general framework for analyzing such interactions.

To do so, we consider a model in which a principal and an agent interact through engagement. The agent engages to collect information about a payoff-relevant decision, while the principal aims solely to maximize the duration of this engagement. This engagement is irreversible and represents a trade-off: it is costly for the agent but valuable for the principal. A key feature of our model is allowing for the agent to have a different prior from the principal which could be publicly observable to the principal (in the spirit of [Aumann \(1976\)](#)'s agree-to-disagree framework) or privately known to the agent. This approach enables us to explore how differences in beliefs and information asymmetry impact optimal engagement strategies.

Our analysis reveals three key patterns that determine optimal information provision in this environment. First, the relative patience of the agent and principal determines whether information is revealed abruptly or gradually. Second, when the agent and the principal agree to disagree on their prior, it is often the case that optimal disclosure features *catering to the bias* – the principal first reveals the state towards which the agent is more biased towards. Third, when the agent's belief is private information, there is a trade-off between speed and quality as information is revealed faster but engagements end with more uncertain beliefs.

When the principal and the agent share the same prior, the main determinant of information disclosure is how the cost of engagement for the agent evolves over time relative to its benefit for the principal. More specifically, suppose that both the principal and the agent use exponential discounting. The principal then values each moment of engagement at an exponentially declining rate, i.e., he values early engagement more than later ones. When the agent is more *impatient* than the principal, exponential discounting implies that the relative cost of engagement decreases over time. This in turn means that by some gradual revelation early on, the principal is able to postpone revelation for a time when the cost of engagement is lower for the agent. Technically, when the discount rate of the principal is lower than that of the agent, the payoff of the agent is convex in the payoff of the principal, and thus the agent and principal both benefit from random or gradual disclosure. This disclosure takes the form of a Poisson arrival of the news. In contrast, when the agent is more *patient*, the cost of being engaged increases for the agent relative to the benefit for the principal, which will lead to abrupt information revelation.

The equilibrium dynamics shift significantly when the principal and agent have different prior beliefs. This divergence introduces a crucial asymmetry in how each party values engagement across different states. Consider a scenario where the agent is more optimistic about state $\omega = 0$ than the principal. In this case, revealing information about $\omega = 0$ carries higher value for the agent but represents a lower cost for the principal. This asymmetry arises because the principal, believing $\omega = 0$ to be less probable, perceives the promise of revelation in this state as less likely

to materialize than the agent anticipates. Consequently, promising to disclose information about $\omega = 0$ becomes a cost-effective strategy for maintaining agent engagement. This is what we refer to as *catering to the bias*.

When the agent is more *patient*, catering to the bias always occurs. There is a phase in which no information is revealed; mirroring the optimal disclosure pattern under shared priors. This is followed by a second phase where the state towards which the agent is more biased is gradually revealed according to a time-varying Poisson rate. Finally, at the end of the second phase, the other state is revealed instantaneously.

With a more *patient* principal, catering to the bias would occur when the difference between the prior of the agent and the principal is sufficiently large. In this case, a frontier steady state pair of beliefs exists such that the beliefs are going to converge to it. When the priors agree, this will be the point of maximum uncertainty and maximum engagement. Along this frontier, the optimal disclosure involves symmetric, constant-rate Poisson disclosure of both states. As such, beliefs remain stable on the frontier over time until a disclosure event occurs.

The evolution of beliefs (probability that $\omega = 1$ for the agent (x-axis) and the principal (y-axis)) away from this steady state is depicted in Figure 1. Below the stationary engagement curve, the state $\omega = 1$ is revealed according to a time-varying Poisson process, and above it the state $\omega = 0$ is going to be revealed until we land on the steady state curve. As shown in Figure 1, the steady state curve crosses the 45° line exactly at the midpoint when $\mu_P = \mu_A = \frac{1}{2}$; which is the special case of starting with the same priors. When the priors do not agree, the point of stationary engagement is no longer at $\mu_A = \frac{1}{2}$. This shift occurs because at $\mu_A = 1/2$, the principal believes one state is less likely than the other, revealing information about that state becomes less costly for the principal while maintaining the same value for the agent. Consequently, the optimal strategy gravitates towards revealing the state the principal deems less likely. The point of steady state is the one that equates costs and benefits of revealing the state given the differences in beliefs.

Building on these insights, we extend our model to examine cases where agents' beliefs are private information. This extension enables us to compare optimal information revelation strategies in a non-personalized setting, where agents with different beliefs are exposed to the same source of information (analogous to mass media in the context of news) with the personalized strategies (similar to personalized social media news feeds) discussed earlier. Our focus is on how these different approaches impact belief evolution, speed of information revelation, and quality of information.

In this extended model, we consider an economy with two agent types, each identified by a different prior belief. To reflect non-personalized communication, all information is public and




Figure 1: Evolution of beliefs under catering to the bias with a patient principal

non-targeted, i.e., types are privately known. This non-personalized environment shares some key features with its personalized counterpart. In both cases, we identify a steady state, in the non-personalized case determined through via a simple constrained concavification.

However, the non-personalized model diverges in a few important aspects. First, the non-personalized setting may see one agent type departing with incomplete information as the principal may wish to keep only one type around while letting the other type leave. This is in contrast with the personalized setting, where agents exit only upon full state revelation. In other words, the principal strategically balances retaining one type while allowing the other to leave. If only one type of agents remains, we revert to personalized strategy. Moreover, in order to make information acquisition incentive compatible, the principal should increase the speed of exit. In other words, there is a trade-off between speed and quality of information arrival.

Finally, we conclude our paper by incorporating more general forms of discounting and adding random exogenous exit for the agents. The model with exit allows us to investigate and compare numerically personalized and non-personalized delivery of news. We observe two main differences between these two modes of communication: their speed of delivery and the quality of delivery.

When news is targeted and personalized, the principal can keep the agent longer and maintain a lower speed of news delivery. In contrast, the principal has to increase the speed of news delivery in the non-personalized case. In the non-personalized case, two forces lead to this increase in speed. First, the principal cannot necessarily keep both types as uncertain as she wishes, given that they start with different priors. Second, upon delivery of news, the principal may wish to keep one type of agent while letting the other go, so the signal will not reveal the state perfectly and therefore the increase in the value for agents is going to be smaller. Both of these forces,

higher steady state value for agents and smaller increase in this value upon arrival of news, lead to higher speed of news delivery to keep them engaged.

The quality of news is consistently higher in the case of personalized news as the principal will never wish to have an agent leave without full information. We also compare the impact of mode of delivery on polarization. We see evidence of polarization in both cases and no clear increase in the personalized case. This suggests that personalization of news delivery may not necessarily exacerbate polarization more than non-personalized approaches.

Ultimately, we view our paper as providing a framework to analyze incentives for attention and its manipulation via engagement. This is especially important as the internet and social media become an integral part of daily life, and engagement and advertising still remain the main source of monetization on the internet. In fact, there is some evidence that our model does capture some behavior of advertisers and web designers, see the discussion in section 2.1. Specifically, various academics and policy makers, concerned by the spread of misinformation on social media, have advocated taxes on digital advertising.¹ Through the lens of our model, a flat proportional tax on advertising does not necessarily change the engagement strategy of the principal as long as it does not remove incentives for this activity altogether. Moreover, our model also suggests that belief polarization is not necessarily worse with personalization of news (due to the trade-off between speed and quality).

Related Literature Our paper relates to several strands of the literature on contracting, mechanism design, and information design. The most relevant is the literature on dynamic Bayesian persuasion as in Ely (2017), Renault et al. (2017), Ely and Szydlowski (2020), Orlov et al. (2020), Smolin (2021), and Che et al. (2023), among many others, which build on the static persuasion model of Kamenica and Gentzkow (2011). The key difference in our setting is that the agent’s final decision does not affect the principal’s payoff, i.e., the principal only cares about the duration of the game. We assume that both players are long-lived, which is different from the myopic agent settings in Ely (2017) and Renault et al. (2017). The myopic setting suggests that gradual revelation of information is optimal. In contrast, Orlov et al. (2020) suggest that having a long-lived agent changes the greedy nature of the principal’s equilibrium strategy, as future information disclosure can be used as an incentive device. Depending on the commitment power of the principal, the agent could be persuaded to wait for future information. In contrast with this literature, we are able to establish results even in settings in which the belief of the agent is private information.

A very recent set of papers have studied problems similar to ours, namely that of communi-

¹In 2021, the state of Maryland put into effect a tax on digital advertising revenues for those above \$100 millions revenue. See also the following articles by Paul Romer in the New York Times, (see <https://www.nytimes.com/...>, accessed on July 11, 2024), and Acemoglu and Johnson (2024) for advocating a flat 50% tax on advertising revenues.

cation when the objective of the principal is to lengthen engagement by the agent. These papers include [Knoepfle \(2020\)](#), [Hébert and Zhong \(2022\)](#), [Koh and Sanguanmoo \(2022\)](#), and [Koh et al. \(2024\)](#). [Knoepfle \(2020\)](#) studies a dynamic model of information transmission where multiple senders compete for the attention of a decision maker by strategically revealing information over time. In their set up the decision maker has a fixed marginal cost of staying each period. They find that in equilibrium senders use simple “all-or-nothing” strategies to reveal information, leading to full information transmission in minimal time. [Hébert and Zhong \(2022\)](#) study this problem assuming that the agent’s time cost is separable and linear in time and that the speed of learning is bounded above – which is modeled as a bound on the change in a generalized entropy function. They show that despite the existence of constraints on learning, the principal is able to reduce the value of the agent to her outside option.

Perhaps the closest paper to ours is that of [Koh and Sanguanmoo \(2022\)](#). In their paper, they assume that the cost of waiting for the agent is linear and separable over time of engagement while the payoff of the principal is a general function of the time-engagement. They provide general principles that determine the optimal mechanisms. There are two key distinctions between their environment and that of ours: 1. The use of linear time cost for the agent which implies that at each point in time, only the expected value of engagement matters to the agent; 2. We allow for mis-specified beliefs and agree-to-disagree which then allows us to talk about catering to the bias as well as communication in the presence of private information. As we show, an important implication of geometric or multiplicative discounting by both parties is that the commitment assumption of the principal is often binding. [Koh et al. \(2024\)](#) also study the general version of this problem (under belief agreements) and extend to the case when commitment is not binding.

Note that our result on the gradual information revelation is purely driven by the convexity of the agent’s time preferences. Similar forces occur in other settings where the precise shape of time preferences and their relationship between the parties involved in contracting matters. These include models of inspection ([Ball and Knoepfle \(2023\)](#)), and choice under uncertainty about the timing of rewards ([DeJarnette et al. \(2020\)](#)), among others. In our setup, the relative curvature of the payoffs, what we refer to as the marginal cost of engagement, and its evolution over time that determines the form of dynamics over time.

Finally, our paper is also related to the extensive literature on experimentation and contracting for discovery. Specifically, in our model, a time-varying Poisson bandit model as in [Keller et al. \(2005\)](#) and [Keller and Rady \(2010\)](#) arises endogenously.² Several papers have studied contracting in models of experimentation ([Guo \(2016\)](#), [Halac et al. \(2016\)](#), [Henry and Ottaviani \(2019\)](#), among others). This literature often takes as given the process for experimentation and studies

²Much of this literature focuses on strategic interactions among several experimenters – see also [Strulovici \(2010\)](#). While this would be relevant to our question, it is beyond the scope of current paper.

ways in which the conflict of interest between the parties can be handled either via delegation or persuasion or by the use of monetary incentives. In our setup, on the other hand, the principal is choosing the process of discovery guided by his interest in maximizing engagement.

The rest of the paper is organized as follows: Section 2 presents our general model, Section 3 provides a simple example, Section 4 analyzes the case of common priors, Section 5 explore scenarios with different priors, Section 6 examines the case of private beliefs, and Section 7 discusses various extensions of our framework.

2 The General Model

In this section, we present our general model. The model consists of a principal (referred to as "he") and an agent (referred to as "she"). The principal provides information to the agent about a payoff-relevant state over time, and the agent collects this information in order to take a final action. Time, denoted by t , is a continuous variable belonging to $[0, \infty)$. The payoff of the agent is given by

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

where $\omega \in \Omega = \{0, 1\}$ represents the underlying state, $a \in A$ is the action taken by the agent, and T is the time spent acquiring information. Throughout our analysis, we assume that the agent prefers to take the action sooner, therefore, δ_A and $\hat{u}(\omega, a)$ are both positive. Additionally, we assume that the payoff of the principal if the agent quits at T is $u_P(T) = \int_0^T e^{-\delta_P t} dt = \frac{1 - e^{-\delta_P T}}{\delta_P}$. In other words, the principal seeks to maximize the engagement time T . The fundamental disagreement between the principal and the agent is on the value of engagement: while the agent desires to acquire information as quickly as possible, the principal prefers longer engagements. As we will show, the key determinant of the principal's strategy is the relative patience of the principal to that of the agent or whether $\delta_A < \delta_P$ or $\delta_A \geq \delta_P$.

Figure 2 depicts the timing of the communication game in a short time period between t and $t + dt$. Conditional on the agent staying engaged up until t , at t the principal sends a signal $s_t \in S_t$ to the agent whose distribution depends on the state ω and the history of signal realizations in the past. This history is represented by the function $s^t = \{s_\tau\}_{\tau \in [0, t]}$. The agent, having observed s_t – and the history of public signals – decides whether to stay engaged after $t + dt$ or quit at $t + dt$ and choose $a_{t+dt} \in A$ to maximize her expected payoff $\mathbb{E}u(t + dt, \omega, a)$ where the expectation is taken with respect to her belief at $t + dt$ about ω .




Figure 2: The timing of the actions within a time interval $[t, t + dt]$

The strategies and learning process of the players can be explained as follows. The principal chooses an information structure: a mapping from the space of history realizations to probability distributions over signals at t . More formally, the principal's strategy is a quadruple $(S_\infty \times \Omega, \mathcal{F}, \mathbb{P}_P, \{\mathcal{F}_{t,P}\}_{t \in \mathbb{R}_+})$ where S_∞ is the set of history of signal realizations, i.e., each member is of the form s^∞ , Σ is a σ -algebra over $S_\infty \times \Omega$, \mathbb{P}_P is its associated probability measure from the principal's perspective, and finally, $\mathcal{F}_{t,P} \subset \mathcal{F}$ is a filtration, i.e., a family of increasing σ -algebras representing the information at time t . Here, the probability measure over the signal realizations is specified from the principal's perspective, which helps to identify the optimal strategies when the players have different prior beliefs over ω . We can think of the filtration $\mathcal{F}_{t,P}(\omega)$ as the σ -algebra induced by the function $\hat{s}^t(s^\infty, \omega)$. In words, this is the function that translates a complete history of the game (s^∞, ω) into history up to t . Let $\mathcal{F}_{t,A}$ be the filtration associated with the agent's information at time t . That is,

$$\mathcal{F}_{t,A} = \{S \subset S_\infty \mid \exists S' \in \mathcal{F}_t, s^\infty \in S \Leftrightarrow (s^\infty, 1) \in S' \text{ or } (s^\infty, 0) \in S'\} \quad (1)$$

The agent's strategy is a stopping time τ , associated with quitting, i.e., $T = \tau$, with respect to the filtration $\{\mathcal{F}_{t,A}\}_{t \in \mathbb{R}_+}$ together with a decision rule $a_\tau : S_\infty \rightarrow A$ which is progressively measurable with respect to the filtration $\mathcal{F}_{t,A}$.

For the learning process, we assume that the principal's and the agent's prior beliefs about ω are given by $\mu_P = \Pr_P(\omega = 1) = \mathbb{P}_P(S_\infty \times \{1\})$ and $\mu_A = \Pr_A(\omega = 1)$, respectively. We thus allow the priors to disagree, but this disagreement is common knowledge.³ Given the priors, the agent uses Bayes' rule to update her belief. Hence, her belief $\mu_A(t)$ is a progressively measurable stochastic process with respect to filtration $\mathcal{F}_{t,A}$ and follows:

$$\mu_A(t) = \mathbb{E}_A[\mathbf{1}[\omega = 1] \mid \mathcal{F}_{t,A}].$$

This conditional expectation operator maps members of $\mathcal{F}_{t,A}$ to $[0, 1]$. Note that since the agent and the principal may disagree on their prior belief about ω , the above expectation is taken with

³In Section 6, we study an extension of this model where the principal is not informed about the type of the agent and cannot send personalized signals.

respect to the agent's probability measure over $S_\infty \times \Omega$, \mathbb{P}_A . This probability measure can be constructed from those of the principal according to

$$\mathbb{P}_A(S) = \mu_A \mathbb{P}_P(S|S_\infty \times \{1\}) + (1 - \mu_A) \mathbb{P}_P(S|S_\infty \times \{0\})$$

This means that even if the agent and the principal disagree on their prior beliefs, they agree on the probability distribution chosen by the principal over signals conditional on the state ω .

Throughout our analysis, we assume that the principal is committed to his strategy, while the agent is not. Hence, a Perfect Bayesian Equilibrium is defined as follows:

Definition 1. A PBE of the game consists of a strategy profile for the principal $(S_\infty, \mathcal{F}, \mathbb{P}_P, \{\mathcal{F}_{t,P}\}_{t \geq 0})$ and a strategy profile for the agent (τ, a_τ) such that:

1. Given (τ, a_τ) , $(S_\infty, \mathcal{F}, \mathbb{P}_P, \{\mathcal{F}_{t,P}\}_{t \geq 0})$ maximizes principal's payoff

$$\mathbb{E}_P [1 - e^{-\delta_P T}]$$

2. Given $(S_\infty, \mathcal{F}, \mathbb{P}_P, \{\mathcal{F}_{t,P}\}_{t \geq 0})$, at any point in time t and conditional on not quitting, the agent chooses $(\tau, a_\tau)_{\tau \geq t}$ to maximize her payoff

$$\mathbb{E}_A [\delta_A e^{-\delta_A \tau} \hat{u}(\omega, a) | \mathcal{F}_{t,A}]$$

where $\mathcal{F}_{t,A}$ is derived from \mathcal{F}_t using (1).

2.1 Examples of the Environment

Our model can be applied to several settings where an expert provides information while maximizing engagement. Here, we discuss a few examples.

Consulting and Legal Services: An important application of our model is in the context of consulting and legal services, where compensation is often a function of the so-called "billable hours," i.e., the time spent on the project by the expert. This structure inherently creates large information asymmetries between the service provider and the customer. In the legal services domain, [Hadfield \(2007\)](#) and [Hadfield \(2022\)](#) argue that the American Bar Association's industry regulations on organizational structure, contract specification, and other areas create perverse incentives. A key inefficiency they highlight is the complexity and opaqueness of contracts, which often do not specify total costs, thereby incentivizing longer engagements. Our results can be interpreted as a theory linking the cost and benefit of engagement to the intertemporal preferences and beliefs of the parties involved in expert services contracts. Given that marginal benefit and

cost of engagement are key determinants of contract structure, and are often influenced by factors such as cost of capital and organizational structures, our model provides testable predictions. Specifically, it suggests examining the effectiveness of expert services in relation to these variables, offering a framework for empirical investigation of contract design and service provision in consulting and legal industries.

Market for News: For individual users, news consumption represents another common scenario of information acquisition with potential conflicts of interest, mirroring the dynamics in our model. News providers, ranging from traditional outlets like television and newspapers to online sources such as Google News, often rely on user engagement and advertisement as primary revenue sources. An important question in the literature on the economics of the media is that of the effect of private incentives on the quality of news and ultimately political competition – see, for example, [Strömberg \(2004\)](#). This issue becomes particularly salient in the context of personalized news delivery, where there is a possibility of news being *catered to the biases* of consumers. Our model can thus be used to shed light on the effect of these new media sources, personalized and catered media as opposed to mass media, on political outcomes. We specifically illustrate this application in Section 6 and 7.2.

Social media and the Internet: Social media has become an increasingly dominant platform for news dissemination. For instance, more than half of US adults now rely on the Internet as their primary source of health information ([Wang et al. \(2021\)](#)). The revenue model of social media platforms, heavily dependent on advertising and user engagement, aligns closely with the dynamics explored in our model, offering insights into the impact of this business approach on information distribution. Empirical evidence from social media algorithm designs appears to corroborate our theoretical predictions. [Levy \(2021\)](#) conducted a large field experiment on Facebook by offering participants to subscribe to outlets with random political attitudes. Their study shows that “... *Facebook’s algorithm is less likely to supply individuals with posts from counter-attitudinal outlets, conditional on individuals subscribing to them,*” suggesting that “*social media algorithms may limit exposure to counter-attitudinal news.*” This observation aligns with our model’s prediction of *catering the news to the bias*, where, in the presence of biased beliefs and personalized news, the information provider (principal) tends to prioritize information about states that the user (agent) considers more likely. Additionally, along the same lines of argument, [Allcott et al. \(2020\)](#) conduct an experiment on Facebook and show that deactivating Facebook significantly reduces the polarization of views on policy issues.

One can also view our model as one in which an advertiser or web designer decides how to deliver content to a user. Designers have a choice of how to deliver content that an interested user is wishing to learn: they can bombard a page with advertising and require a lot of scrolling until the end or randomly choose instances during a video to advertise, etc. According to IAB UK, an

industry body for digital advertising, 21% of impressions on the web were via the so-called “made for advertising” websites.⁴ While design experts often advise against requiring users to watch or scroll through large segments of advertising, providers still use such strategies. According to our model, this behavior can be explained by the relative patience of the provider to that of the user. This is in contrast to providing “... *salient information within a page’s initially viewable area* ..” together with “... *while placing the most important stuff on top, don’t forget to put a nice morsel at the very bottom*...”.⁵ Through the lens of our model, such behavior, interpreted as gradual revelation of information with the most likely information arriving first, is associated with a more patient principal.

2.2 Solving the Model

In this section, we describe the technical results required to characterize the solution of the principal’s optimal choice of communication.

First notice that similar to the formulation of the standard Bayesian Persuasion model of [Kamenica and Gentzkow \(2011\)](#), it is sufficient to describe the evolution of beliefs of the agent from the perspective of the principal, given that the agent and the principal could disagree on their priors. More specifically, an application of Bayes’ rule implies that

$$\frac{\mu_{t,A}(s^t)}{1 - \mu_{t,A}(s^t)} \frac{1 - \mu_{t,P}(s^t)}{\mu_{t,P}(s^t)} = \frac{\mu_A}{1 - \mu_A} \frac{1 - \mu_P}{\mu_P}$$

In words, since both the principal and the agent use the same signal structure to update their beliefs, the relative likelihood ratio of the agent’s belief to that of the principal remains constant over time. Let us define ℓ as this ratio

$$\ell = \frac{\mu_A}{1 - \mu_A} \frac{1 - \mu_P}{\mu_P}$$

Given ℓ and the belief of the principal, the belief of the agent can be calculated using the above and is given by

$$\mu_{t,A}(s^t) = \frac{\ell \mu_{t,P}(s^t)}{\ell \mu_{t,P}(s^t) + 1 - \mu_{t,P}(s^t)} \quad (2)$$

⁴See <https://www.iabuk.com/news-article/...> (accessed on July 11, 2024) for more guidelines on how to design a website for effective digital marketing. According to this article, “*Made for Advertising sites will often have unusual navigation and user journeys in order to maximize ad exposure. For example, they may tempt users to view ‘20 actors from the nineties, you won’t believe what number 17 looks like now!’ The content for an article like this will be laid out with one actor per page, so the user must click through 17 pages, being exposed to multiple ads on each page of the journey ...*”.

⁵Quoted from <https://www.nngroup.com/...>, a User Experience (UX) Research company (accessed on July 11, 2024).

An implication of this stationarity is that it is sufficient to describe the strategy of the principal via a stochastic process over his beliefs $\sigma_{t,P}(\mu_{t,P}|\mu_P^{t-})$ and rewrite the payoff of the agent in terms of that of the principal. Under this reformulation, we can refer to the history of signals as μ_P^t or history of beliefs for the principal. The strategy of the principal then is simply choosing a distribution over such histories given by $\sigma_P(\mu_P^t)$.

Subsequently, we can define the value of the agent upon exiting as a function of her belief, $v(\mu_A)$, as

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

It is evident that $v(\mu_A)$ is a convex function of μ_A . For convenience, we make the following assumption about $v(\mu_A)$:

Assumption 1. *The payoff function $v(\mu_A)$ is strictly convex, differentiable, and symmetric around $\mu_A = 1/2$.*

We should note that this assumption is rather innocuous and allows us to conveniently characterize optimal strategies of the principal via first-order conditions. One can always normalize the payoffs of the principal and the agent to make v symmetric. Moreover, any convex symmetric function can be approximated by a sequence of strictly convex and differentiable functions. Assumption 1 allows us to take derivatives, which streamlines the analysis.

Our first result concerns the beliefs of the agent upon exit.

Lemma 1. *Suppose that Assumption 1 holds. If the agent exits after the belief history μ_P^{t-} , then $\mu_P(t) = \mathbb{E}_P[\omega|\mu_P^{t-}] \in \{0, 1\}$ almost surely.*

The intuition behind this result is straightforward. Consider a scenario where $\mu_P(t) \notin \{0, 1\}$ for a positive measure of histories in which the agent exits. In this case, the principal can improve the information revelation strategy by splitting the signal into two fully revealing signals, $\{0, 1\}$, with probabilities $1 - \mu_P(t)$ and $\mu_P(t)$ respectively. This modification is a mean preserving spread of the beliefs. Given that $v(\cdot)$ is strictly convex (Assumption 1), this spread increases the agent's expected payoff without inducing earlier exits. Hence, such a strategy can be combined with an initial period of no information revelation at the beginning of the game, thus inducing a profitable deviation for the principal.

Given our reformulation of the problem, we can apply the Carathéodory theorem and show that three signals are enough for each period. This together with Lemma 1 implies that these signals are given by $\Omega \cup \{\text{No News}\}$. Therefore, in each period, the agent either quits with full information or updates her belief based on the fact that the state is not revealed.⁶ As a result, we

⁶This feature of optimal communication is similar to models of Poisson experimentation a la Keller and Rady (2010) where no news leads to a gradual change in beliefs.

can summarize the strategy of the principal by the use of two distribution functions:

$$\begin{aligned}
G_{P,1}(t) &= \mathbb{P}_P(\text{exit} \geq t, \omega = 1) \\
G_{P,0}(t) &= \mathbb{P}_P(\text{exit} \geq t, \omega = 0) \\
\mu_P(t) &= \mathbb{P}_P(\{\omega = 1\} | \text{stay until } t) \\
&= \frac{G_{P,1}(t)}{G_{P,1}(t) + G_{P,0}(t)} = \frac{G_{P,1}(t)}{G_P(t)}
\end{aligned}$$

where $G_{P,1}(0) = \mu_P = 1 - G_{P,0}(0)$, and $G_P(t) = G_{P,1}(t) + G_{P,0}(t)$. Note that $G_{P,\omega}(t)$'s are decumulative distribution functions and are thus decreasing over time. We will refer to $G_{P,\omega}(t)$'s as *engagement probability functions*.

This simplification of the problem allows us to rewrite the payoffs of the agent and the principal in simpler forms:

$$\begin{aligned}
u_P &= - \int_0^\infty \frac{1 - e^{-\delta_P t}}{\delta_P} dG_P(t) \\
u_{A,t} &= - \int_t^\infty e^{-\delta_A(s-t)} v(1) \frac{d(\ell G_{P,1}(s) + G_{P,0}(s))}{\ell G_{P,1}(t) + G_{P,0}(t)}
\end{aligned}$$

In the above, the negative signs represent the fact that $G_{p,\omega}(t)$ is a decumulative probability function. Moreover, the probability used for the agent is adjusted to account for their difference in their priors while we have used the fact that $v(0) = v(1)$. Additionally, $u_{A,t}$ is the continuation payoff of the agent in the event that she does not exit until t . Using integration by parts, we can write these payoffs as

$$\begin{aligned}
u_P &= \int_0^\infty e^{-\delta_P t} G_P(t) dt \\
u_{A,t} &= v(1) - \frac{\delta_A v(1) \int_t^\infty e^{-\delta_A(s-t)} (\ell G_{P,1}(s) + G_{P,0}(s)) ds}{\ell G_{P,1}(t) + G_{P,0}(t)}
\end{aligned}$$

The above calculations imply that the principal's optimal information provision problem, which yields $V(\mu_P, \ell)$, is given by:

$$V(\mu_P, \ell) = \max_{G_{P,0}, G_{P,1}} \int_0^\infty e^{-\delta_P t} [G_{P,1}(t) + G_{P,0}(t)] dt \quad (\text{P})$$

subject to

$$\begin{aligned}
v(1)G_A(t) - v(1)\delta_A \int_t^\infty G_A(s) e^{-\delta_A(s-t)} ds &\geq G_A(t) v\left(\frac{\ell G_{P,1}(t)}{\ell G_{P,1}(t) + G_{P,0}(t)}\right), \forall t \quad (3) \\
G_{P,\omega}(t) &: \text{non-increasing}, \omega \in \{0, 1\} \\
G_{P,1}(0) &= 1 - G_{P,0}(0) = \mu_P
\end{aligned}$$

where in the above, $G_A(t) = \ell G_{P,1}(t) + G_{P,0}(t)$ is the probability of engagement until t from the agent's perspective and is used for brevity.

In the above problem, when the agent and the principal agree on their prior, i.e., $\ell = 1$, the gain from information revelation in each state is given by $v(1) - v(\mu_A)$. Given our symmetry assumption, this is maximized at $\mu_A = 1/2$, i.e., when agent is maximally uncertain. Thus, depending on the evolution of Marginal Cost of Engagement (MCE), the principal wishes to maximize the amount of time spent at this belief or as close as possible to this level.

In contrast, when $\ell < 1$, the agent is more optimistic about the state $\omega = 0$ relative to the principal. In this case, while the gain from revelation from the agent's perspective is still maximized at $\mu_A = 1/2$, since the principal disagrees with the agent on her beliefs, this maximum gain does not necessarily lead to longer engagement from the principal's perspective.

Finally, we should note that given Assumption 1, the function $(\ell G_{P,1} + G_{P,0}) v\left(\frac{\ell G_{P,1}}{\ell G_{P,1} + G_{P,0}}\right)$ is strictly convex. As a result, the constraint set in (P) is convex. Together with the fact that the objective is linear in $G_{P,1}(t)$ and $G_{P,0}(t)$, the standard results from convex optimization (See, for example, Luenberger (1997), sections 8 and 9) apply.⁷ Therefore, we can set \mathcal{G} to be the set of decumulative distribution functions $(G_{P,1}, G_{P,0})$ with $G_{P,1}(0) = 1 - G_{P,0}(0) = \mu_P$. Then, we can use integration by parts to write the Lagrangian associated with P as:

$$\begin{aligned}
\mathcal{L} &= \int_0^\infty e^{-\delta_P t} [G_{P,1}(t) + G_{P,0}(t)] dt \\
&+ \int_0^\infty e^{-\delta_P t} G_A(t) \left[v(1) - v\left(\frac{\ell G_{P,1}(t)}{\ell G_{P,1}(t) + G_{P,0}(t)}\right) \right] d\Lambda(t) \\
&- \delta_A v(1) \int_0^\infty G_A(t) e^{-\delta_A t} [\Lambda(t) - \Lambda(0)] dt
\end{aligned}$$

where $\Lambda(t)$ is a weakly increasing function which is constant at t whenever the incentive constraint is slack.

With this reformulation, $d\Lambda(t)$ is the value of gains from information revelation, and the last

⁷The Luenberger's results cannot readily apply because the constraint set does not have a non-empty interior. In the appendix, we first show that when time is finite, the constraint set has a non-empty interior. We then send the time horizon to infinity, show that multipliers converge and use Berge's maximum Theorem.

term captures the benefits of revelation at t on all previous periods' incentives. If $\partial\mathcal{L}(G_1, G_0; h)$ is the Gateau derivative of \mathcal{L} along the direction $h \in \mathcal{G}$, we can use the result on convex programming in [Luenberger \(1997\)](#) and state the following.

Theorem 1. *If $(G_{P,1}^*, G_{P,0}^*)$ is the solution to P , then there exists an increasing function $\Lambda(t)$ such that*

$$\begin{aligned}\partial\mathcal{L}(G_{P,1}^*, G_{P,0}^*; G_{P,1}^*, G_{P,0}^*) &= 0 \\ \partial\mathcal{L}(G_{P,1}^*, G_{P,0}^*; h) &\leq 0, \forall h \in \mathcal{G}.\end{aligned}$$

Conversely, if for some increasing $\Lambda(t)$, $(G_{P,1}^, G_{P,0}^*)$ satisfies the above conditions, then it is the solution to P .*

We should note that the above does not imply that the monotonicity constraints on G_ω 's are not binding. Rather, it implies that sometimes ironing is needed. Our solution technique is similar to a strand of the literature on Bayesian Persuasion that uses linear and convex programming techniques in infinite-dimensional spaces. See [Dworczak and Martini \(2019\)](#), [Kleiner et al. \(2021\)](#), and [Saeedi and Shourideh \(2023\)](#) among many others. Note that since our constraint set is strictly convex, it is often the case that the solution to the optimization problem P is unique.

3 A Simplified Example

In order to build intuition for our first set of results, we illustrate some of the basic ideas in a simple version of the model discussed above, where the principal can only choose to reveal the state or not.

Suppose that the value to the agent of engagement until T is given by

$$\frac{1}{2}e^{-\delta_A T} (1 + \mathbf{1}[\text{informed by } T]),$$

where δ_A is the agent's discount rate. The value to the principal is given by T .

The principal can commit to a revelation strategy, while the agent cannot commit to an engagement strategy. In this example, we have abstracted from the details of the agent's value of information. Additionally, we have only allowed the principal to reveal the state or not.

Given that the principal is able to commit to any, possibly random, revelation strategy, we can represent his strategy as $F(T)$, an increasing function of time that represents the cumulative




Figure 3: A mean preserving spread of revelation time

probability of revelation until time T . Hence, the principal's problem can be written as

$$\begin{aligned} \max_{F \in \mathcal{F}} \mathbb{E}_F T \\ \text{s.t. } \mathbb{E}_F [e^{-\delta_A T} | T > t] \geq \frac{1}{2} e^{-\delta_A t}, \forall t \geq 0 \end{aligned} \quad (4)$$

where \mathcal{F} is the set of increasing functions over $\mathbb{R}_+ \cup \{0\}$ with $F(0) = 0, \lim_{T \rightarrow \infty} F(T) = 1$. Moreover, the constraint is the incentive constraint of the agent that remains uninformed at t and needs to be incentivized to stay engaged.

To understand the forces at play in solving the principal's problem, consider the naive strategy of delivering the information at a predetermined period \hat{T} . For this to be incentive compatible for the agent, we must have

$$e^{-\delta_A \hat{T}} \geq \frac{1}{2} \Rightarrow \hat{T} \leq \frac{\log 2}{\delta_A}.$$

Therefore, the optimal time is $\hat{T} = \frac{\log 2}{\delta_A}$. The convexity of the function $e^{-\delta_A T}$ in T implies that a small mean-preserving spread of \hat{T} increases the expected value for the agent at all times while keeping $\mathbb{E}(T)$ unchanged. Moreover, if the spread is small enough, the agent is willing to remain engaged. This implies that it can be accompanied by an increase in all values of T , consequently increasing the principal's value. This is depicted in Figure 3.

Such a perturbation is always feasible when the incentive constraint is slack for any interval of time t , as it can be implemented within that interval. Consequently, in the optimal solution, the incentive constraint must bind for all values of t . A multiplication of 4 by $1 - F(t)$ and differentiation with respect to t yields:

$$-e^{-\delta_A t} f(t) = -\frac{1}{2} e^{-\delta_A t} f(t) - \delta e^{-\delta_A t} (1 - F(t))$$

which in turn implies that $F(t) = 1 - e^{-\delta_A t}$. In other words, the information is revealed according to a Poisson process with arrival rate δ_A .

The key property driving this result is the convexity of the agent's payoff function with respect to the principal's payoff. To see this, consider the case in which the principal discounts future engagements at rate δ_P , where $0 \leq \delta_A < \delta_P$, i.e., the principal's payoff is given by $u_P = \int_0^T e^{-\delta_P t} dt$, while the agent's payoff remains as before. We can write the agent's payoff as a function of the payoff of the principal:

$$u_P = \frac{1 - e^{-\delta_P T}}{\delta_P} \rightarrow T = \frac{-\log(1 - \delta_P u_P)}{\delta_P} \rightarrow u_A = \frac{1}{2} (1 - \delta_P u_P)^{\frac{\delta_A}{\delta_P}} (1 + \mathbf{1}[\text{informed by } T]).$$

Since u_P is an increasing function T , choosing a distribution of T , $F(T)$, is equivalent to choosing a distribution of u_P , $G(u_P)$. Thus, the principal's problem of delivering the information is given by:

$$\begin{aligned} & \max_{G \in \mathcal{G}} \mathbb{E}_G [u_P], \\ & \text{s.t. } \mathbb{E}_G \left[(1 - \delta_P u_P)^{\frac{\delta_A}{\delta_P}} \mid u_P > v \right] \geq \frac{1}{2} (1 - \delta_P v)^{\frac{\delta_A}{\delta_P}}, \forall v \in [0, 1]. \end{aligned}$$

The key distinction between this case and the previous one lies in the relationship between δ_P and δ_A . Since $\delta_P > \delta_A$, the function $(1 - \delta_P u_P)^{\frac{\delta_A}{\delta_P}}$ is decreasing and concave. As a result, starting from any distribution of u_P , a change to revelation at \hat{T} given by

$$\frac{1 - e^{-\delta_P \hat{T}}}{\delta_P} = \mathbb{E}_G [u_P],$$

does not change the payoff of the principal while it increases the payoff of the agent at all times. This implies that the best strategy of the principal is to reveal the information with certainty at $\hat{T} = -\log 2 / \delta_A$. This perturbation is depicted in Figure 4 for a two point distribution of u_P .

The key distinction between these the two scenarios lies in relative changes in the cost of engagement for the agent relative to its marginal benefit for the principal. The value of extending engagement from T to $T + dT$ for the principal is $e^{-\delta_P T} dT$, while the cost of postponing revelation for the agent is $\delta_A e^{-\delta_A T} dT$. We define the ratio of these costs to benefits as the *Marginal Cost of Engagement* (MCE):

$$\text{MCE} = \delta_A e^{-(\delta_A - \delta_P)T}.$$

In the first case, when $\delta_P = 0 < \delta_A$, the MCE decreases over time. Loosely speaking, the




Figure 4: A mean-preserving contraction of Principal's payoff

agent becomes relatively more patient as engagement duration increases. Thus, the principal benefits from divulging some information early on to incentivize the agent to wait until later periods when maintaining engagement becomes easier.

In contrast, in the second case, when $0 < \delta_A < \delta_P$, the MCE increases over time, suggesting that the agent becomes relatively less patient. Hence, early partial revelation of information is not beneficial for the principal since to keep the agent engaged later, information must be released faster. In the rest of the paper, we illustrate how these two strategies can be applied to the more complex information and utility structure.

4 Communication under Belief Agreement

In this section, we analyze optimal communication strategies when the agent and the principal share the same prior. This implies that $\ell = 1$ and that the agent does not have any bias towards any state relative to the principal. This shared prior setting serves as a baseline for our subsequent extensions.

More Patient Agent

When the agent is more patient than the principal ($\delta_P > \delta_A$) and they share the same prior, the optimal disclosure strategy takes a simple form. Similar to the example in section 3, in this case, optimal information revelation occurs abruptly. This result is formalized in the following proposition:

Proposition 1. *When $\delta_P > \delta_A$ and $\mu_P = \mu_A$, the optimal solution of the principal's problem is*

$$\begin{aligned} G_{A,1}(t) &= \mu_A \mathbf{1}[t < t^*] \\ G_{A,0}(t) &= (1 - \mu_A) \mathbf{1}[t < t^*] \\ e^{-\delta_A t^*} v(1) &= v(\mu_A) \end{aligned}$$

In words, Proposition 1 states that optimal information provision, in the case of concave discounting or when the Marginal Cost of Engagement (MCE) increases over time, leads to abrupt revelation of information at t^* . At $t = 0$, the agent is indifferent between this revelation at t^* , and no revelation at all or not participating. The intuition behind this result is that when MCE increases with time, any gradual information revelation before t^* must be accompanied by even faster revelation later to maintain the agent's engagement. Hence, optimal revelation is abrupt.

More Patient Principal

When the principal is more patient than the agent ($\delta_P < \delta_A$), the endogeneity of beliefs leads to more complex dynamics in the optimal provision of information compared to the simple example. As mentioned above, without belief disagreement, the gain from information revelation is maximized at $\mu_A = 1/2$. However, the initial belief may differ from half, i.e., $\mu_A(0) \neq 1/2$. In such cases, the principal aims to shift beliefs towards $1/2$. In this case, the principal wants to shift the beliefs toward half. Additionally, this transition should be smooth due to the strict convexity of the agent's payoff in $\mu_{t,A}$ and the agent's preference for gradual belief changes.

We show that it is optimal to first reveal the state towards which initial beliefs are closer. For instance, if $\mu_A(0) > 1/2$, $\omega = 1$ should be revealed first. This revelation continues until beliefs reach $\mu_{t,A} = 1/2$. Once the agent reaches this point of maximum uncertainty (and thus maximum engagement), the principal optimally reveals both states at a rate that keeps the agent indifferent between engaging and quitting. This strategy ensures that when the state is not revealed, beliefs remain at the point of maximum engagement where the gain from revelation is highest. We formalize these findings in the following proposition:

Proposition 2. *When $\delta_A > \delta_P$ and the priors of the agent and principal agree, the optimal solution to (P) has two phases:*

1. *A phase, $0 \leq t < t^*$, where only state one (zero) is revealed when $\mu_P > 1/2$ ($\mu_P < 1/2$) and*

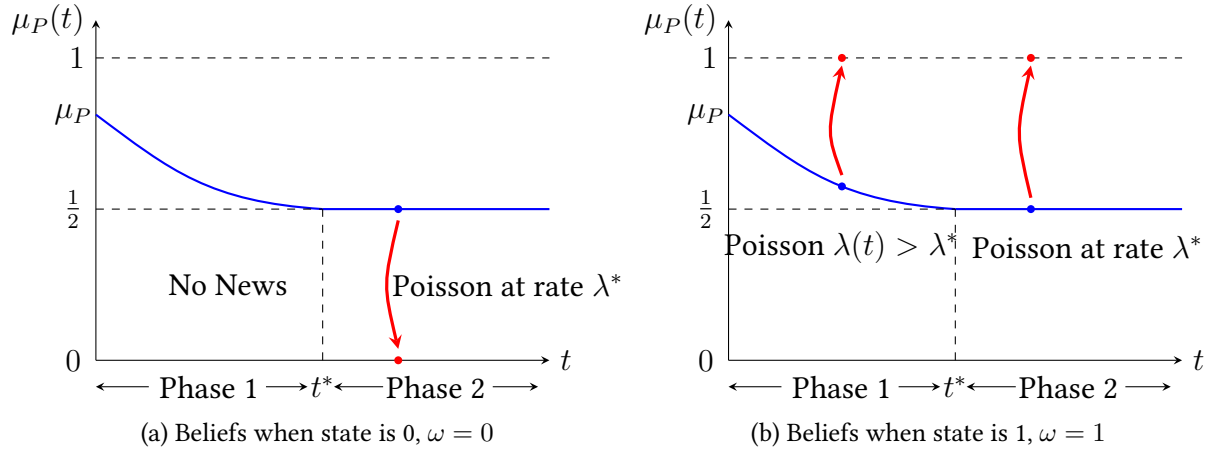


Figure 5: Evolution of beliefs in each state when the agent is more impatient and $\mu_P > 1/2$

beliefs upon staying satisfy

$$\begin{aligned} \mu_P(0) > 1/2 : -\delta_A &= \frac{\mu'_P(t)}{v(\mu_P(t))} \left[\frac{v(1) - v(\mu_P(t))}{1 - \mu_P(t)} - v'(\mu_P(t)) \right] \\ \mu_P(0) < 1/2 : \delta_A &= \frac{\mu'_P(t)}{v(\mu_P(t))} \left[\frac{v(1) - v(\mu_P(t))}{\mu_P(t)} + v'(\mu_P(t)) \right] \end{aligned}$$

2. A phase, $t \geq t^*$, where both states are revealed at rate λ^* with

$$\lambda^* = \delta_A \frac{v(1/2)}{v(1) - v(1/2)}$$

and beliefs are at the point of maximum uncertainty, $\mu_P = 1/2$.

In this case, phase 1 consists of revealing the state towards which beliefs are more optimistic, at a time-varying poisson rate. For example, if $\mu_P > 1/2$ or state $\omega = 1$ is more likely, during phase 1, there is only a possibility of receiving a signal if the true state is 1. If no signal was observed, both the agent and the principal update their beliefs downwards. This process continues until beliefs reach $1/2$, the point of maximum uncertainty. At this point, both states are revealed at a rate λ^* . Figure 5 illustrates this process.

The rate of information revelation and the length of phase 1 depends on the curvature of the value function $v(\cdot)$. The higher the curvature, the longer phase 1 becomes. This is due to the strong preferences for smoothing of beliefs which implies that in order to keep the agent engaged, it is optimal for beliefs to decline more gradually. Consequently, this lengthens phase 1, albeit at the cost of reducing the principal's value.

Proposition 2 highlights one of the key mechanisms in our model: the principal wishes to reach the state of *maximum engagement* as fast as possible by revealing information. Under belief agreement, this state is given by the state of maximum uncertainty or $\mu_A = 1/2$.

It’s worth emphasizing that the value of δ_P does not affect how the revelation happens beyond being gradual or abrupt. This invariance is partly due to the principal’s ability to capture all surplus in both scenarios.

5 Catering to the Bias

In this section, we illustrate the “catering to the bias” force that emerges when the principal and the agent disagree on their prior beliefs. As we discussed in Section 2.2, the problem of optimal communication under belief disagreement can be formulated in terms of the distribution of beliefs from the principal’s perspective. While our primary focus is on interpreting the model with differing priors, this framework can be equivalently interpreted as a scenario where the principal and agent share the same prior, but the agent exhibits a preference for information acquisition in one state over the other—a form of confirmation bias.

More Patient Agent

We begin our analysis of optimal communication under belief disagreement by examining the case where the agent is more patient than the principal ($\delta_A < \delta_P$). Recall that with shared priors, an impatient principal would opt to reveal all information abruptly. However, we demonstrate that this strategy changes when priors differ.

Consider the scenario where $\ell < 1$, indicating that the agent believes $\omega = 0$ is more likely than the principal does.⁸ In this case, the agent puts a higher weight on $G_{P,0}(t)$ than principal, as shown in Equation (3). This implies that the agent attributes greater value to information about $\omega = 0$ than the cost perceived by the principal. The principal, believing $\omega = 0$ to be relatively less likely, perceives this cost as low and considers it unlikely to result in a high probability of exit. This disparity creates an incentive for the principal to first reveal the state towards which the agent is biased.

However, despite this logic, the principal remains less patient than the agent, and forces favoring postponed revelation persist in this case. This tension results in an optimal strategy comprising an initial phase of no information revelation, followed by a phase of catering to the bias. The following proposition characterizes the optimal information revelation strategy in this case:

⁸The case of $\ell > 1$ is the mirror of this case.

Proposition 3. *Suppose that $\delta_P > \delta_A$ and $\ell < 1$. Then the solution to (P) consists of two phases and two instantaneous revelations times:*

1. *In phase 1, $t \in [0, t_1^*)$, no information is revealed.*
2. *At t_1^* , $\omega = 0$ is revealed with a positive probability.*
3. *In phase 2, $t \in [t_1^*, t_2^*]$, $\omega = 0$ is revealed gradually according to a Poisson process at a rate such that the agent's beliefs satisfy the following ODE:*

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

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

Moreover, the length of phase 2 is always positive, i.e., $t_1^ < t_2^*$.*

The dynamics implied by Proposition 3 are depicted in Figure 6. The optimal communication strategy consists of two distinct phases: 1. An initial phase where no information is provided, and thus beliefs stay unchanged. 2. A subsequent phase that begins with an abrupt, but not certain, revelation of $\omega = 0$, followed by a gradual revelation of $\omega = 0$ according to a Poisson process at a rate determined by the ODE specified in the proposition. This phase concludes when the belief reaches 1, at which point the true state is fully revealed. Furthermore, in the second phase, news arrive increasingly faster over time as the poisson revelation rate increases over time.

Figure 6 illustrates these dynamics under two scenarios: Sub-figure (a) depicts the case where the true state is $\omega = 0$. In this scenario, information is revealed either abruptly at the end of phase 1 or gradually revealed during phase 2. Sub-figure (b) shows the case where the true state is $\omega = 1$. Here, all information is abruptly revealed at t_2^* .

Proposition 3 illustrates the “*catering-to-the-bias*” force present in our model. Due to belief disagreement, the optimal information provision problem becomes equivalent to one in which the principal prefers the agent to stay engaged in state $\omega = 1$ more than in state $\omega = 0$. With bias in beliefs, the solution is to reveal the state $\omega = 0$ gradually. This gradual revelation of $\omega = 0$ serves as a low cost method of keeping the agent engaged when the true state is $\omega = 1$, exploiting the agent’s biased beliefs towards $\omega = 0$. This allows the principal to postpone the revelation of $\omega = 1$ as late as possible.

While the statement that $t_1^* < t_2^*$ implies that catering to the bias always occurs under the optimal engagement policy, it does not rule out the possibility of phase 1 having zero length (i.e.,

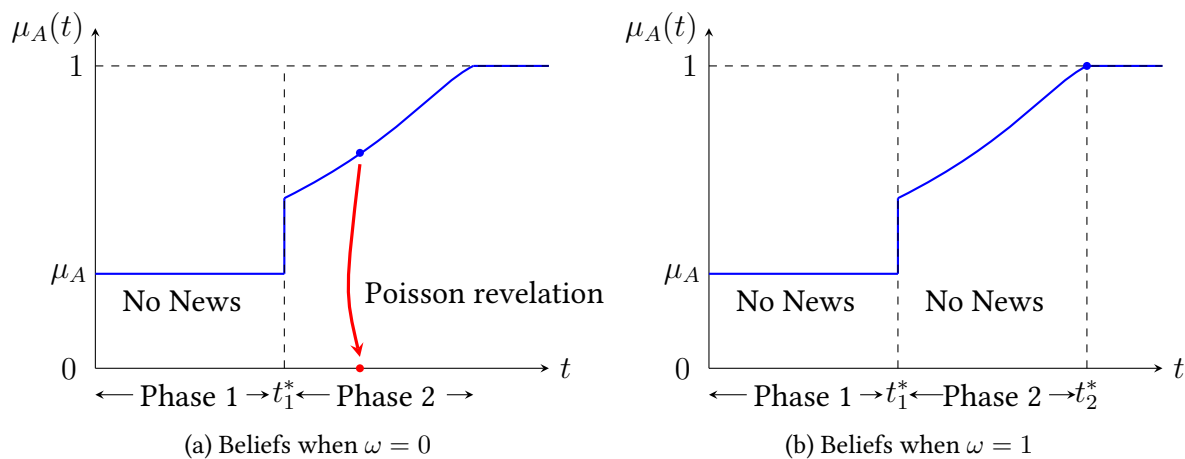


Figure 6: Catering to the bias with an Impatient Principal

$t_1^* = 0$). The following corollary provides conditions under which catering to the bias becomes so extreme that it eliminates the no-news phase entirely. Notably, it indicates that phase 1 is more likely to occur when belief disagreement is small.

Corollary 1. *Given the parameters of the model, there exists $\underline{\mu}_P$ and $\bar{\mu}_P$ such that when $\mu_A \notin [\underline{\mu}_P, \bar{\mu}_P]$, phase 1 has zero length.*

More Patient Principal

Suppose now that the principal is more patient than the agent ($\delta_A > \delta_P$). Recall that under belief agreement, the principal reveals the state gradually, with the more likely state being revealed first, to maximize the time the agent spends at a belief with maximum uncertainty. With belief disagreement, the spirit of the strategy of driving the agent to a stationary belief stays the same, yet it takes a form of “catering to the bias” as discussed earlier and this stationary belief does not maximize engagement.

More specifically, we start by showing that the point of maximum uncertainty ($\mu_A = 1/2$) is not going to be steady state anymore due to the disagreement in beliefs. To see this, suppose that $\mu_A = 1/2$ and $\mu_P < 1/2$, for this point to be steady state, both states should be revealed at a same rate λ , for the agent’s belief to remain at $\mu_A = 1/2$. The variable λ is chosen to keep the agent indifferent between staying engaged and quitting at any point in time. Next, we show a profitable deviation for the principal.

Consider an alternative strategy where the principal reveals only $\omega = 1$ for a short interval of length dt at a rate 2λ , before reverting to symmetric Poisson revelation. This brief asymmetric revelation lowers the agent’s belief slightly. While the symmetric Poisson rate should increase to maintain indifference, the change in the agent’s payoff is only second-order in dt due to $v'(1/2) =$

0. Thus, to first order, λ in the symmetric revelation case remains unchanged. However, since the principal believes the probability of $\omega = 1$ is $\mu_P < 1/2$, revealing only disclosing $\omega = 1$ for the period between 0 and dt has first order gain for the principal. This implies that keeping the agent at $\mu_A = 1/2$ is not optimal for the principal.

We can continue with the above heuristic argument, to find a steady state pair of (μ_A^*, μ_P^*) . To keep the agent indifferent between quitting and staying, her payoff from quitting at any point in time must equal $v(\mu_A)$. Under this strategy, the agent's payoff from staying engaged is:

$$\frac{\lambda}{\lambda + \delta_A} v(1) = v(\mu_A). \quad (5)$$

The principal's payoff from this strategy is $1/(\lambda + \delta_P)$. Now, consider a deviation where the principal reveals state $\omega = 1$ at rate q over the time period $[0, \Delta]$ and reverts to gradual revelation of both states at rate $\lambda + d\lambda$ after $t = \Delta$. The agent's belief at the end of period Δ is:

$$\mu_A + d\mu_A = \mu_A - \mu_A(1 - \mu_A)q\Delta + o(\Delta^2).$$

Thus, the principal can choose $\lambda + d\lambda$ such that:

$$\frac{d\lambda}{(\lambda + \delta_A)^2} v(1) = v'(\mu_A) d\mu_A.$$

The agent's payoff from this strategy, up to a second-order approximation, is:

$$q\mu_A [v(1) - v(\mu_A)] \Delta - \delta_A v(\mu_A) \Delta + \frac{d\lambda}{(\lambda + \delta_A)^2} v(1).$$

Here, the first term represents the gain from revealing $\omega = 1$ with probability $q\Delta\mu_A$ over the time interval $[0, \Delta]$. The second term is the change in value from discounting, while the last term is the change in the agent's value due to the belief change. Equating this to $v(\mu_A)$, implies a lower bound on q :

$$q = \frac{\delta_A}{\mu_A} \frac{v(\mu_A)}{v(1) - v(\mu_A) - (1 - \mu_A) v'(\mu_A)}.$$

The change in the principal's payoff from this deviation, to first-order approximation, is:

$$-\frac{1}{\lambda + \delta_P} \mu_P q \Delta - \frac{d\lambda}{(\lambda + \delta_P)^2} + \frac{\lambda}{\lambda + \delta_P} \Delta.$$

The first term represents losses from revealing $\omega = 1$ over $[0, \Delta]$. The second term is the change in payoff from postponing symmetric revelation at rate $\lambda + d\lambda$, and the last term is the gain from




Figure 7: Catering to the Bias with a Patient Principal

abandoning symmetric revelation at rate λ over the interval $[0, \Delta]$. While revelation of $\omega = 1$ at rate q benefits the agent, it is costly for the principal, with this cost decreasing for lower μ_P . Thus, for sufficiently high μ_P , this deviation is not profitable, defining a threshold μ_P^* :

$$\mu_P^*(\mu_A) = \mu_A + \frac{\mu_A(1 - \mu_A)v'(\mu_A)}{v(\mu_A) + \frac{\delta_P}{\delta_A - \delta_P}v(1)}.$$

For $\mu_P < \mu_P^*(\mu_A)$, an initial phase of revealing $\omega = 1$ is optimal with both $\mu_A(t)$ and $\mu_P(t)$ decreasing until the above threshold is reached. This is depicted in Figure 7. As time progresses, the likelihood ratios, $\ell = \mu_A / (1 - \mu_A) \times (1 - \mu_P) / \mu_P$, remains constant. Eventually, the above holds, at which point both states are revealed symmetrically and the beliefs conditional on not quitting remain constant. In contrast, when $\mu_P > \mu_P^*(\mu_A)$, the opposite happens where $\omega = 0$ is revealed first until the steady state levels are reached.

The threshold defined by the above is an increasing function of μ_A provided that its value is between 0 and 1. For high values of μ_A the threshold can exceed 1, and for low values, it can be less than 0. Thus, we can define the threshold as follows:

$$\mu_P^*(\mu_A) = \max \left\{ \min \left\{ \mu_A + \frac{\mu_A(1 - \mu_A)v'(\mu_A)}{v(\mu_A) + \frac{\delta_P}{\delta_A - \delta_P}v(1)}, 1 \right\}, 0 \right\} \quad (6)$$

Given this discussion, we have the following result:

Proposition 4. *Suppose that $\delta_A > \delta_P$. Then there exists a threshold $\mu_P^*(\mu_A)$ given by (6) such that the solution to (P) consists of two phases:*

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

$$\delta_A = \frac{\mu'_A(t) v(1) - v(\mu_A(t)) + v'(\mu_A(t)) \mu_A(t)}{\mu_A(t) v(\mu_A(t))} \quad (7)$$

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

$$-\delta_A = \frac{\mu'_A(t) v(1) - v(\mu_A(t)) - v'(\mu_A(t)) (1 - \mu_A(t))}{1 - \mu_A(t) v(\mu_A(t))} \quad (8)$$

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

Proposition 4 and Figure 7 highlight the key property of the catering to the bias effect. Specifically, for very high and very low values of μ_A where μ_P^* is 1 or 0, respectively, our result implies that irrespective of the principal's initial belief, μ_P , it is always optimal to first reveal the state towards which the agent is biased.

We should note that catering to the bias can become extreme, such that no information about one of the states ever arrives before the end of communication. This occurs when the function $\mu_P^*(\mu_A)$ at $\mu_A = 0, 1$ is not flat. It implies that for high or low enough values of ℓ , the constant ℓ curves – loci of points with the same relative likelihood ratio – only intersect $\mu_P^*(\mu_A)$ at $\mu_A = 0, 1$. This implies that catering to the bias occurs during the entire communication process. We thus have the following proposition:

Proposition 5. *Suppose that $\delta_A v(0) + (\delta_A - \delta_P) v'(0) > 0$. Then for all values of initial beliefs μ_A, μ_P such that $\ell = \frac{\mu_A}{1 - \mu_A} \frac{1 - \mu_P}{\mu_P} \leq 1 + \frac{(\delta_A - \delta_P) v'(0)}{\delta_A v(0)}$, only $\omega = 1$ is revealed during the course of communication. Accordingly, when $\frac{1}{\ell} \geq 1 + \frac{(\delta_A - \delta_P) v'(0)}{\delta_A v(0)}$, only $\omega = 0$ is revealed during the course of communication.*

6 Non-Personalized vs. Personalized News

In this section, we extend our general model to investigate the impact of news personalization on information transmission. Specifically, we consider a scenario where the agent's belief is privately known and unobserved by the principal.

We analyze a variant of our model where the agent'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 . Let $\alpha^j = \Pr(\mu_A^j)$, $j \in \{L, H\}$ denote the probability of the agent being of type j . The agent discounts future payoffs at rate $\delta_A > 0$, while the principal's discount rate is $\delta_P < \delta_A$, thus the principal is more patient than the agent.

Our goal in this section is to compare the personalized benchmark of previous sections with a non-personalized case where all communications are public and the agent type is privately known. As we will show, while due to the presence of private information, engagement has to necessarily be shorter, the principal can trade this off with quality of information and thus still capture the full surplus from the agent.

General Structure of Optimal Communication

Our key departure from the personalized news case is that communication between the principal and agent is public and is thus not fully targeted. This assumption implies that we can formulate optimal communication simply as a pure recommendation mechanism.

At first glance, this appears to be a complicated game with persistent private information. Specifically, if an equilibrium calls for an agent of type j to quit at t , upon deviation and staying, she persistently holds superior information over the principal, and her future incentives need to be respected. However, due to the particular principal's objective, such deviations need not be considered. That is because if staying instead of quitting is profitable, the principal's initial strategy is suboptimal since the principal prefers staying to quitting. Thus, we can describe the best equilibrium for the principal as a recommendation strategy that only needs to be obedient. This leads to the following lemma:

Lemma 2. *The best equilibrium of the game from the principal's perspective can be described by a signal structure $(S_\infty, \mathcal{F}, \mathbb{P}_P, \{\mathcal{F}_{t,P}\}_{t \geq 0})$ 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 $v(\mu_A^j(s^t))$,*
2. *If type j is recommended to stay following signal history s^t , the value of staying engaged for j is not lower than $v(\mu_A^j(s^t))$,*

where $\mu_A^j(s^t)$ is the agent of type j 's belief induced by the signal structure $(S_\infty, \mathcal{F}, \mathbb{P}_P, \{\mathcal{F}_{t,P}\}_{t \geq 0})$.

Lemma 2 allows us to significantly simplify the principal's constrained optimization by only imposing obedience constraint for each type, similar to those in Equation (3).

To characterize non-personalized optimal communication, we leverage the solution of the personalized case. Lemma 2 implies that once one type quits, the solution for the remaining type reverts to the personalized solution. Recall that the relative likelihood ratios, $\ell^j = \mu_A^j(1 - \mu_P) / \mu_P(1 - \mu_A^j)$, remain constant. Hence, the principal's value in the personalized game for type j agent is $\alpha^j V(\mu_P(t), \ell^j)$, as in Equation (P).




Figure 8: The two stage structure of optimal communication

This formulation allows us to focus on the principal’s strategy when both agent types are engaged, as the principal’s value upon type $-j$ quitting is $\alpha_j V(\mu_P(t), \ell^j)$. Moreover, analogous to the personalized case, the value for the remaining agent of type j will be $v(\mu_A^j(t))$, while the value for the quitting agent of type $-j$ will be $v(\mu_A^{-j}(t))$. Thus, the payoff of agent of type j is always $v(\mu_A^j(t))$ upon realization of a transition signal.

With two types, a recommendation mechanism needs at most 4 recommendations per period, each associated with a pair of actions from the two types of agents. Moreover, if both types of agents are recommended to quit in a period, then an equivalent of Lemma 1 holds: it is still optimal to fully reveal the state to keep them engaged longer. However, if only one type is quitting, full revelation might be sub-optimal as it could cause the other type to quit as well. Let us define $\{p_\sigma(t), \mu_{\sigma,P}(t)\}_{\sigma \in \Sigma}$ to be the distribution of posteriors induced by the realization of $\sigma \in \Sigma$ for the principal at the beginning of phase 2, i.e., conditional on transition.

The above discussion implies that the principal’s strategy can be thought of as having two phases:

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. Each transitional signal realization σ is associated with $j(\sigma) \in \{L, H\}$, the type recommended to stay engaged. When recommending both agents to quit, the game ends at the end of phase 1, while with recommending only one type to stay we transition to phase 2. With one type engaged, we revert to the personalized case.

Figure 8 depicts the two phases of communication.

Stationary Engagement

Having established the above general structure and those of the relevant incentive constraints, we can analyze optimal communication similar to that in section 5. Before providing this characterization, we discuss one more property that will prove useful:




Figure 9: Value of the principal in phase 2

Lemma 3. *In any period at the beginning of phase 2 and for any signal realization σ , $j(\sigma) = H$ if and only if $\alpha^H V(\mu^\sigma; \ell^H) \geq \alpha^L V(\mu^\sigma; \ell^L)$ and vice versa. Moreover, it is sufficient to restrict the number of transition signals, $|\Sigma|$, to at most 4.*

The first part of Lemma 3 is simply a result of the fact that when the agent of type $j \in \{L, H\}$ exits, her payoff is always her outside option of $v(\mu_A^j(t))$. This implies that if a type is chosen to stay engaged, it should be the type that delivers the higher value to the principal. Figure 9 depicts two possible cases that can occur. In the first case (left panel), type L is never chosen since their fraction α^L is small enough while in the second case, L is chosen for high beliefs and H is chosen for low values. We can thus use $\hat{V}(\mu)$ instead of $\alpha^{j(\mu)} V(\mu; \ell^{j(\mu)})$ in the principal's payoff, where

$$\hat{V}(\mu) = \max \{ \alpha^H V(\mu; \ell^H), \alpha^L V(\mu; \ell^L) \}.$$

Moreover, an application of Fenchel and Bunt's theorem (see [Kamenica and Gentzkow \(2010\)](#) or Theorem 1.3.7. in [Hiriart-Urruty and Lemaréchal \(2004\)](#)) implies that we only need 4 signals.⁹

To draw a parallel with the personalized model of section 5, we can start by characterizing the point of long-run engagement or the level of beliefs for the principal μ_P^* , for which the strategy of the principal is stationary. Note that a stationary strategy for the principal involves an arrival rate of the transition signal λ^* from phase 1, and a distribution of posteriors, $\{p_\sigma^*, \mu_\sigma^*\}_{\sigma \in \Sigma}$. Let $\mu_{P,E}$ be the beliefs at the beginning of phase 2 – following the realization of transition signal –

⁹This can be seen by considering the convex hull of the set $S = \left\{ \left(\mu_P, \hat{V}(\mu_P), v(\mu_A^H), v(\mu_A^L) \right) : \mu_P \in [0, 1] \right\} \subset \mathbb{R}^4$ where $\mu_A^j / (1 - \mu_A^j) = \ell^j \mu_P / (1 - \mu_P)$. This set is connected since any two convex combinations of these points can be connected with a path which connects the weights in the appropriate simplex that contains them. Fenchel and Bunt's theorem then implies that any points in the convex hull can be written as a convex combination of 4 points in S .

but before the realization of σ which determines which type stays. Bayes plausibility implies that

$$\mu_{P,E} = \sum_{\sigma \in \Sigma} p_{\sigma}^* \mu_{\sigma}^*$$

Moreover, since μ_P is to stay constant over time, we must have that $\mu_{P,E}^* = \mu_P^*$. Finally, an argument akin to duality shown in Theorem 2 in the Appendix implies that Lagrange multipliers $\Lambda_j \geq 0$ should exist so that $\{p_{\sigma}^*, \mu_{\sigma}^*\}_{\sigma \in \Sigma}$ should maximize

$$\sum_{\sigma \in \Sigma} p_{\sigma} \left[\hat{V}(\mu_{\sigma}) + \sum_{j=L,H} \Lambda_j v_j(\mu_{\sigma}) \right] \quad (9)$$

where $v_j(\mu) = (\ell^j \mu_{\sigma} + 1 - \mu_{\sigma}) v\left(\frac{\ell^j \mu_{\sigma}}{\ell^j \mu_{\sigma} + 1 - \mu_{\sigma}}\right)$ is a change in payoff function that allows us to write everything from the perspective of the principal. Since $\mu_{P,E}$ satisfies the Bayes plausibility above, the distribution of posteriors upon transition to phase 2 is given by a concavification a la [Kamenica and Gentzkow \(2011\)](#) of the function $\hat{V} + \sum_{j=L,H} \Lambda_j v_j$.¹⁰ Let us refer to the concavification of the above as $\hat{V}^{\text{cav}}(\mu_{P,E}; \{\Lambda_j\})$. Finally, optimality of choice of transition probabilities from phase 1 and optimality $\mu_{P,E}$ implies that

$$\delta_P \frac{\partial \hat{V}^{\text{cav}}}{\partial \mu_{P,E}}(\mu_P^*; \Lambda_L, \Lambda_H) = (\delta_A - \delta_P) \sum_{j=L,H} \Lambda_j v_j'(\mu_P^*), \quad (10)$$

$$1 - \delta_P \hat{V}^{\text{cav}}(\mu_P^*; \Lambda_L^*, \Lambda_H^*) = (\delta_A - \delta_P) \sum_{j=L,H} \Lambda_j v_j(\mu_P^*) \quad (11)$$

The above are cost and benefits of changes in belief at the start of phase 2, $\mu_{P,E}$, and probability of exit from phase 1. In the equation 10, the right-hand-side (RHS) is the marginal cost of raising $\mu_P(t)$ and the left-hand-side (LHS) is marginal benefit of increase $\mu_{P,E}$ in the future. Note that an increase in future $\mu_{P,E}$ should be accompanied by an increase in current μ_P since beliefs have to be a martingale. One can think about this as a form of belief smoothing. Similarly, in 11, the left hand side is the benefit of lengthening phase 1, while the right hand side is the incentive cost of doing so. We summarize this in the following Proposition:

Proposition 6. *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 $\lambda^*, \Lambda_L^*, \Lambda_H^* \geq 0$ exists that satisfy:*

- (1) *The Belief Smoothing equation 10 holds,*
- (2) *The phase 1 optimality 11 hold,*

¹⁰Our problem is essentially a constrained Bayesian Persuasion problem – constrained by the agent’s incentive constraints. For a related paper, see [Doval and Skreta \(2024\)](#) for a general class of constrained Bayesian Persuasion problems.

(3) The following incentive compatibility and complementary slackness conditions are satisfied:

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

The above Proposition allows us to characterize the steady state of the non-personalized model. In the following example, we compare the model with personalized and non-personalized news in order to shed light on the key trade-offs.

Example 1. Suppose that $v(\mu) = 1 - 2\mu(1 - \mu)$, $\delta_P = 0$ and $\delta_A = 0.2$. In order to compare with the personalized model of section 5, we fix ℓ^L at the value of 0.3 and vary $\ell^H \geq 0.3$ and find the steady state. We assume that the fraction of high types is 0.6.

Figure 10 illustrates the idea of speed versus quality. In Figure 10a, for given likelihood ratios $\ell^L = 0.3$ and $\ell^H \geq 0.3$, we depict the stationary level of principal's belief μ_P^* (the green-crossed line) together with the beliefs that are induced in phase 2 (red and blue lines), μ_{σ}^* . For comparison, we have also included the stationary belief of the principal in the personalized case for $\ell^H \geq 0.3$. Not surprisingly, the non-personalized stationary belief is higher – since the existence of low type agent has to be taken into account. More importantly, the beliefs induced upon transition to phase 2 are different from certainty and sometimes fairly close to highest level of uncertainty. This implies that the type that receives the exit recommendation (High type for the red line and low type for the blue line) is exiting with a low quality of information, especially when ℓ^H is low. In other words, in order to keep the high type engaged, low type may exit at the end of phase 1 with fairly poor information.

In Figure 10b, we depict λ^* , the rate at which phase 1 ends in the non-personalized model together with the rate of arrival of exit in the personalized model. As this figure illustrates, the speed with which transition to phase 2 happens is significantly higher in the presence of private information about beliefs.

The above discussions highlight the key differences and similarities between personalized and non-personalized communication strategies. In the personalized scenario, the principal can exert perfect control over the agent's beliefs upon exit. In contrast, under non-personalized communication, while phase 1 of the interaction bears similarities to the personalized model, it differs in that the value of quitting from this phase is no longer zero for the principal. Furthermore, to prevent one type from quitting, it may sometimes be optimal to allow some types to leave without perfect information. Despite these distinctions, the fundamental principle governing the non-personalized optimal mechanism remains similar to that of the personalized one.

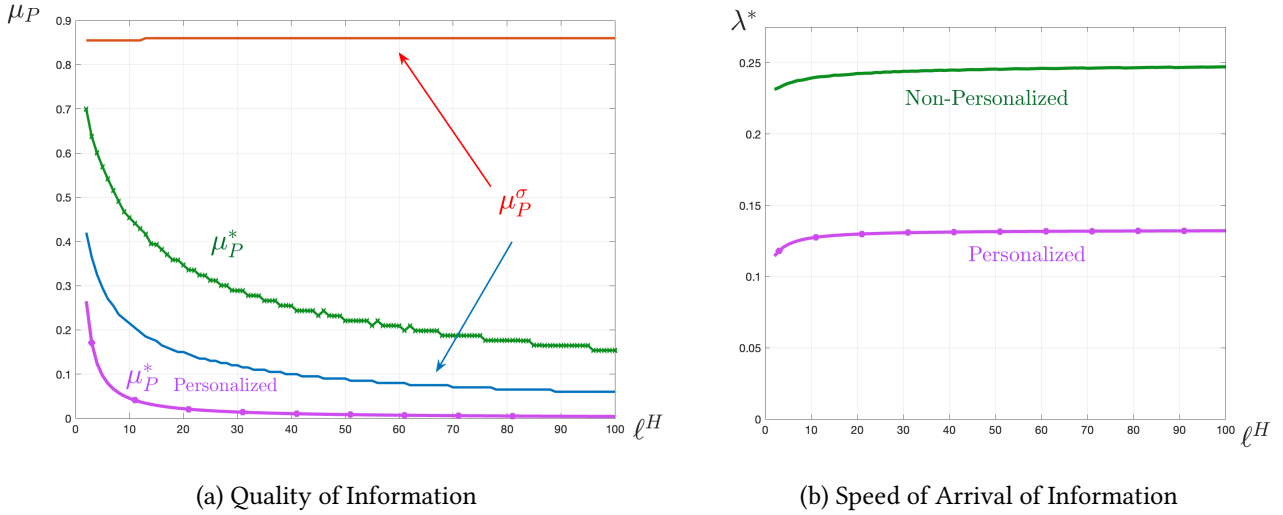


Figure 10: Speed vs. Quality

7 Extensions and Applications

In this section, we explore several extensions to our base model. We begin by incorporating random exit into our framework. We then investigate the implications of personalization on belief polarization, using our model to shed light on the ongoing debate about algorithmic news feeds and their impact on political discourse. Finally, we consider more general forms of time preferences, demonstrating how our model can accommodate various discounting structures beyond simple exponential discounting.

7.1 Random Exit

Throughout our analysis, our assumption has been that the agent can perfectly control her acquisition of information. In reality, however, attention is not perfect, and it is possible that audiences of news randomly exit. This is also in line with the extensive literature on limited attention or *rational inattention* that models attention as a noisy process where reducing the noise is costly; see [Sims \(2003\)](#) and the extensive literature following it.¹¹

To allow for random exit, suppose that if the agent is engaged at t , with probability $e^{-\rho(\tau-t)}$ she is still engaged at $\tau > t$. That is, she will exogenously exit in an interval $[\tau, \tau + d\tau]$ with probability $e^{-\rho(\tau-t)} - e^{-\rho(\tau+d\tau-t)}$. Under this modification and absent any discounting by the

¹¹[Morris and Strack \(2019\)](#) provide a foundation for the static costly noise reduction model of [Sims \(2003\)](#) with a dynamic decision maker. Our agent essentially behaves like theirs and solves an optimal stopping problem.

principal, his payoff is given by

$$-\int_0^\infty t d(e^{-\rho t} G_P(t)) = \int_0^\infty e^{-\rho t} G_P(t) dt$$

where G_P is the decumulative probability that the principal recommends staying. In this case, unlike the baseline model, exit before t can also occur exogenously and with probability $1 - e^{-\rho t}$. Hence, the probability has been adjusted accordingly.

The payoff of the agent is somewhat more complicated, as we need to take into account the occasions in which the agent exits exogenously and her payoff upon such exit. If the payoff of the agent upon recommended exit is $v_A(t)$, then her payoff at time t conditional on being engaged is given by

$$-\int_t^\infty e^{-(\delta_A + \rho)(\tau - t)} v_A(\tau) d\frac{G_A(\tau)}{G_A(t)} + \rho \int_t^\infty e^{-(\delta_A + \rho)(\tau - t)} v(\mu_A(\tau)) \frac{G_A(\tau)}{G_A(t)} d\tau.$$

When recommendation is personalized, we have $v_A(\tau) = v(1)$ which then implies that the above can be written as:

$$v(1) - \int_t^\infty e^{-(\delta_A + \rho)(\tau - t)} [(\rho + \delta_A) v(1) - \rho v(\mu_A(\tau))] \frac{G_A(\tau)}{G_A(t)} d\tau.$$

Therefore, the incentive constraint of the agent ensures that the agent's payoff above is higher than $v(\mu_A(t))$.

The following Proposition establishes that the case with exit is similar (but not identical) to that of the unequal discounting case with $\delta_P > \delta_A$:

Proposition 7. *Optimal communication with exit satisfies the following properties:*

1. *There exists $\mu_P^* \in [0, 1]$ such that $G_{P,\omega}(t) = G_{P,\omega}(0) e^{-\lambda t}$ and $\mu_A(t) = \mu_A$ for all t where λ and μ_P^* satisfy*

$$\mu_P^* = \max \left\{ \min \left\{ \mu_A + \frac{(\delta_A - \rho) \mu_A (1 - \mu_A) v'(\mu_A)}{(\delta_A - \rho) v(\mu_A) + \rho v(1)}, 1 \right\}, 0 \right\}$$

$$\lambda = \frac{\delta_A v(\mu_A)}{v(1) - v(\mu_A)}$$

2. *If $\delta_A > \rho$, then if $\mu_P > \mu_P^*$ ($\mu_P < \mu_P^*$), $\omega = 1$ ($\omega = 0$) is initially revealed to the agent until $t = t^*$ where $\mu_P^*(\mu_A(t^*)) = \mu_A(t^*)$.*
3. *If $\delta_A < \rho$, then if $\mu_P > \mu_P^*$ ($\mu_P < \mu_P^*$), $\omega = 0$ ($\omega = 1$) is initially revealed to the agent until $t = t^*$ where $\mu_P^*(\mu_A(t^*)) = \mu_A(t^*)$.*

4. When $t \geq t^*$, $\mu_A(t) = \mu_A(t^*)$ and both states are revealed at a rate $\delta_A v(\mu_A(t^*)) / (v(1) - v(\mu_A(t^*)))$.

7.2 Personalization and Belief Polarization

One of the much debated issues related to political polarization is the conflict of interest between algorithmic news feeds and news consumers. The concern is that social media platforms, whose main revenue source is targeted advertising, aim to maximize engagement rather than provide the most accurate information, potentially leading to political polarization.¹²

Using an example, we demonstrate how our model can be used as a framework to study the effect of targeting news. We consider the model with random exit discussed in the previous section and compare the optimal personalized communication outcomes with those of non-personalized communication outcomes. This comparison allows us to examine the potential impacts of news targeting strategies on information dissemination.

More specifically, we assume the following parametric assumptions. The agent's payoff as a function of her belief is $v(\mu) = 1 - 2\mu(1 - \mu)$, her discount rate is $\delta_A = 0.25$, and her exit rate is $\rho = 0.05$. Moreover, we consider two types of agents with different prior beliefs: high type and low type. We assume that their prior beliefs are $\mu_A^H = 0.647$, $\mu_A^L = 0.196$, respectively, and the fraction of high types is $\alpha^H = 0.6$. The prior belief of the principal is assumed to be $\mu_P = 0.55$ and its discount rate to be zero.

For this specification, we compare the evolution of beliefs under personalized and non-personalized communication. In this example, the principal's prior belief falls between those of the high and low agent types, specifically being less than the high type's prior and greater than the low type's prior. In the personalized case, the principal's strategy differs for each type. For the high type, the principal initially reveals information about $\omega = 1$, as this is the state toward which the high type is more biased. Conversely, for the low type, the principal initially reveals information about $\omega = 0$, reflecting the low type's bias. If no signals arrive, this strategy gradually shifts the agents' beliefs toward the opposite state. The steady state points for these two types are 0.71 for the low type and 0.43 for the high type. Once the agents reach their respective steady state points, the principal reveals both states at the same rate. This process continues until either a signal arrives or the agent exits due to the exogenous exit probability.

The optimal non-personalized communication can be calculated using the arguments in Section 6. As mentioned before, there are two phases in this process. In phase 1, the principal attempts to keep both agents engaged and during this phase the beliefs merge to a steady state level. This phase continues until a transition signal arrives which results in at least one of the agents leaving. We first numerically solve the constrained optimization problem that character-

¹²The term *filter bubble* was coined by Pariser (2011) to describe content controlled by algorithms that can create bias.




Figure 11: Evolution of beliefs and engagement probabilities

Notes: (a) The variable $\mu_P(t)$ is principal's belief if phase 1 continues. Variable $\mu_{P,l}$ ($\mu_{P,h}$) is principal's belief if the low-type (high-type) agent stays for phase 2. (b) The variable p_l (p_h) is the probability that the low-type (high-type) agent stays for phase 2 (personalized phase).

izes the steady state in phase 1. Next, we guess with an initial value for first order derivative of transition signals at time zero. Using the optimality conditions and this initial guess, we can solve the rest of parameters in the model. Finally, we checked that the resulting solution satisfies the time 0 incentive constraints, if not we modify the guess and reiterate.

During phase 1, the principal's belief changes until it converges to a steady state level, as depicted in Figure 11a. This change occurs because the transitional signal may arrive with different probabilities given the state. The principal's belief increases from 0.55 to a long-run rate of 0.63 in phase 1. The beliefs of the two types of agents also shift, as the likelihood ratio of the beliefs remains constant.

The transition signal can take different forms, either instructing only one type to stay or delivering full information and allowing both types of agents to leave. Figure 11b illustrates the probability of these different outcomes. Given the parameters in this example, it is never optimal to reveal the state and let both types leave. The probabilities for which type of agent stays and which leaves vary by the arrival time of the signal, reflecting changes in beliefs. These probabilities are calculated using the concavification method mentioned in Section 6. They are associated with the arrival of one of two different types of signals that discontinuously change the principal's and agents' beliefs. Figure 11b also illustrates the principal's beliefs after the arrival of each of these signals. The signal associated with the exit of the high type and the stay of the low




Figure 12: Distribution of exit beliefs under optimal personalized and non-personalized communication

Notes: The red solid line is the empirical cumulative distribution function (CDF) of agents' beliefs under personalized communication strategies. The blue dashed line is the empirical CDF of agents' beliefs under non-personalized communication strategies. The left two panels are for the low-type agent and the right two panels are for the high-type agent. The upper two figures are under true state $\omega = 1$ and the lower two figures are under true state $\omega = 0$. The empirical CDF comes from simulations with 10,000 belief paths for each case.

type is indicated by the belief $\mu_{P,l}$. After this signal arrives, the belief of the high type becomes very close to one, while the belief of the low type falls below the depicted line.¹³ After receiving this signal, the high-type agent will quit while the low-type agent chooses to stay and gather more information and we arrive to phase 2. The opposite occurs when the signal for keeping the high type arrives.

Figure 12 depicts the CDF of beliefs for each state and agent type as time approaches infinity. Due to exogenous exit rates, some agents may leave without full information even in the personalized case, resulting in non-degenerate belief distributions. In the personalized case, prolonged interactions may result in agents leaving without perfect information. We observe that

¹³To keep the figures tidy, we only depict the beliefs of the principal, the likelihood ratio of the beliefs stay the same, so beliefs of high type is always above the belief of principal and belief of low type is always below it.




Figure 13: Distribution of exit time under optimal personalized and non-personalized communication

Notes: Assuming that the distribution of the true state matches the principal’s prior $\mu_p = 0.55$, the red solid line is the CDF of exit time of agents under personalized recommendation strategies, and the blue dashed line is the CDF of exit time of agents under non-personalized recommendation strategies.

if the state is 1, 70% of high-type agents exit with this information compared to 30% of low-type agents, while if the state is 0, 80% of low-type agents leave with this information compared to 50% of high-type agents. This asymmetry arises from the principal catering to agents’ biases. If the true state aligns with an agent’s bias, they are likely to receive this information quickly. Otherwise, they may exit before receiving a signal revealing the opposite state and stay uninformed and biased toward the other state.

In the non-personalized case, alongside exogenous exits, the principal may choose to cater to only one agent type, prompting the other type to leave without full information. Comparing the distribution of agents across the four scenarios, we observe that in the non-personalized case, agents’ beliefs tend to be more dispersed and generally less aligned with the true state. For instance, when the state is 1, approximately 60% of low-type agents still strongly believe that the low state is likely ($\mu < 0.2$), and only about 25% of these agents learn that the state is actually 1. Similarly, fewer high-type agents discover that the state is 1 compared to the personalized case. These observations indicate that personalization leads to higher quality information dissemination, as agents are more likely to learn the true state under personalized communication strategies.

Another significant difference between personalized and non-personalized cases lies in the speed of information provision to agents. Figure 13 illustrates the CDF of exit times for agents in both personalized and non-personalized scenarios. We consistently observe earlier and faster exits in the non-personalized case, driven by two key forces. First, in the non-personalized case, during phase 1, the principal must keep both agent types engaged simultaneously, unable to cater

to them specifically. In contrast, the personalized case allows the principal to guide each agent to a point where the cost of maintaining engagement is lowest, thus prolonging interaction. This strategy is generally not feasible when dealing with two distinct agent types simultaneously. Second, in the personalized case, the arrival of a signal leads to state revelation, increasing the reward for staying engaged. In the non-personalized case the state is not going to be perfectly revealed as it prompts both types to quit. Consequently, given the lower reward in the non-personalized scenario, the principal must increase the frequency of rewards and, therefore, the speed of information arrival. This strategy ultimately results in quicker agent exits in the non-personalized case.

These findings highlight the opposing forces present in our model, illustrating a trade-off between information quality and speed of delivery. Personalization results in slower information transmission but higher quality, while non-personalized communication leads to faster transmission but lower quality. In both scenarios, we observe that agents starting with lower priors tend to maintain their belief that their preferred state is more likely more often than those agents who initially believe the opposite state is more likely. This persistence of beliefs suggests that polarization can occur in both personalized and non-personalized communication settings. Importantly, our results indicate that personalized communication does not necessarily lead to more polarization than non-personalized communication. This insight may help explain the mixed results found in studies of the “filter bubble” phenomenon.¹⁴ However, it is crucial to note that a more comprehensive quantitative investigation is needed to fully understand these findings.

7.3 Other forms of Discounting

It is possible to extend our analysis to more general cases of time preferences. Specifically, suppose that principal’s preferences are given by T and that of the agent is $D(T) \hat{u}(\omega, a)$. Below, we show how this encompasses several popular models of time preferences.

7.3.1 The Nature of Time Cost for the Agent

Here, we show how different commonly used setups for inter-temporal discounting map into the setting described above:

1. **Exponential Discounting.** The case of exponential discounting that we have so far studied is a special case of this extension. In this case the principal payoff is $\delta_P \int_0^T e^{-\delta_p t} dt$ and

¹⁴Recent research offers conflicting evidence on this topic. For instance, [Jones-Jang and Chung \(2024\)](#) use survey data to show that social media, in contrast to traditional media, leads to less polarization. Conversely, [Bakshy et al. \(2015\)](#) conduct a randomized experiment and find that algorithmically ranked news feeds result in 15% lower exposure to cross-cutting content. These contrasting findings underscore the complexity of the issue and the need for further research in this area.

the agent payoff is $e^{-\delta_a \hat{T}} \hat{u}(\omega, a)$, where \hat{T} is the engagement time by the agent. We can let T be the payoff of the principal and in this case the payoff of the agent is given by

$$u(T, \omega, a) = (1 - T)^{\frac{\delta_a}{\delta_P}} \hat{u}(\omega, a)$$

whose concavity and convexity is determined by how δ_A/δ_P compares to 1.

2. **Habit Formation a la Becker and Murphy (1988):** Another example of inter-temporal preferences is the habit formation model of [Becker and Murphy \(1988\)](#) (also used in [Allcott et al. \(2022\)](#)). In this model, preferences are a function of habit – a stock variable that increases exponentially with consumption. In the context of digital addiction setting of [Allcott et al. \(2022\)](#), consumption can be interpreted as social media engagement. Let H_t be the stock of habit which satisfies

$$\frac{d}{dt} H_t = \lambda H_t - h \rightarrow H_t = e^{\lambda t} H_0 + h/\lambda$$

and suppose that the agent's utility is given by

$$H_{\hat{T}} e^{-\delta_A \hat{T}} = e^{(\lambda - \delta_A) \hat{T}} H_0 + e^{-\delta_A \hat{T}} h/\lambda$$

When $\lambda < \delta_A$, this is equivalent to a mixed-discounting model and can similarly be expressed in terms of the payoff of the principal $T = 1 - e^{-\delta_P \hat{T}}$. As depicted in [Figure 14a](#), when habits are weak and the agent is less patient than the principal ($\delta_P < \delta_A - \lambda$), MCE is decreasing over time while when the agent is less patient than the principal, MCE is increasing over time. In the intermediate case when habits are strong enough to make $\delta_A - \lambda < \delta_P$, MCE initially decreases over time and as time passes habit formation becomes dominant and MCE starts increasing over time.

3. **A Reduced Form Model of Boredom:** Another way to think about time preferences for the agent is through the lens of a reduced form version of the dual-self control model of [Fudenberg and Levine \(2006\)](#). We can think of the agent as consisted of two selves: a long-run planner and a short-run decider. The decider's costly choice depends on the total time spent being engaged which can be a stock variable associated with boredom. The long-run planner in turn has to make contingent planning taking into account the cost of boredom for the short-run player. For example, the agent could have a utility of the form

$$D(T) = e^{-\delta_A T - \gamma T^2}$$

where γT^2 captures the cost of boredom. In this case, MCE initially increases while after a certain period it decreases. This is depicted in Figure 14b.¹⁵




Figure 14: Examples of Relative Marginal Cost of Engagement

7.3.2 Optimal Communication with General Discounting

The case of concave and convex $D(T)$ are qualitatively the same as their associated cases with exponential discounting. Here, we describe what happens when $D(T)$ is either a concave-convex or a convex-concave function. In these cases, perhaps not surprisingly, the solution is often determined by superimposing the Poisson gradual revelation and abrupt revelations. In what follows, we describe the nature of this superimposition and key properties of optimal revelation.

Proposition 8. Convex-Concave Discounting. *Suppose that there exists a threshold t^i such that above $D'(t)$ is increasing for values of $t \leq t^i$ and decreasing for values of $t > t^i$. Then optimal information provision has three phases:*

1. A phase for $0 \leq t < t_1^* < t^i$ where only one state is revealed (towards which the agent is more optimistic) gradually, i.e. according to a time varying Poisson process. When $\mu_0 > 1/2$, the evolution of beliefs is given by the ODE:

$$\frac{\mu'(t)}{\mu(t)(1-\mu(t))} \frac{v(1) - v(\mu(t)) - v'(\mu(t))(1-\mu(t))}{v(\mu(t))} = \frac{D'(t)}{D(t)}.$$

2. A second phase for $t \in [t_1^*, t_2^*]$ with $t_2^* < t^i$ where both states are gradually revealed according to a Poisson process with arrival rate $-D'(t)/D(t)$.

¹⁵See also Wojtowicz et al. (2023) for an extended version of this model with information acquisition by the dual selves.

3. A final phase for $t \in [t_2^*, t_3^*]$ with $t^i < t_3^*$ in which no information is provided and all information is revealed at t_3^* .

As Proposition 8 illustrates, the convex-concave case is a superimposition of the two cases discussed before. Considering the habit formation example of section 7.3.1, when $0 < \delta_A - \lambda < \delta_P < \delta_A$, as the agent is initially impatient relative to the principal, information must be revealed gradually and in the second phase at rate

$$-\frac{D'(t)}{D(t)} = \delta_A + (\lambda - \delta_A) \frac{\alpha}{\alpha + \beta e^{-\lambda t}}$$

for some $\alpha, \beta > 0$. This gradual revelation is slowed down as time goes by and the agent accumulates habit. Once the force of habit formation becomes strong enough, the value of gradual revelation disappears and instead abrupt revelation is used.

8 Conclusion

This paper examines the strategic interaction between an expert (principal) seeking to maximize engagement and an agent aiming to collect information swiftly. Our analysis reveals several key insights into optimal information provision strategies in such environments.

First, we find that the relative patience of the agent and principal is a crucial determinant of the optimal information revelation strategy. In the case of common priors, a more patient agent leads to gradual information disclosure, while an impatient agent results in more abrupt revelation. This underscores the importance of considering time preferences in designing engagement strategies.

Second, our model demonstrates that when the principal and agent have different prior beliefs, optimal disclosure often involves “catering to the bias.” The principal tends to initially reveal information aligned with the agent’s bias, exploiting this asymmetry to prolong engagement. This finding has significant implications for understanding how information providers may strategically utilize belief disagreements.

Third, in scenarios with private agent beliefs, we identify two phases in communication strategies. During phase 1, the principal attempts to keep both agent types engaged. A transition signal during this phase determines which type of agent should stay and which should leave. Notably, the principal may allow one agent to depart without full information, contrasting with the personalized case, to keep the other agent engaged. Once only one agent type remains, the strategy reverts to the personalized case.

Our comparison of personalized and non-personalized information revelation strategies highlights important trade-offs. Personalization enables longer engagement and slower information

delivery, while non-personalized approaches lead to faster information transmission but may result in lower quality information overall.

References

- ACEMOGLU, D. AND S. JOHNSON (2024): “The Urgent Need to Tax Digital Advertising,” *Network Law Review*, Spring. 5
- ALLCOTT, H., L. BRAGHIERI, S. EICHMEYER, AND M. GENTZKOW (2020): “The Welfare Effects of Social Media,” *American Economic Review*, 110, 629–76. 10
- ALLCOTT, H., M. GENTZKOW, AND L. SONG (2022): “Digital addiction,” *American Economic Review*, 112, 2424–2463. 39
- AUMANN, R. J. (1976): “Agreeing to Disagree,” *The Annals of Statistics*, 4, 1236–1239. 2
- BAKSHY, E., S. MESSING, AND L. A. ADAMIC (2015): “Exposure to ideologically diverse news and opinion on Facebook,” *Science*, 348, 1130–1132. 38
- BALL, I. AND J. KNOEPFLE (2023): “Should the timing of inspections be predictable?” *Available at SSRN 4419716*. 6
- BECKER, G. S. AND K. M. MURPHY (1988): “A Theory of Rational Addiction,” *Journal of Political Economy*, 96, 675–700. 39
- CHE, Y.-K., K. KIM, AND K. MIERENDORFF (2023): “Keeping the Listener Engaged: a Dynamic Model of Bayesian Persuasion,” *Journal of Political Economy*, 131, null. 5
- DEJARNETTE, P., D. DILLENBERGER, D. GOTTLIEB, AND P. ORTOLEVA (2020): “Time lotteries and stochastic impatience,” *Econometrica*, 88, 619–656. 6
- DOVAL, L. AND V. SKRETA (2024): “Constrained information design,” *Mathematics of Operations Research*, 49, 78–106. 30
- DUNFORD, N. AND J. T. SCHWARTZ (1988): *Linear operators, part 1: general theory*, vol. 10, John Wiley & Sons. 47
- DWORCZAK, P. AND G. MARTINI (2019): “The simple economics of optimal persuasion,” *Journal of Political Economy*, 127, 1993–2048. 15
- ELY, J. C. (2017): “Beeps,” *American Economic Review*, 107, 31–53. 5

- ELY, J. C. AND M. SZYDLOWSKI (2020): “Moving the Goalposts,” *Journal of Political Economy*, 128, 468–506. 5
- FUDENBERG, D. AND D. K. LEVINE (2006): “A Dual-Self Model of Impulse Control,” *American Economic Review*, 96, 1449–1476. 39
- GUO, Y. (2016): “Dynamic delegation of experimentation,” *American Economic Review*, 106, 1969–2008. 6
- HADFIELD, G. K. (2007): “Legal barriers to innovation: The growing economic cost of professional control over corporate legal markets,” *Stan. L. Rev.*, 60, 1689. 9
- (2022): “Legal markets,” *Journal of Economic Literature*, 60, 1264–1315. 9
- HALAC, M., N. KARTIK, AND Q. LIU (2016): “Optimal contracts for experimentation,” *The Review of Economic Studies*, 83, 1040–1091. 6
- HÉBERT, B. AND W. ZHONG (2022): “Engagement Maximization,” . 6
- HENRY, E. AND M. OTTAVIANI (2019): “Research and the approval process: The organization of persuasion,” *American Economic Review*, 109, 911–955. 6
- HIRIART-URRUTY, J.-B. AND C. LEMARÉCHAL (2004): *Fundamentals of convex analysis*, New York: Springer. 29
- JONES-JANG, S. M. AND M. CHUNG (2024): “Can we blame social media for polarization? Counter-evidence against filter bubble claims during the COVID-19 pandemic,” *New Media & Society*, 26, 3370–3389. 38
- KAMENICA, E. AND M. GENTZKOW (2010): “Bayesian persuasion: Web appendix,” *American Economic Review*. 29
- (2011): “Bayesian Persuasion,” *American Economic Review*, 101, 2590–2615. 5, 11, 30
- KELLER, G. AND S. RADY (2010): “Strategic Experimentation with Poisson Bandits,” *Theoretical Economics*, 5, 275–311. 6, 12
- KELLER, G., S. RADY, AND M. CRIPPS (2005): “Strategic experimentation with exponential bandits,” *Econometrica*, 73, 39–68. 6
- KLEINER, A., B. MOLDOVANU, AND P. STRACK (2021): “Extreme Points and Majorization: Economic Applications,” *Econometrica*, 89, 1557–1593. 15

- KNOEPFLE, J. (2020): “Dynamic Competition for Attention,” Tech. rep., University of Bonn and University of Mannheim, Germany. 6
- KOH, A. AND S. SANGUANMOO (2022): “Attention capture,” *arXiv preprint arXiv:2209.05570*. 6
- KOH, A., S. SANGUANMOO, AND W. ZHONG (2024): “Persuasion and Optimal Stopping,” *arXiv preprint arXiv:2406.12278*. 6
- LEVY, R. (2021): “Social Media, News Consumption, and Polarization: Evidence from a Field Experiment,” *American Economic Review*, 111, 831–70. 10
- LUENBERGER, D. G. (1997): *Optimization by Vector Space Methods*, John Wiley & Sons. 14, 15, 45, 48
- MORRIS, S. AND P. STRACK (2019): “The wald problem and the relation of sequential sampling and ex-ante information costs,” *Available at SSRN 2991567*. 32
- ORLOV, D., A. SKRZYPACZ, AND P. ZRYUMOV (2020): “Persuading the Principal to Wait,” *Journal of Political Economy*, 128, 2542–2578. 5
- PARISER, E. (2011): *The filter bubble: What the Internet is hiding from you*, penguin UK. 34
- RENAULT, J., E. SOLAN, AND N. VIEILLE (2017): “Optimal dynamic information provision,” *Games and Economic Behavior*, 104, 329–349. 5
- SAEEDI, M. AND A. SHOURIDEH (2023): “Optimal Rating Design,” *Working Paper*. 15
- SIMS, C. A. (2003): “Implications of rational inattention,” *Journal of monetary Economics*, 50, 665–690. 32
- SMOLIN, A. (2021): “Dynamic Evaluation Design,” *American Economic Journal: Microeconomics*, 13, 300–331. 5
- STRÖMBERG, D. (2004): “Mass media competition, political competition, and public policy,” *The Review of Economic Studies*, 71, 265–284. 10
- STRULOVICI, B. (2010): “Learning while voting: Determinants of collective experimentation,” *Econometrica*, 78, 933–971. 6
- WANG, X., J. SHI, AND H. KONG (2021): “Online Health Information Seeking: A Review and Meta-Analysis,” *Health Communication*, 36, 1163–1175, PMID: 32290679. 10
- WOJTOWICZ, Z., N. CHATER, AND G. LOEWENSTEIN (2023): “Boredom and Flow: An Opportunity Cost Theory of Attention-Directing Motivational States,” *Available at SSRN 3339123*. 40

A Proofs

A.1 Proof of Theorem 1

Proof. We prove a more general version of the theorem. As we discussed in Section 7.3.2, our model is equivalent to one in which time is given by $T \in [0, \bar{T}]$ with $\bar{T} < \infty$ and the payoff of the principal is given by T while the payoff of the agent is given by

$$D(T) v(\mu_A(T))$$

where $D(T)$ is a strictly decreasing and positive function with possibly satisfying $D(\bar{T}) = 0$. Note that when exponential discounting, $T = \frac{1-e^{-\delta_P t}}{\delta_P}$ and $D(T) = (1 - \delta_P T)^{\frac{\delta_A}{\delta_P}}$ with $\bar{T} = 1/\delta_P$.

For completeness and easier reference, let us restate Theorem 1 on page 217 of [Luenberger \(1997\)](#):

Theorem 2. *Let X be a linear vector space, Z a normed space, Ω a convex subset of X , and P a positive cone in Z . Suppose that P contains an interior point. Let f be a real-valued convex functional on Ω and H a convex mapping from Ω into Z . Assume the existence of $x_1 \in \Omega$ for which $H(x_1) < 0$, i.e., $H(x_1)$ is an interior point of $-P$. Let*

$$\mu_0 = \inf_{x \in \Omega} f(x) \text{ subject to } H(x) \leq 0 \quad (12)$$

and assume μ_0 is finite. Then there is an element $z_0^* \geq 0$ in Z^* such that

$$\mu_0 = \inf_{x \in \Omega} f(x) + \langle z_0^*, H(x) \rangle. \quad (13)$$

Furthermore, if the infimum is achieved in (12) by an $x_0 \in \Omega$, $H(x_0) \leq 0$, it is achieved by x_0 in (13) and

$$\langle z_0^*, H(x_0) \rangle = 0$$

We will verify that our optimization problem satisfies the conditions in Theorem 2.

A key assumption in Theorem 2 is the existence of x_1 ; existence of $x_1 - H(x_1) \in \text{int}(P)$. This proves difficult in general. Therefore, we consider two cases: 1. when $D(\bar{T}) > 0$ and 2. when $D(\bar{T}) = 0$. In the first case, we can use Theorem 2, and in the second case, we use a limiting argument.

Case 1. Suppose that $D(\bar{T}) < \infty$. We define the objects in Theorem 2 as follows:

$$\begin{aligned}
X &= \{x = (G_{P,1}(T), G_{P,0}(T)) \mid T \in [0, \bar{T}], G_{P,\omega} : \text{bounded}\} = L^\infty([0, \bar{T}])^2 \\
Z &= L^\infty([0, \bar{T}]) \times \mathbb{R}^2 \\
\Omega &= \{G_{P,\omega} : \text{non-increasing}, G_{P,\omega} \geq 0\} \\
P &= \{h(T) \geq 0, \forall T \in [0, \bar{T}]\} \times \mathbb{R}_+^2 \\
H(x) &= (H_1(x), G_{P,1}(0) - \mu_P, G_{P,0}(0) - 1 + \mu_P) \\
H_1(x)(T) &= D(T) (\ell G_{P,1}(T) + G_{P,0}(T)) \left[v \left(\frac{\ell G_{P,1}(T)}{\ell G_{P,1}(T) + G_{P,0}(T)} \right) - v(1) \right] \\
&\quad - v(1) \int_T^{\bar{T}} (\ell G_{P,1}(s) + G_{P,0}(s)) D'(s) ds \\
f(x) &= - \int_0^{\bar{T}} [G_{P,1}(T) + G_{P,0}(T)] dT
\end{aligned}$$

where in the above \mathbb{R}_+ is the set of real numbers that are non-negative. Trivially, X and Z satisfy the hypothesis in Theorem 2. The set Ω is a convex set, and P is a convex cone. Finally, since v is a convex function,

$$\begin{aligned}
\frac{\frac{1}{2}(x + \ell y) v \left(\frac{x}{x + \ell y} \right) + \frac{1}{2}(x' + \ell y') v \left(\frac{x'}{x' + \ell y'} \right)}{\frac{1}{2}(x + \ell y) + \frac{1}{2}(x' + \ell y')} &\geq v \left(\frac{\frac{1}{2}(x + \ell y) \frac{x}{x + \ell y} + \frac{1}{2}(x' + \ell y') \frac{x'}{x' + \ell y'}}{\frac{1}{2}(x + \ell y) + \frac{1}{2}(x' + \ell y')} \right) \\
&= v \left(\frac{\frac{1}{2}x + \frac{1}{2}x'}{\frac{1}{2}(x + \ell y) + \frac{1}{2}(x' + \ell y')} \right)
\end{aligned}$$

which then implies that

$$\begin{aligned}
\frac{1}{2}(x + \ell y) v \left(\frac{x}{x + \ell y} \right) + \frac{1}{2}(x' + \ell y') v \left(\frac{x'}{x' + \ell y'} \right) &\geq \\
\left(\frac{1}{2}(x + \ell y) + \frac{1}{2}(x' + \ell y') \right) v \left(\frac{\frac{1}{2}x + \frac{1}{2}x'}{\frac{1}{2}(x + \ell y) + \frac{1}{2}(x' + \ell y')} \right). &
\end{aligned}$$

In other words, the function $(x + \ell y) v \left(\frac{x}{x + \ell y} \right)$ is convex. As a result, the function $H(x)$ is convex.

Let $G_{P,1}(T) = (1 - \varepsilon) \mu_P D(T)^\lambda$, $G_{P,0}(T) = (1 - \varepsilon) (1 - \mu_P) D(T)^\lambda$ where

$$\lambda = \frac{v(\mu_A)}{v(1) - v(\mu_A)} > 0$$

and the strict concavity of v implies that if $\mu_A, \mu_P \in (0, 1)$, the above is positive. Let x_1 be the

associated member of $L^\infty([0, \bar{T}])$. We have that:

$$\begin{aligned}
G_{P,1}(0) - \mu_P &= -\varepsilon\mu_P, \\
G_{P,0}(0) + 1 - \mu_P &= -\varepsilon(1 - \mu_P) \\
H_1(x_1)(T) &= D(T)^{1+\lambda}(1 - \varepsilon)(\ell\mu_P + 1 - \mu_P)[v(\mu_A) - v(1)] \\
&\quad - v(1)(1 - \varepsilon)(\ell\mu_P + 1 - \mu_P) \int_T^{\bar{T}} D(s)^\lambda D'(s) ds \\
&= (1 - \varepsilon) \frac{1 - \mu_P}{1 - \mu_A} \left([v(\mu_A) - v(1)] D(T)^{1+\lambda} - v(1) \frac{D(\bar{T})^{1+\lambda} - D(T)^{1+\lambda}}{1 + \lambda} \right) \\
&= (1 - \varepsilon) \frac{1 - \mu_P}{1 - \mu_A} \left(-\frac{v(1)}{1 + \lambda} D(T)^\lambda - v(1) \frac{D(\bar{T})^{1+\lambda} - D(T)^{1+\lambda}}{1 + \lambda} \right) \\
&= - (1 - \varepsilon) \frac{1 - \mu_P}{1 - \mu_A} v(1) \frac{D(\bar{T})^{1+\lambda}}{1 + \lambda}
\end{aligned}$$

If we let $\bar{\varepsilon} = \min \left\{ \varepsilon\mu_P, \varepsilon(1 - \mu_P), (1 - \varepsilon)v(1) \frac{1 - \mu_P}{1 - \mu_A} \frac{D(\bar{T})^{1+\lambda}}{1 + \lambda} \right\} > 0$, then we have that

$$\max \left\{ \sup_{T \leq \bar{T}} H_1(x_1)(T), G_{P,1}(0) - \mu_P, G_{P,0}(0) - 1 + \mu_P \right\} \leq -\bar{\varepsilon}$$

and therefore, a ball of radius $\bar{\varepsilon}/2$ around x_1 is a subset of $-P$. Moreover, f is linear in x and its value is bounded below by 0 and above by $\bar{T} < \infty$. Finally, since the functions $G_{P,\omega}$ are bounded, the set of x 's such that $x \in \Omega, H(x) \leq 0$ is compact. This means that the optimization $\inf_{x \in \Omega, H(x) \leq 0} f(x)$ achieves its value, and this confirms all the hypotheses of Theorem 2.

By standard Representation Theorems (Riesz-Kakutani-Markov), see Theorem IV.16 in [Dunford and Schwartz \(1988\)](#) as an example, the dual of the space of bounded functions is isomorphic to the space of finitely additive measures of bounded variation or in turn the space of functions of bounded variations that vanish at $t = 0$. This implies that there must exist $z^* \in Z^* = (\Lambda(T), \lambda_1, \lambda_0)$ where $\lambda_1, \lambda_0 \geq 0$, $\Lambda(t)$ is an increasing function of bounded variation¹⁶ with $\Lambda(0) = 0$ such that $\hat{x} = (\hat{G}_{P,1}, \hat{G}_{P,0})$ is the solution to $\max_{x \in \Omega, H(x) \leq 0} -f(x)$ if

$$\begin{aligned}
\hat{x} &\in \arg \max_{x \in \Omega} -f(x) - \int_0^{\bar{T}} H_1(x)(T) d\Lambda(T) - \lambda_1(G_{P,1}(0) - \mu_P) - \lambda_0(G_{P,0}(0) - 1 + \mu_P) \\
&= \arg \max_{x \in \Omega} L(x)
\end{aligned}$$

¹⁶The fact that $\Lambda^T(t)$ must be (weakly) increasing is implied by $z^* \geq 0$ which means that $z^* \in P^*$, the dual of the cone P which is defined as members of Z^* , such that $\langle z^*, z \rangle \geq 0$ for all $z \in P$.

Note that Ω is the set of positive, bounded, and non-increasing functions. Thus, it is a convex cone in X . Hence, we can apply the following Lemma from [Luenberger \(1997\)](#) – Lemma 1, page 227, section 8.7:

Lemma 4. *Let f be a Fréchet differentiable convex functional on a real normed space X . Let P be a convex cone in X . A necessary and sufficient condition for $x_0 \in P$ to minimize f over P is that*

$$\begin{aligned}\partial f(x_0; x) &\geq 0, \forall x \in P \\ \partial f(x_0; x_0) &= 0.\end{aligned}$$

Applying the above to L , we have

$$\begin{aligned}\partial L(\hat{x}; x) &= \int_0^{\bar{T}} (G_{P,1}(T) + G_{P,0}(T)) dT - \tag{14} \\ &\int_0^{\bar{T}} [v(\mu_A(T)) - v(1)] (\ell G_{P,1}(T) + G_{P,0}(T)) D(T) d\Lambda(T) - \\ &\int_0^{\bar{T}} [v'(\mu_A(T))(1 - \mu_A(T)) \ell G_{P,1}(T) - v'(\mu_A(T)) \mu_A(T) G_{P,0}(T)] D(T) d\Lambda(T) + \\ &v(1) \int_0^{\bar{T}} \int_T^{\bar{T}} D'(s) (\ell G_{P,1}(s) + G_{P,0}(s)) ds d\Lambda(T) - \\ &\lambda_1 G_{P,1}(0) - \lambda_0 G_{P,0}(0) \leq 0\end{aligned}$$

$$\partial L(\hat{x}; \hat{x}) = 0 \tag{15}$$

where in the above $x = (G_{P,1}(t), G_{P,0}(t))$ is an arbitrary member of Ω and $\mu_A(T)$ is the belief of the agent given by $\ell G_{P,1}(T) / (\ell G_{P,1}(T) + G_{P,0}(T))$ calculated at the optimum. This proves the claim for case 1.

To show the result for case 2, i.e., $D(\bar{T}) = 0$, we use continuity arguments. Since $D(\hat{T}) > 0$, we can apply the above logic to $T \in [0, \hat{T}]$ with $\hat{T} < \bar{T}$. We then consider the solution $\hat{x}_{\hat{T}}$ and its associated multiplier $\Lambda_{\hat{T}}$ that satisfy the above condition. We extend $\Lambda_{\hat{T}}$ to the entire interval $[0, \bar{T}]$ by setting $\Lambda_{\hat{T}}(T) = \Lambda_{\hat{T}}(\hat{T}), \forall T \geq \hat{T}$. We use the following lemma to characterize the properties of a limit point of the solutions $\hat{x}_{\hat{T}}$ and multipliers $\Lambda_{\hat{T}}$.

Lemma 5. *There exists a constant $C > 0$ such that for all values of $\hat{T} < \bar{T}$, $D(\hat{T})^{\frac{v(1)}{v(1)-v(1/2)}} \Lambda_{\hat{T}}(\hat{T}) \leq C$.*

We relegate the proof to Section [A.7](#).

Since each function $\Lambda_{\hat{T}}(\cdot)$ is increasing, Lemma 5 together with the compactness of bounded

functions in the space of functions with bounded variations¹⁷ imply that there exists a bounded function with bounded variation $\bar{\Gamma}(T)$ and a subsequence of $\Lambda_{\hat{T}}$'s, given by Λ_k 's, so that

$$\lim_{k \rightarrow \infty} \sup_{T \in (0, \bar{T})} \left| D(T)^{\frac{v(1)}{v(1)-v(1/2)}} \Lambda_k(T) - \bar{\Gamma}(T) \right| = 0 \quad (16)$$

Moreover, since each Λ_k is increasing so is $\bar{\Gamma}(T) D(T)^{-\frac{v(1)}{v(1)-v(1/2)}}$.

Additionally, note that $\frac{v(1)}{v(1)-v(1/2)} > 1$, $\int_0^{\bar{T}} D(T)^{-\frac{v(1)}{v(1)-v(1/2)}} dT < \infty$. This implies that the measure $\bar{\Lambda}$ associated with the function $\bar{\Gamma}(T) D(T)^{-\frac{v(1)}{v(1)-v(1/2)}}$, i.e., $\bar{\Gamma}(T) D(T)^{-\frac{v(1)}{v(1)-v(1/2)}}$ is the CDF of $\bar{\Lambda}$, is a finite measure. This implies that $\int f d\Lambda_k \rightarrow \int f d\bar{\Lambda}$ for all functions with bounded derivatives.

Since $x_{\hat{T}}$'s are bounded, they must converge as $\hat{T} \rightarrow \bar{T}$. By Berge's maximum theorem, the limit of $x_{\hat{T}}$ should be an optimizer for \bar{T} . The result then follows from the continuity of the expressions in (14). \square

¹⁷By the Banach-Alaoglu theorem.

A.2 Proof of Proposition 1

Proof. We use Theorem 1 and construct multiplier so that abrupt revelation is optimal.

$$\Lambda(t) = \begin{cases} 0 & t = 0 \\ \frac{e^{(\delta_A - \delta_P)t^*}}{v(1)\delta_A} & t > 0 \end{cases}.$$

$$\lambda_1 = \frac{1 - e^{-\delta_P t^*}}{\delta_P} + \frac{e^{(\delta_A - \delta_P)t^*}}{\delta_A v(1)} (v(1) e^{-\delta_A t^*} - v(\mu_A) - v'(\mu_A)(1 - \mu_A))$$

$$\lambda_0 = \frac{1 - e^{-\delta_P t^*}}{\delta_P} + \frac{e^{(\delta_A - \delta_P)t^*}}{\delta_A v(1)} (v(1) e^{-\delta_A t^*} - v(\mu_A) + v'(\mu_A)\mu_A)$$

where

$$t^* = \frac{1}{\delta_A} \log \frac{v(1)}{v(\mu_A)}$$

Under abrupt revelation, the distribution of exit rates are

$$\hat{G}_{P,1}(t) = \mu_P \mathbf{1}[t < t^*]$$

$$\hat{G}_{P,0}(t) = (1 - \mu_P) \mathbf{1}[t < t^*]$$

As we show in the proof of Theorem 1, $\partial L(\hat{x}; \hat{x}) = 0$ is equivalent to

$$\int_0^{t^*} e^{-\delta_P t} dt - \lambda_1 \mu_P - \lambda_0 (1 - \mu_P) = 0$$

which is satisfied by construction.

Next, we need to show that for all non-increasing positive functions $x = (G_{P,0}(t), G_{P,1}(t))$, $\partial L(\hat{x}; x) \leq 0$ is satisfied. We have

$$\begin{aligned} \partial L(\hat{x}; x) &= \int_0^\infty (G_{P,1}(t) + G_{P,0}(t)) e^{-\delta_P t} dt - \\ &\quad [v(\mu_A) - v(1)] (G_{P,1}(0) + G_{P,0}(0)) \Lambda(0+) - \\ &\quad [v'(\mu_A)(1 - \mu_A) G_{P,1}(0) - v'(\mu_A)\mu_A(0) G_{P,0}(0)] \Lambda(0+) + \\ &\quad - v(1) \int_0^\infty \delta_A e^{-\delta_A s} (G_{P,1}(s) + G_{P,0}(s)) ds \Lambda(0+) - \\ &\quad \lambda_1 G_{P,1}(0) - \lambda_0 G_{P,0}(0) \end{aligned}$$

The above is linear in x . Thus, if we show that for any positive and decreasing function $f(t)$, $\partial L(\hat{x}; (f, 0)) \leq 0$ and $\partial L(\hat{x}; (0, f)) \leq 0$, this proves the desired result. We will show that

$\partial L(\hat{x}; (f, 0)) \leq 0$; the second inequality can be shown similarly. We have

$$\begin{aligned}
\partial L(\hat{x}; x) &= \int_0^\infty f(t) e^{-\delta_P t} dt - \\
&\quad f(0) [v(\mu_A) - v(1) + v'(\mu_A)(1 - \mu_A)] \Lambda(0+) - \\
&\quad - v(1) \int_0^\infty \delta_A e^{-\delta_A s} f(s) ds \Lambda(0+) - \lambda_1 f(0) \\
&= \int_0^\infty f(t) e^{-\delta_P t} dt - \\
&\quad f(0) [v(\mu_A) - v(1) + v'(\mu_A)(1 - \mu_A)] \Lambda(0+) - \\
&\quad - v(1) \int_0^\infty \delta_A e^{-\delta_A s} f(s) ds \Lambda(0+) \\
&\quad - f(0) \frac{1 - e^{-\delta_P t^*}}{\delta_P} - f(0) \Lambda(0+) (v(1) e^{-\delta_A t^*} - v(\mu_A) - v'(\mu_A)(1 - \mu_A)) \\
&= \int_0^\infty f(t) e^{-\delta_P t} dt - f(0) \frac{1 - e^{-\delta_P t^*}}{\delta_P} \\
&\quad + \Lambda(0+) \left[v(1) f(0) - v(1) e^{-\delta_A t^*} f(0) - v(1) \int_0^\infty \delta_A e^{-\delta_A s} f(s) ds \right] \\
&\leq \int_0^\infty (f(t) - f(0)) e^{-\delta_P t} dt + f(0) \frac{e^{\delta_P t^*}}{\delta_P} \\
&\quad + \Lambda(0+) [v(1) f(0) - v(1) e^{-\delta_A t^*} f(0) - v(1) f(0)] \\
&= \int_0^\infty (f(t) - f(0)) e^{-\delta_P t} dt + \left[\frac{e^{\delta_P t^*}}{\delta_P} - \Lambda(0+) v(1) e^{-\delta_A t^*} \right] f(0) \\
&= \int_0^\infty (f(t) - f(0)) e^{-\delta_P t} dt + \left[\frac{e^{\delta_P t^*}}{\delta_P} - v(1) e^{-\delta_A t^*} \frac{e^{(\delta_A - \delta_P) t^*}}{v(1) \delta_A} \right] f(0) \\
&= \int_0^\infty (f(t) - f(0)) e^{-\delta_P t} dt + \left[\frac{e^{-\delta_P t^*}}{\delta_P} - v(1) e^{-\delta_A t^*} \frac{e^{(\delta_A - \delta_P) t^*}}{v(1) \delta_A} \right] f(0) \\
&= \int_0^\infty (f(t) - f(0)) e^{-\delta_P t} dt + e^{-\delta_P t^*} \left[\frac{1}{\delta_P} - \frac{1}{\delta_A} \right] f(0) \leq 0
\end{aligned}$$

where the above is negative since $\delta_P \geq \delta_A$ and $f(t) \leq f(0)$. Finally, we should note that $\Lambda(\cdot)$ satisfies the complementary slackness condition as $d\Lambda(t) = 0$ for all $t > 0$ and the incentive constraint is slack, while $d\Lambda(0) = \Lambda(0+) > 0$ and the incentive constraint is binding. This concludes the proof. \square

A.3 Proof of Proposition 2

Proof. For the proof, we will focus on the $\mu_P \geq 1/2$ case. The proof for the case with $\mu_P < 1/2$ is almost identical. Let $\mu_P^*(t)$ be the solution to the following differential equation:

$$-\delta_A = \frac{\mu_P'(t)}{v(\mu_P(t))} \left[\frac{v(1) - v(\mu_P(t))}{1 - \mu_P(t)} - v'(\mu_P(t)) \right]$$

with initial value given by $\mu_P^*(0) = \mu_P \geq 1/2$. This differential equation is derived from the binding incentive constraint 3 using the fact that $G_{P,0}(t) = 1 - \mu_P$ in the first phase. Simple calculations allow us to show that the solution of the above satisfies

$$-\delta_A t = \int_{\mu_P}^{\mu_P^*(t)} \frac{v(1) - v(\mu) - (1 - \mu)v'(\mu)}{(1 - \mu)v(\mu)} d\mu$$

and thus $\mu_P^*(t)$ exists, and it is unique and decreasing. Moreover, we can define t^* as the first time that $\mu_P^*(t^*) = 1/2$ given by

$$\delta_A t^* = \int_{1/2}^{\mu_P} \frac{v(1) - v(\mu) - (1 - \mu)v'(\mu)}{(1 - \mu)v(\mu)} d\mu.$$

We extend $\mu_P^*(t)$ beyond t^* by setting it equal to $1/2$. In addition, the function $\mu_P^*(t)$ defines the engagement probability functions given by

$$G_{P,1}^*(t) = \begin{cases} \frac{\mu_P^*(t)}{\mu_P^*(t) + 1 - \mu_P} & t \leq t^* \\ (1 - \mu_P) e^{-\lambda^*(t-t^*)} & t \geq t^* \end{cases}$$

$$G_{P,0}^*(t) = \begin{cases} 1 - \mu_P & t \leq t^* \\ (1 - \mu_P) e^{-\lambda^*(t-t^*)} & t \geq t^* \end{cases}$$

where $\lambda^* = \delta_A v(1) / v(1/2) - \delta_A$. The variable λ^* is the arrival rate of revelation in the second phase and is calculated using the 3. Additionally, for all values of t , let $\Lambda(t)$ be defined by

$$\Delta_1(t) \Lambda'(t) - \delta_A v(1) \Lambda(t) = -e^{-(\delta_P - \delta_A)t}$$

where $\Delta_1(t) = v(1) - v(\mu_P^*(t)) - (1 - \mu_P^*(t))v'(\mu_P^*(t))$. This is a simple ordinary differential equation whose solution is given by

$$\begin{aligned}\Lambda(t) &= \Gamma(t) \Lambda(0+) - \Gamma(t) \int_0^t \frac{e^{(\delta_A - \delta_P)s}}{\Gamma(s) \Delta_1(s)} ds, \\ \log \Gamma(t) &= \delta_A v(1) \int_0^t \frac{ds}{\Delta_1(s)}.\end{aligned}\tag{17}$$

Finally, we set

$$\begin{aligned}\Lambda(0+) &= \int_0^\infty \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t) \Delta_1(t)} dt \\ \lambda_1 &= \Lambda(0+) \Delta_1(0) > 0 \\ \lambda_0 &= \Lambda(0+) \Delta_0(0) \\ &+ \int_0^{t^*} (e^{-\delta_P t} + e^{-\delta_A t} \Delta_0(t) \Lambda'(t) - \delta_A e^{-\delta_A t} v(1) \Lambda(t)) dt\end{aligned}$$

where $\Lambda(0+)$ exists because for large values of t , $\frac{d}{dt} \log \Gamma(t) = \frac{\delta_A v(1)}{v(1) - v(1/2)} > \delta_A$ and $\Delta_1(t) = v(1) - v(1/2)$.

We can rewrite the derivative of the Lagrangian using integration by part as:

$$\begin{aligned}\partial L(\hat{x}; x) &= \int_0^\infty (G_{P,1}(t) + G_{P,0}(t)) e^{-\delta_P t} dt + \\ &\int_0^\infty e^{-\delta_A t} (\Delta_1(t) G_{P,1}(t) + \Delta_0(t) G_{P,0}(t)) d\Lambda(t) \\ &- v(1) \int_0^\infty \delta_A e^{-\delta_A t} (G_{P,1}(t) + G_{P,0}(t)) \Lambda(t) dt \\ &- \lambda_1 G_{P,1}(0) - \lambda_0 G_{P,0}(0)\end{aligned}$$

where in the above $\Delta_0(t) = v(1) - v(\mu_P^*(t)) + \mu_P^*(t) v'(\mu_P^*(t))$. Since $\mu_P^*(t) = 1/2$ for all $t \geq t^*$ which implies that $\Delta_0(t) = \Delta_1(t) = v(1) - v(1/2)$ for all values of $t \geq t^*$. We can thus

write

$$\begin{aligned}
\partial L(\hat{x}; x) &= \int_0^\infty G_{P,1}(t) [e^{-\delta_P t} + e^{-\delta_A t} \Delta_1(t) \Lambda'(t) - v(1) \delta_A e^{-\delta_A t} \Lambda(t)] dt \\
&\quad \int_0^\infty G_{P,0}(t) [e^{-\delta_P t} + e^{-\delta_A t} \Delta_0(t) \Lambda'(t) - v(1) \delta_A e^{-\delta_A t} \Lambda(t)] dt \\
&\quad + \Delta_1(0) G_{P,1}(0) \Lambda(0+) + \Delta_0(0) G_{P,0}(0) \Lambda(0+) \\
&\quad - \lambda_1 G_{P,1}(0) - \lambda_0 G_{P,0}(0) \\
&= \int_0^{t^*} G_{P,0}(t) [e^{-\delta_P t} + e^{-\delta_A t} \Delta_0(t) \Lambda'(t) - v(1) \delta_A e^{-\delta_A t} \Lambda(t)] dt \\
&\quad + \Delta_1(0) G_{P,1}(0) \Lambda(0+) + \Delta_0(0) G_{P,0}(0) \Lambda(0+) - \lambda_1 G_{P,1}(0) - \lambda_0 G_{P,0}(0) \\
&= \int_0^{t^*} G_{P,0}(t) [e^{-\delta_P t} + e^{-\delta_A t} \Delta_0(t) \Lambda'(t) - v(1) \delta_A e^{-\delta_A t} \Lambda(t)] dt \\
&\quad - G_{P,0}(0) \int_0^{t^*} (e^{-\delta_P t} + e^{-\delta_A t} \Delta_0(t) \Lambda'(t) - \delta_A e^{-\delta_A t} v(1) \Lambda(t)) dt \\
&= \int_0^{t^*} (G_{P,0}(t) - G_{P,0}(0)) [e^{-\delta_P t} + e^{-\delta_A t} \Delta_0(t) \Lambda'(t) - v(1) \delta_A e^{-\delta_A t} \Lambda(t)] dt \\
&\leq \int_0^{t^*} (G_{P,0}(t) - G_{P,0}(0)) \underbrace{[e^{-\delta_P t} + e^{-\delta_A t} \Delta_1(t) \Lambda'(t) - v(1) \delta_A e^{-\delta_A t} \Lambda(t)]}_{=0} dt \\
&= 0
\end{aligned}$$

In the above, the inequality follows from two observations: first, since $G_{P,0}(t)$ is non-increasing, it must be that $G_{P,0}(t) - G_{P,0}(0) \leq 0$; second, $\mu_P^*(t) \geq 1/2$ which then implies that $\Delta_0(t) \geq \Delta_1(t)$ and as we will show later $\Lambda'(t) \geq 0$.

To finish the proof, we need to show that $\Lambda'(t)$, $\lambda_0 \geq 0$. To see these, recall from (17) that

$$\begin{aligned}
\Lambda'(t) &= \Gamma'(t) \Lambda(0+) - \Gamma'(t) \int_0^t \frac{e^{(\delta_A - \delta_P)s}}{\Gamma(s) \Delta_1(s)} ds - \frac{e^{(\delta_A - \delta_P)t}}{\Delta_1(t)} \\
\frac{\Lambda'(t)}{\Gamma(t)} \Delta_1(t) &= \delta_A v(1) \Lambda(0+) - \delta_A v(1) \int_0^t \frac{e^{(\delta_A - \delta_P)s}}{\Gamma(s) \Delta_1(s)} ds - \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} \\
&= \delta_A v(1) \int_0^\infty \frac{e^{(\delta_A - \delta_P)s}}{\Gamma(s) \Delta_1(s)} ds - \delta_A v(1) \int_0^t \frac{e^{(\delta_A - \delta_P)s}}{\Gamma(s) \Delta_1(s)} ds - \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} \\
&= \delta_A v(1) \int_t^\infty \frac{e^{(\delta_A - \delta_P)s}}{\Gamma(s) \Delta_1(s)} ds - \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} \\
&= \int_t^\infty \frac{\Gamma'(s) e^{(\delta_A - \delta_P)s}}{\Gamma(s)^2} ds - \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} = - \int_t^\infty e^{(\delta_A - \delta_P)s} d\left(\frac{1}{\Gamma(s)}\right) - \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} \\
&= - \lim_{s \rightarrow \infty} \frac{e^{(\delta_A - \delta_P)s}}{\Gamma(s)} + \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} + (\delta_A - \delta_P) \int_t^\infty \frac{e^{(\delta_A - \delta_P)s}}{\Gamma(s)} ds - \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} \\
&= (\delta_A - \delta_P) \int_t^\infty \frac{e^{(\delta_A - \delta_P)s}}{\Gamma(s)} ds > 0
\end{aligned}$$

where in the above, we have used the definition of $\Gamma(t)$ and the fact that as $t \rightarrow \infty$, $\Gamma(t) \propto e^{\delta_A \frac{v(1)}{v(1)-v(1/2)} t}$. Finally,

$$\begin{aligned}
\lambda_0 &= \Lambda(0+) \Delta_0(0) + \int_0^{t^*} (e^{-\delta_P t} + e^{-\delta_A t} \Delta_0(t) \Lambda'(t) - \delta_A e^{-\delta_A t} v(1) \Lambda(t)) dt \\
&\geq \int_0^{t^*} (e^{-\delta_P t} + e^{-\delta_A t} \Delta_1(t) + (\Delta_0(t) - \Delta_1(t)) \Lambda'(t) - \delta_A e^{-\delta_A t} v(1) \Lambda(t)) dt \\
&= \int_0^{t^*} (\Delta_0(t) - \Delta_1(t)) \Lambda'(t) dt > 0
\end{aligned}$$

where the inequality follows from the fact that $\Delta_0(t) \geq \Delta_1(t) \geq 0$ for all values of $\mu_P^*(t) \geq 1/2$. This concludes the proof. \square

A.4 Proof of Proposition 3

Proof. As in the statement of the Proposition, we focus on $\ell < 1$; the case with $\ell > 1$ is the mirror of this case. To prove the claim of the proposition, we first define a class of exit probability functions parametrized by $\mu_1 \in [\mu_A, 1]$ – the agent’s belief at the end of phase 1. Given μ_1 , the length of phase 1, t_1 , and end of engagement, t_2 as well as the agent’s belief is given by

$$\begin{aligned} \hat{\mu}_A(t) &= \mu_A : 0 \leq t < t_1 \\ e^{-\delta_A t_1} \left(\frac{\mu_A}{\mu_1} v(\mu_1) + \left(1 - \frac{\mu_A}{\mu_1}\right) v(0) \right) &= v(\mu_A) \\ \frac{d\hat{\mu}_A(t) v(1) - v(\hat{\mu}_A(t)) + \hat{\mu}_A(t) v'(\hat{\mu}_A(t))}{dt} &= \delta_A, t_2 > t \geq t_1 \\ \hat{\mu}_1(t_1) = \mu_1, \hat{\mu}_A(t_2) &= 1. \end{aligned} \tag{18}$$

When $\mu_1 = \mu_A$ the length of phase 1 is 0 while its maximum length is $t^* = \frac{1}{\delta_A} \log \frac{v(1)}{v(\mu_A)}$ when $\mu_1 = 1$. These beliefs are associated with engagement probability functions given by

$$\begin{aligned} \hat{G}_{P,1}(t) &= \begin{cases} \mu_P & 0 \leq t < t_2 \\ 0 & t_2 \leq t \end{cases} \\ \hat{G}_{P,0}(t) &= \begin{cases} 1 - \mu_P & 0 \leq t < t_1 \\ \ell \mu_P \frac{1 - \hat{\mu}_A(t)}{\hat{\mu}_A(t)} & t_1 \leq t \leq t_2 \end{cases} \end{aligned}$$

For each parameter value μ_1 , we also define the following Lagrange multipliers given by

$$\Lambda(t) = \begin{cases} 0 & t = 0 \\ \frac{e^{(\delta_A - \delta_P)t_1}}{\delta_A v(1)} & 0 < t \leq t_1 \end{cases}$$

and for values of $t \in [t_1, t_2]$ the ODE

$$e^{(\delta_A - \delta_P)t} + \Delta_0(t) \Lambda'(t) - \delta_A v(1) \Lambda(t) = 0 \tag{19}$$

holds where $\Delta_0(t) = v(1) - v(\hat{\mu}_A(t)) + \hat{\mu}_A(t) v'(\hat{\mu}_A(t))$ and finally, we set $\Lambda(t) = \Lambda(t_2)$ for values of $t \geq t_2$. Let μ_1^* be the value of the parameter μ_1 such that either the above multiplier satisfies $\Lambda(t_2) = e^{(\delta_A - \delta_P)t_2} / (\ell \delta_A v(1))$ and if such value does not exist, we set $\mu_1^* = \mu_A$. Later, we will show that there exists at most one such value of μ_1 and that $\mu_1^* < 1$ has to hold. In other words, if $\mu_1 = 1$, we will show that $\Lambda(t_2) < e^{(\delta_A - \delta_P)t_2} / (\ell \delta_A v(1))$. Let t_1^* and t_2^* be the associated

end of phase 1 and phase 2. Finally, we define λ_1, λ_0 as

$$\begin{aligned}\lambda_0 &= \int_0^{t_1^*} e^{-\delta_P s} ds + \Lambda(0+) (v(1) e^{-\delta_A t_1^*} - v(\mu_A) + \mu_{AV}'(\mu_A)) \\ \lambda_1 &= \int_0^{t_2^*} e^{-\delta_P s} ds + \ell \Lambda(0+) (v(1) - v(\mu_A) - (1 - \mu_A) v'(\mu_A)) \\ &\quad + \ell \int_0^{t_2^*} e^{-\delta_A s} (\Delta_1(s) \Lambda'(s) - \delta_{AV}(1) \Lambda(s)) ds\end{aligned}$$

For this allocation, we show that $\partial L(\hat{x}; \hat{x}) = 0$ and $\partial L(\hat{x}; x) \leq 0$ for all non-increasing and positive pair of functions $x = (G_{P,0}(t), G_{P,1}(t))$. Moreover, the multipliers $\Lambda, \lambda_1, \lambda_0$ need to satisfy the complementary slackness and $\Lambda(t)$ has to be non-decreasing.

- 1. Complementary Slackness (CS):** The incentive constraint is binding when $t = 0$ and when $t \in [t_1^*, t_2^*]$. Since $\Lambda(t)$ is flat for $t \in (0, t_1^*)$, for CS to be satisfied, we need $\Lambda(0+) = \frac{e^{(\delta_A - \delta_P)t_1^*}}{\delta_{AV}(1)}$ to be positive, which holds, and $\Lambda(t)$ to be increasing when $t \in [t_1^*, t_2^*]$. The ODE in 19 has an exact solution given by

$$\begin{aligned}\Lambda(t) &= \Gamma(t) \left[\frac{e^{(\delta_A - \delta_P)t_1^*}}{\delta_{AV}(1)} - \int_{t_1}^t \frac{e^{(\delta_A - \delta_P)s}}{\Delta_0(s) \Gamma(s)} ds \right] \\ \log \Gamma(t) &= \int_{t_1}^t \frac{\delta_{AV}(1)}{\Delta_0(s)} ds\end{aligned} \tag{20}$$

Differentiating $\Lambda(t)$, we have

$$\begin{aligned}\Lambda'(t) &= \frac{\delta_{AV}(1)}{\Delta_0(t)} \Gamma(t) \left[\frac{e^{(\delta_A - \delta_P)t_1^*}}{\delta_{AV}(1)} - \int_{t_1}^t \frac{e^{(\delta_A - \delta_P)s}}{\Delta_0(s) \Gamma(s)} ds \right] \\ &\quad - \frac{e^{(\delta_A - \delta_P)t}}{\Delta_0(t)} \\ \rightarrow \Lambda'(t) \frac{\Delta_0(t)}{\Gamma(t) \delta_{AV}(1)} &= \frac{e^{(\delta_A - \delta_P)t_1^*}}{\delta_{AV}(1)} - \int_{t_1}^t \frac{e^{(\delta_A - \delta_P)s}}{\Delta_0(s) \Gamma(s)} ds - \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t) \delta_{AV}(1)} \\ \rightarrow \frac{d}{dt} \left(\Lambda'(t) \frac{\Delta_0(t)}{\Gamma(t) \delta_{AV}(1)} \right) &= - \frac{e^{(\delta_A - \delta_P)t}}{\Delta_0(t) \Gamma(t)} - (\delta_A - \delta_P) \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t) \delta_{AV}(1)} + \frac{\Gamma'(t)}{\Gamma(t)} \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t) \delta_{AV}(1)} \\ &= - \frac{e^{(\delta_A - \delta_P)t}}{\Delta_0(t) \Gamma(t)} - (\delta_A - \delta_P) \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t) \delta_{AV}(1)} + \frac{\delta_{AV}(1)}{\Delta_0(t)} \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t) \delta_{AV}(1)} \\ &= (\delta_P - \delta_A) \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t) \delta_{AV}(1)} \geq 0\end{aligned}$$

Since $\Delta_0(t), \Gamma(t) \geq 0$ for all values of t , the above implies that if $\Lambda'(t_1^*) \geq 0$, then for all

values of $t \geq t_1$, $\Lambda'(t) \geq 0$. At $t = t_1^*$ we have

$$\begin{aligned} e^{(\delta_A - \delta_P)t_1^*} + \Delta_0(t_1^*) \Lambda'(t_1^*) - \delta_A v(1) \Lambda(t_1^*) &= 0 \\ \rightarrow e^{(\delta_A - \delta_P)t_1^*} + \Delta_0(t_1^*) \Lambda'(t_1^*) - \delta_A v(1) \frac{e^{(\delta_A - \delta_P)t_1^*}}{\delta_A v(1)} &= 0 \\ \Delta_0(t_1^*) \Lambda'(t_1^*) &= 0 \end{aligned}$$

Since $\mu_1^* \geq \mu_A > 0$, we must have that $\Delta_0(t_1^*) > 0$ and hence $\Lambda'(t_1^*) = 0$.

Additionally, we have to show that $\lambda_1, \lambda_0 \geq 0$. We have

$$\begin{aligned} \lambda_0 &= \int_0^{t_1^*} e^{-\delta_P s} ds + \Lambda(0+) (v(1) e^{-\delta_A t_1^*} - v(\mu_A) + \mu_A v'(\mu_A)) \\ &= \frac{1 - e^{-\delta_P t_1^*}}{\delta_P} + \frac{e^{(\delta_A - \delta_P)t_1^*}}{\delta_A v(1)} (v(1) e^{-\delta_A t_1^*} - v(1) + v(1) - v(\mu_A) + \mu_A v'(\mu_A)) \\ &= \frac{1 - e^{-\delta_P t_1^*}}{\delta_P} + \frac{e^{-\delta_P t_1^*} - e^{(\delta_A - \delta_P)t_1^*}}{\delta_A} + e^{(\delta_A - \delta_P)t_1^*} \frac{v(1) - v(\mu_A) + \mu_A v'(\mu_A)}{\delta_A v(1)} \\ &= \frac{1 - e^{-\delta_P t_1^*}}{\delta_P} + e^{-\delta_P t_1^*} \frac{1 - e^{\delta_A t_1^*}}{\delta_A} + e^{(\delta_A - \delta_P)t_1^*} \frac{v(1) - v(\mu_A) + \mu_A v'(\mu_A)}{\delta_A v(1)} \\ &= e^{-\delta_P t_1^*} \left(\frac{e^{\delta_P t_1^*} - 1}{\delta_P} - \frac{e^{\delta_A t_1^*} - 1}{\delta_A} \right) + e^{(\delta_A - \delta_P)t_1^*} \frac{v(1) - v(\mu_A) + \mu_A v'(\mu_A)}{\delta_A v(1)} \end{aligned}$$

The first expression above is positive since $\delta_P > \delta_A$ and the second expression is positive since v is convex and positive. Therefore, $\lambda_0 \geq 0$. Additionally,

$$\begin{aligned} \lambda_1 &= \int_0^{t_2^*} e^{-\delta_P s} ds + \ell \Lambda(0+) (v(1) - v(\mu_A) - (1 - \mu_A) v'(\mu_A)) \\ &\quad + \ell \int_0^{t_2^*} e^{-\delta_A s} (\Delta_1(s) \Lambda'(s) - \delta_A v(1) \Lambda(s)) ds \\ &= \int_0^{t_1^*} (e^{-\delta_P s} - \ell \delta_A v(1) \Lambda(0+) e^{-\delta_A s}) ds + \ell \Lambda(0+) \Delta_1(0) \\ &\quad + \int_{t_1^*}^{t_2^*} [e^{-\delta_P s} + \ell e^{-\delta_A s} (\Delta_1(s) \Lambda'(s) - \delta_A v(1) \Lambda(s))] ds \\ &= e^{-\delta_P t_1^*} \int_0^{t_1^*} \left(\underbrace{e^{\delta_P(t_1^* - s)} - \ell e^{\delta_A(t_1^* - s)}}_{\geq 0} \right) ds + \ell \Lambda(0+) \Delta_1(0) \\ &\quad + \int_{t_1^*}^{t_2^*} e^{-\delta_P s} [1 + \ell e^{(\delta_P - \delta_A)s} (\Delta_1(s) \Lambda'(s) - \delta_A v(1) \Lambda(s))] ds \end{aligned}$$

The first two terms in the above expression are positive since $\delta_P > \delta_A, \ell < 1$. Additionally,

since Λ satisfies the ODE (19), we can write the integrand in the last integral as

$$\begin{aligned}
& 1 + \ell e^{(\delta_P - \delta_A)s} (\Delta_1(s) \Lambda'(s) - \delta_A v(1) \Lambda(s)) = \\
& 1 + \ell e^{(\delta_P - \delta_A)s} \left(\frac{\Delta_1(s)}{\Delta_0(s)} (\delta_A v(1) \Lambda(s) - e^{(\delta_A - \delta_P)s}) - \delta_A v(1) \Lambda(s) \right) = \\
& 1 - \ell + \ell e^{(\delta_P - \delta_A)s} \left(\frac{\Delta_1(s)}{\Delta_0(s)} (\delta_A v(1) \Lambda(s) - e^{(\delta_A - \delta_P)s}) + e^{(\delta_P - \delta_A)s} - \delta_A v(1) \Lambda(s) \right) = \\
& 1 - \ell + \ell \frac{\Delta_1(s) - \Delta_0(s)}{\Delta_0(s)} (e^{(\delta_P - \delta_A)s} \delta_A v(1) \Lambda(s) - 1) = \\
& 1 - \ell - \ell \frac{v'(\hat{\mu}_A(s))}{\Delta_0(s)} (e^{(\delta_P - \delta_A)s} \delta_A v(1) \Lambda(s) - 1) \quad (21)
\end{aligned}$$

In the above, since $\delta_P > \delta_A$, $\Lambda'(s) \geq 0$, and $\Lambda(t_1) \delta_A v(1) e^{(\delta_P - \delta_A)s} = 1$, $e^{(\delta_P - \delta_A)s} \delta_A v(1) \Lambda(s) - 1$ is positive and increasing. If $\hat{\mu}_A \leq 1/2$, then $v' \leq 0$ which implies that the above is positive. Finally, when $\hat{\mu}_A \geq 1/2$, $v' \geq 0$

$$\frac{d}{ds} \frac{v'(\hat{\mu}_A(s))}{\Delta_0(s)} = \frac{v''(\hat{\mu}_A(s)) (v(1) - v(\hat{\mu}_A(s)))}{\Delta_0(s)^2} \hat{\mu}'_A(s) \geq 0.$$

As a result, $\frac{v'(\hat{\mu}_A(s))}{\Delta_0(s)} (e^{(\delta_P - \delta_A)s} \delta_A v(1) \Lambda(s) - 1)$ is increasing, and therefore the integrand in (21) is decreasing in s . Its value at t_2^* is given by

$$1 - \ell - \ell \frac{v'(1)}{v'(1)} (e^{(\delta_P - \delta_A)t_2^*} \delta_A v(1) \Lambda(t_2^*) - 1) = 1 - \ell e^{(\delta_P - \delta_A)t_2^*} \delta_A v(1) \Lambda(t_2^*).$$

By construction of μ_1^* , we must have that $1 - \ell e^{(\delta_P - \delta_A)t_2^*} \delta_A v(1) \Lambda(t_2^*) \geq 0$. Since the integrand in (21) is decreasing in s , it must always be positive. As a result, the last integral in the expression for λ_1 is positive and $\lambda_1 \geq 0$.

2. **Optimality Condition (14):** Since $\partial L(\hat{x}; x)$ is linear in x , it is sufficient to show $\partial L(\hat{x}; x) \leq 0$ for $x = (f(t), 0)$ and $x = (0, f(t))$ with f an arbitrary non-increasing and non-negative function. For such values, we have

$$\begin{aligned}
\partial L(\hat{x}; (f(t), 0)) &= \int_0^\infty f(t) e^{-\delta_P t} dt + \int_0^\infty e^{-\delta_A t} \Delta_0(t) f(t) d\Lambda(t) \\
&\quad - v(1) \int_0^\infty \delta_A e^{-\delta_A t} f(t) \Lambda(t) dt - \lambda_0 f(0)
\end{aligned}$$

Given our construction of the allocation and the multipliers, we can rewrite the above as

$$\begin{aligned}
\partial L(\hat{x}; (f(t), 0)) &= \int_0^\infty f(t) e^{-\delta_P t} dt + \Delta_0(0) f(0) \Lambda(0+) + \\
&\quad \int_{t_1^*}^{t_2^*} e^{-\delta_A t} \Lambda'(t) \Delta_0(t) f(t) dt - \delta_{AV}(1) \int_0^\infty e^{-\delta_A t} f(t) \Lambda(t) dt \\
&\quad - \Lambda(0+) (v(1) e^{-\delta_A t_1^*} - v(\mu_A) + \mu_A v'(\mu_A)) f(0) - \int_0^{t_1^*} e^{-\delta_P s} ds f(0) \\
&= \int_0^{t_2^*} f(t) e^{-\delta_P t} dt + (v(1) - v(\mu_A) + \mu_A v'(\mu_A)) f(0) \Lambda(0+) \\
&\quad + \int_{t_1^*}^{t_2^*} e^{-\delta_A t} \Lambda'(t) \Delta_0(t) f(t) dt - \delta_{AV}(1) \int_{t_1^*}^{t_2^*} e^{-\delta_A t} f(t) \Lambda(t) dt \\
&\quad - \delta_{AV}(1) \int_0^{t_1^*} e^{-\delta_A t} f(t) dt \Lambda(0+) \\
&\quad - \Lambda(t_1) (v(1) e^{-\delta_A t_1^*} - v(\mu_A) + \mu_A v'(\mu_A)) f(0) \\
&\quad + \int_{t_2^*}^\infty f(t) [e^{-\delta_P t} - e^{-\delta_A t} \delta_{AV}(1) \Lambda(t)] dt \\
&= \int_0^{t_1^*} [e^{-\delta_P t} (f(t) - f(0)) + \Lambda(0+) e^{-\delta_A t} (\delta_{AV}(1) f(0) - \delta_{AV}(1) f(t))] dt \\
&\quad + \int_{t_1^*}^{t_2^*} f(t) \left[\underbrace{e^{-\delta_P t} + e^{-\delta_A t} (\Lambda'(t) \Delta_0(t) - \delta_{AV}(1) \Lambda(t))}_{=0 \text{ by ODE (19)}} \right] dt \\
&\quad + \int_{t_2^*}^\infty e^{-\delta_P t} f(t) \left[1 - e^{(\delta_P - \delta_A)t} \delta_{AV}(1) \underbrace{\Lambda(t)}_{=\Lambda(t_2^*)} \right] dt \\
&= \int_0^{t_1^*} (f(t) - f(0)) e^{-\delta_P t} \left[1 - \frac{e^{(\delta_A - \delta_P)t_1^*}}{\delta_{AV}(1)} e^{(\delta_P - \delta_A)t} \delta_{AV}(1) \right] dt + \\
&\quad \int_{t_2^*}^\infty e^{-\delta_P t} f(t) \left[1 - e^{(\delta_P - \delta_A)t} \delta_{AV}(1) \frac{e^{(\delta_A - \delta_P)t_2^*}}{\ell \delta_{AV}(1)} \right] dt \\
&= \int_0^{t_1^*} \left(\underbrace{f(t) - f(0)}_{\leq 0} \right) e^{-\delta_P t} \left[1 - \underbrace{e^{(\delta_P - \delta_A)(t - t_1^*)}}_{\leq 1} \right] dt \\
&\quad + \int_{t_2^*}^\infty e^{-\delta_P t} f(t) \left[1 - \underbrace{\frac{e^{(\delta_P - \delta_A)(t - t_2^*)}}{\ell}}_{\geq 1} \right] dt \leq 0
\end{aligned}$$

where the last inequality follows from $\delta_P > \delta_A$ and $\ell < 1$.

Similarly, for $x = (0, f(t))$, we can write

$$\begin{aligned}
\partial L(\hat{x}; (0, f(t))) &= \int_0^\infty f(t) e^{-\delta_P t} dt + \\
&\quad \ell \Delta_1(0) f(0) \Lambda(0+) + \ell \int_{t_1^*}^{t_2^*} e^{-\delta_A t} \Lambda'(t) \Delta_1(t) f(t) dt \\
&\quad - \ell \delta_A v(1) \int_0^\infty e^{-\delta_A t} f(t) \Lambda(t) dt \\
&\quad - f(0) \int_0^{t_2^*} e^{-\delta_P s} ds - f(0) \ell \Lambda(0+) (v(1) - v(\mu_A) - (1 - \mu_A) v'(\mu_A)) \\
&\quad - f(0) \int_0^{t_2^*} [e^{-\delta_P s} + \ell e^{-\delta_A s} (\Delta_1(s) \Lambda'(s) - \delta_A v(1) \Lambda(s))] ds \\
&= \int_0^{t_2^*} e^{-\delta_P t} (f(t) - f(0)) dt + \\
&\quad \ell \int_0^{t_2^*} e^{-\delta_A t} [\Lambda'(t) \Delta_1(t) - \delta_A v(1) \Lambda(t)] (f(t) - f(0)) dt \\
&\quad + \int_{t_2^*}^\infty e^{-\delta_P t} f(t) [1 - \ell e^{(\delta_P - \delta_A)t} \delta_A v(1) \Lambda(t)] dt
\end{aligned}$$

Since $\Lambda(t) = \Lambda(t_2) = e^{(\delta_A - \delta_P)t_2^*} / (\delta_A \ell v(1))$, the last integral is negative. It is thus sufficient to show that the first two terms add up to a negative value. We can use integration by parts to write

$$\begin{aligned}
&\int_0^{t_2^*} e^{-\delta_P t} (f(t) - f(0)) dt + \\
&\ell \int_0^{t_2^*} e^{-\delta_A t} [\Lambda'(t) \Delta_1(t) - \delta_A v(1) \Lambda(t)] (f(t) - f(0)) dt = \\
&\int_0^{t_2^*} (f(0) - f(t)) d \left(\int_t^{t_2^*} [e^{-\delta_P s} + \ell e^{-\delta_A s} (\Lambda'(s) \Delta_1(s) - \delta_A v(1) \Lambda(s))] ds \right) = \\
&\int_0^{t_2^*} \int_t^{t_2^*} [e^{-\delta_P s} + \ell e^{-\delta_A s} (\Lambda'(s) \Delta_1(s) - \delta_A v(1) \Lambda(s))] ds df(t)
\end{aligned}$$

Earlier we have shown that $e^{-\delta_P s} + \ell e^{-\delta_A s} (\Lambda'(s) \Delta_1(s) - \delta_A v(1) \Lambda(s)) \geq 0$, see the discussion following (21). Hence, since $df(t) \leq 0$, that is, f is decreasing, the above integral is negative, which is the desired result.

3. **Optimality Condition (15):** In the proof of Theorem 1, we showed that Equation (15) boils down to checking $\lambda_1 \mu_P + \lambda_0 (1 - \mu_P) = \int_0^{t_2^*} e^{-\delta_P t} G_P(t) dt$. Given the values of λ_1

and λ_0 , we have

$$\begin{aligned}\lambda_0 &= \int_0^{t_1^*} e^{-\delta_P s} ds + \Lambda(0+) (v(1) - v(\mu_A) + \mu_A v'(\mu_A)) - \Lambda(0+) \delta_A v(1) \int_0^{t_1^*} e^{-\delta_A s} ds \\ \lambda_1 &= \int_0^{t_2^*} e^{-\delta_P s} ds + \ell \Lambda(0+) (v(1) - v(\mu_A) - (1 - \mu_A) v'(\mu_A)) \\ &\quad + \ell \int_0^{t_2^*} e^{-\delta_A s} (\Delta_1(s) \Lambda'(s) - \delta_A v(1) \Lambda(s)) ds\end{aligned}$$

Since $G_{P,0}(t) = 1 - \mu_P, \forall t \leq t_1^*, G_{P,1}(t) = \mu_P, \forall t \leq t_2^*$, we can use the above to write

$$\begin{aligned}(1 - \mu_P) \lambda_0 + \mu_P \lambda_1 &= \int_0^{t_1^*} G_{P,0}(s) e^{-\delta_P s} ds + \Lambda(0+) (v(1) - v(\mu_A) + \mu_A v'(\mu_A)) (1 - \mu_P) \\ &\quad - \Lambda(0+) \delta_A v(1) \int_0^{t_1^*} G_{P,0}(s) e^{-\delta_A s} ds \\ &\quad + \int_0^{t_2^*} G_{P,1}(s) e^{-\delta_P s} ds + \ell \Lambda(0+) (v(1) - v(\mu_A) - (1 - \mu_A) v'(\mu_A)) \mu_P \\ &\quad + \ell \int_0^{t_2^*} G_{P,1}(s) e^{-\delta_A s} (\Delta_1(s) \Lambda'(s) - \delta_A v(1) \Lambda(s)) ds\end{aligned}$$

Recall the definition of $\ell = \frac{\mu_A}{1 - \mu_A} \frac{1 - \mu_P}{\mu_P}$ which implies that $\ell(1 - \mu_A) \mu_P = (1 - \mu_P) \mu_A$ and hence, we can write the above as

$$\begin{aligned}(1 - \mu_P) \lambda_0 + \mu_P \lambda_1 &= \int_0^{t_1^*} G_{P,0}(s) e^{-\delta_P s} ds + \Lambda(0+) (v(1) - v(\mu_A)) (\ell G_{P,1}(0) + G_{P,0}(0)) \\ &\quad - \Lambda(0+) \delta_A v(1) \int_0^{t_1^*} G_{P,0}(s) e^{-\delta_A s} ds + \int_0^{t_2^*} G_{P,1}(s) e^{-\delta_P s} ds \\ &\quad + \ell \int_0^{t_2^*} G_{P,1}(s) e^{-\delta_A s} (\Delta_1(s) \Lambda'(s) - \delta_A v(1) \Lambda(s)) ds\end{aligned}$$

Now, given that $\Lambda(t)$ satisfies the ODE (19), we can multiply the ODE by $G_{P,0}(t)$ and integrate over time. The resulting expression is 0 and thus can be added to the above.

Hence,

$$\begin{aligned}
(1 - \mu_P) \lambda_0 + \mu_P \lambda_1 &= \int_0^{t_1^*} G_{P,0}(s) e^{-\delta_P s} ds + \Lambda(0+) (v(1) - v(\mu_A)) (\ell G_{P,1}(0) + G_{P,0}(0)) \\
&\quad - \Lambda(0+) \delta_A v(1) \int_0^{t_1^*} G_{P,0}(s) e^{-\delta_A s} ds + \int_0^{t_2^*} G_{P,1}(s) e^{-\delta_P s} ds \\
&\quad + \ell \int_0^{t_2^*} G_{P,1}(s) e^{-\delta_A s} (\Delta_1(s) \Lambda'(s) - \delta_A v(1) \Lambda(s)) ds + \\
&\quad \int_{t_1^*}^{t_2^*} G_{P,0}(s) e^{-\delta_P s} ds + \int_{t_1^*}^{t_2^*} G_{P,0}(s) e^{-\delta_A s} (\Delta_0(s) \Lambda'(s) - \delta_A v(1) \Lambda(s)) ds \\
&= \int_0^{t_2^*} G_P(s) e^{-\delta_A s} ds \\
&\quad + \int_0^{t_2^*} (\ell G_{P,1}(s) + G_{P,0}(s)) e^{-\delta_A s} (v(1) - v(\hat{\mu}_A(s))) d\Lambda(s) \\
&\quad - \delta_A v(1) \int_0^{t_2^*} e^{-\delta_A s} \Lambda(s) (\ell G_{P,1}(s) + G_{P,0}(s)) ds \\
&= \int_0^{t_2^*} G_P(s) e^{-\delta_A s} ds \\
&\quad + \int_0^{t_2^*} (\ell G_{P,1}(s) + G_{P,0}(s)) e^{-\delta_A s} (v(1) - v(\hat{\mu}_A(s))) d\Lambda(s) \\
&\quad - \delta_A v(1) \int_0^{t_2^*} \int_t^{t_2^*} e^{-\delta_A s} (\ell G_{P,1}(s) + G_{P,0}(s)) ds d\Lambda(t)
\end{aligned}$$

Using integration by part the above simplifies to:

$$\begin{aligned}
&\int_0^{t_2^*} G_P(s) e^{-\delta_A s} ds + \\
&\underbrace{\int_0^{t_2^*} \left[G_A(t) e^{-\delta_A t} (v(1) - v(\hat{\mu}_A(s))) - \delta_A v(1) \int_t^{t_2^*} G_A(s) e^{-\delta_A s} ds \right] d\Lambda(s)}_{=0 \text{ by complementary slackness}} \\
&= \int_0^{t_2^*} G_P(s) e^{-\delta_A s} ds
\end{aligned}$$

Finally, to finish the proof, we will show that there is at most one μ_1 for which $\Lambda(t_2) = \frac{e^{(\delta_A - \delta_P)t_2}}{\delta_A \ell v(1)}$ holds and that if such μ_1 exist it must be less than 1. To see this, suppose that $\mu_1 = 1$. In this case, $t_1 = t_2 = \frac{1}{\delta_A} \log \frac{v(1)}{v(\mu_A)}$. Note under this assumption, engagement strategy in phase 2 does not exist and hence we must have that $\Lambda(t_1) = \Lambda(t_2) = \frac{e^{(\delta_A - \delta_P)t_1}}{\delta_A v(1)}$. Since $\ell < 1$, this is less than $\frac{e^{(\delta_A - \delta_P)t_2}}{\delta_A \ell v(1)}$ which implies that if μ_1^* exists, it cannot be equal to 1. In order to show the

existence of μ_1^* , we will show that as we increase μ_1 , $\frac{e^{(\delta_A - \delta_P)t_2}}{\delta_A \ell v(1)} - \Lambda(t_2)$ decreases. As a result, either there is a unique value μ_1^* under which $\frac{e^{(\delta_A - \delta_P)t_2}}{\delta_A \ell v(1)} - \Lambda(t_2) = 0$ or that for all values of μ_1 , $\frac{e^{(\delta_A - \delta_P)t_2}}{\delta_A \ell v(1)} - \Lambda(t_2) > 0$ in which case $\mu_1^* = \mu_A$. To avoid further lengthening the proof of this proposition, we show this in the proof of Corollary 1 below. \square

A.5 Proof of Corollary 1

Proof. We first focus on the case with $\ell < 1$ or $\mu_A < \mu_P$. In this case, we first show that $\frac{1}{\delta_A \ell v(1)} - \frac{\Lambda(t_2)}{e^{(\delta_A - \delta_P)t_2}}$ is strictly decreasing in μ_1 . As a result, to have no phase 1, we would need $\frac{e^{(\delta_A - \delta_P)t_2}}{\delta_A \ell v(1)} - \Lambda(t_2) > 0$ for $\mu_1 = \mu_A$. For a given μ_1 , since $\hat{\mu}_A$ is strictly increasing over phase 2, we can do a change of variable and write the Lagrange multiplier $\Lambda(t)$ as a function of beliefs so that $L(\hat{\mu}_A(t)) = \Lambda(t)$. The law of motion for $\hat{\mu}_A(t)$ in phase 2 can be written as:

$$\frac{d\mu}{dt} \frac{v(1) - v(\mu) + \mu v'(\mu)}{\mu v(\mu)} = \delta_A.$$

This implies that

$$\Psi(\mu) = \int_{\mu}^1 \frac{v(1) - v(x) + xv'(x)}{xv(x)} dx = \delta_A(t_2 - t) \Rightarrow t = t_2 - \frac{\Psi(\mu)}{\delta_A}$$

Given the above, we have

$$e^{\delta_A t_1} = e^{\delta_A t_2 - \Psi(\mu_1)} = \frac{(\mu_1 - \mu_A)v(1) + \mu_A v(\mu_1)}{\mu_1 v(\mu_A)}$$

We can use this change of variable and write the solution for $\Lambda(t)$ in (20) as

$$L(\mu) = \hat{\Gamma}(\mu) \left[\frac{e^{(\delta_A - \delta_P)(t_2 - \Psi(\mu_1)/\delta_A)}}{\delta_A v(1)} - \int_{\mu_1}^{\mu} \frac{e^{(\delta_A - \delta_P)(t_2 - \Psi(x)/\delta_A)}}{v(1) - v(x) + xv'(x)} \frac{1}{\hat{\Gamma}(x)} d\left(-\frac{\Psi(x)}{\delta_A}\right) \right]$$

$$\log \hat{\Gamma}(\mu) = \int_{\mu_1}^{\mu} \frac{\delta_A v(1)}{v(1) - v(x) + xv'(x)} d\left(-\frac{\Psi(x)}{\delta_A}\right)$$

Using the definition of Ψ , we can rewrite the above as

$$L(\mu) = \hat{\Gamma}(\mu) \left[\frac{e^{(\delta_A - \delta_P)(t_2 - \Psi(\mu_1)/\delta_A)}}{\delta_A v(1)} - \int_{\mu_1}^{\mu} \frac{e^{(\delta_A - \delta_P)(t_2 - \Psi(x)/\delta_A)}}{\delta_A xv(x)} \frac{1}{\hat{\Gamma}(x)} dx \right]$$

$$\log \hat{\Gamma}(\mu) = \int_{\mu_1}^{\mu} \frac{v(1)}{xv(x)} dx$$

Hence,

$$\frac{\Lambda(t_2)}{e^{(\delta_A - \delta_P)t_2}} = \frac{L(1)}{e^{(\delta_A - \delta_P)t_2}} = \hat{\Gamma}(1) \left[\frac{e^{-(\delta_A - \delta_P)\Psi(\mu_1)/\delta_A}}{\delta_A v(1)} - \int_{\mu_1}^1 \frac{e^{-(\delta_A - \delta_P)\Psi(x)/\delta_A}}{\delta_A xv(x)} \frac{1}{\hat{\Gamma}(x)} dx \right]$$

Taking a derivative of the above with respect to μ_1 , we have

$$\begin{aligned}
\frac{d}{d\mu_1} \frac{\Lambda(t_2)}{e^{(\delta_A - \delta_P)t_2}} &= -\frac{v(1)}{\mu_1 v(\mu_1)} \hat{\Gamma}(1) \left[\frac{e^{-(\delta_A - \delta_P)\Psi(\mu_1)/\delta_A}}{\delta_A v(1)} - \int_{\mu_1}^1 \frac{e^{-(\delta_A - \delta_P)\Psi(x)/\delta_A}}{\delta_A x v(x)} \frac{1}{\hat{\Gamma}(x)} dx \right] \\
&+ \hat{\Gamma}(1) \frac{v(1) - v(\mu_1) + \mu_1 v'(\mu_1)}{\mu_1 v(\mu_1)} \frac{\delta_A - \delta_P}{\delta_A} \frac{e^{-(\delta_A - \delta_P)\Psi(\mu_1)/\delta_A}}{\delta_A v(1)} \\
&+ \hat{\Gamma}(1) \frac{e^{-(\delta_A - \delta_P)\Psi(\mu_1)/\delta_A}}{\delta_A \mu_1 v(\mu_1)} \\
&- \hat{\Gamma}(1) \frac{v(1)}{\mu_1 v(\mu_1)} \int_{\mu_1}^1 \frac{e^{-(\delta_A - \delta_P)\Psi(x)/\delta_A}}{\delta_A x v(x)} \frac{1}{\hat{\Gamma}(x)} dx \\
&= \hat{\Gamma}(1) \frac{v(1) - v(\mu_1) + \mu_1 v'(\mu_1)}{\mu_1 v(\mu_1)} \underbrace{\frac{\delta_A - \delta_P}{\delta_A}}_{<0} \frac{e^{-(\delta_A - \delta_P)\Psi(\mu_1)/\delta_A}}{\delta_A v(1)} < 0
\end{aligned}$$

Hence, either there exists a unique $\mu_1^* \in (\mu_A, 1)$ such that $\frac{\Lambda(t_2)}{e^{(\delta_A - \delta_P)t_2}} = \frac{1}{\delta_A \ell v(1)}$ or that

$$\frac{\Lambda(t_2)}{e^{(\delta_A - \delta_P)t_2}} \leq \frac{1}{\delta_A \ell v(1)} \text{ at } \mu_1 = \mu_A$$

in which case $\mu_1^* = \mu_A$ and phase 1 is of length 0. That is

$$e^{\int_{\mu_A}^1 \frac{v(1)}{xv(x)} dx} \left[\frac{e^{-(\delta_A - \delta_P)\Psi(\mu_A)/\delta_A}}{\delta_A v(1)} - \int_{\mu_A}^1 \frac{e^{-(\delta_A - \delta_P)\Psi(x)/\delta_A}}{\delta_A x v(x)} e^{-\int_{\mu_A}^x \frac{v(1)}{zv(z)} dz} dx \right] \leq \frac{1}{\delta_A \ell v(1)}$$

or

$$e^{\int_{\mu_A}^1 \frac{v(1)}{xv(x)} dx} e^{-(\delta_A - \delta_P)\Psi(\mu_A)/\delta_A} - v(1) \int_{\mu_A}^1 \frac{e^{-(\delta_A - \delta_P)\Psi(x)/\delta_A}}{xv(x)} e^{\int_x^1 \frac{v(1)}{zv(z)} dz} dx \leq 1/\ell \quad (22)$$

The derivative of the above with respect to μ_A is

$$\begin{aligned}
&\left(\frac{\delta_A - \delta_P}{\delta_A} \frac{v(1) - v(\mu_A) + \mu_A v'(\mu_A)}{\mu_A v(\mu_A)} - \frac{v(1)}{\mu_A v(\mu_A)} \right) e^{\int_{\mu_A}^1 \frac{v(1)}{xv(x)} dx} e^{-(\delta_A - \delta_P)\Psi(\mu_A)/\delta_A} + \\
&\quad v(1) \frac{e^{-(\delta_A - \delta_P)\Psi(\mu_A)/\delta_A}}{\mu_A v(\mu_A)} e^{\int_{\mu_A}^1 \frac{v(1)}{zv(z)} dz} = \\
&\quad \frac{\delta_A - \delta_P}{\delta_A} \frac{v(1) - v(\mu_A) + \mu_A v'(\mu_A)}{\mu_A v(\mu_A)} e^{\int_{\mu_A}^1 \frac{v(1)}{xv(x)} dx} e^{-(\delta_A - \delta_P)\Psi(\mu_A)/\delta_A} < 0
\end{aligned}$$

Additionally an increase in μ_A decreases ℓ . This implies that the inequality (22) holds for low values of μ_A .

When $\mu_A > \mu_P$ or $\ell > 1$, the law of motion for $\hat{\mu}_A(t)$ in phase 2 can be written as

$$\frac{d\mu}{dt} \frac{v(1) - v(\mu) - (1-\mu)v'(\mu)}{(1-\mu)v(\mu)} = -\delta_A$$

This implies that

$$\Phi(\mu) = \int_0^\mu \frac{v(1) - v(x) - (1-x)v'(x)}{(1-x)v(x)} dx = \delta_A(t_2 - t) \Rightarrow t = t_2 - \frac{\Phi(\mu)}{\delta_A}$$

Given the above, we have

$$e^{\delta_A t_1} = e^{\delta_A t_2 - \Phi(\mu_1)} = \frac{(\mu_A - \mu_1)v(1) + (1 - \mu_A)v(\mu_1)}{(1 - \mu_1)v(\mu_A)}$$

At the same time, the ODE characterizing the Lagrange multiplier in phase 2 is given by

$$e^{(\delta_A - \delta_P)t} + \ell \Lambda'(t) \Delta_1(t) - \ell \delta_A v(1) \Lambda(t) = 0$$

with $\Lambda(t_1) = e^{(\delta_A - \delta_P)t_1} / (\ell \delta_A v(1))$. In this case, the solution of this ODE is given by

$$\Lambda(t) = \Omega(t) \left[\frac{e^{(\delta_A - \delta_P)t_1}}{\ell \delta_A v(1)} - \int_{t_1}^t \frac{1}{\Omega(s)} \frac{e^{(\delta_A - \delta_P)s}}{\ell \Delta_1(s)} ds \right]$$

$$\log \Omega(t) = \int_{t_1}^t \frac{\delta_A v(1)}{\Delta_1(s)} ds.$$

We can use a change of variable as above and write

$$\frac{L(\mu)}{e^{(\delta_A - \delta_P)t_2}} = \hat{\Omega}(\mu) \left[\frac{e^{-(\delta_A - \delta_P)\Phi(\mu_1)/\delta_A}}{\ell \delta_A v(1)} - \int_\mu^{\mu_1} \frac{1}{\hat{\Omega}(x)} \frac{e^{-(\delta_A - \delta_P)\Phi(x)/\delta_A}}{\ell \delta_A (1-x)v(x)} dx \right]$$

$$\log \hat{\Omega}(\mu) = \int_\mu^{\mu_1} \frac{v(1)}{(1-x)v(x)} dx.$$

A similar argument as before establishes that $\frac{L(0)}{e^{(\delta_A - \delta_P)t_2}}$ is increasing in μ_1 and thus for phase 1 to be of length 0, we need

$$\mu_1 = \mu_A : \frac{L(0)}{e^{(\delta_A - \delta_P)t_2}} \leq \frac{1}{\delta_A v(1)} \Rightarrow$$

$$e^{\int_0^{\mu_A} \frac{v(1)}{(1-x)v(x)} dx} \left[\frac{e^{-(\delta_A - \delta_P)\Phi(\mu_A)/\delta_A}}{\ell \delta_A v(1)} - \int_0^{\mu_A} \frac{1}{\hat{\Omega}(x)} \frac{e^{-(\delta_A - \delta_P)\Phi(x)/\delta_A}}{\ell \delta_A (1-x)v(x)} dx \right] \leq \frac{1}{\delta_A v(1)}$$

A similar argument as before shows that the LHS of the above is decreasing in μ_A and thus holds

for high enough values of μ_A and thus establishing the result. \square

A.6 Proof of Proposition 4

Proof. We start by characterization of the steady-state level, i.e., $\mu_P = \mu_P^*(\mu_A)$ where $\mu_P^*(\cdot)$ satisfies (6). Given its definition in (6), when $\mu_A < 1/2$, then, $\mu_A > \mu_P^*(\mu_A)$ or $\ell > 1$ while when $\mu_A > 1/2$, then $\mu_A < \mu_P^*(\mu_A)$ or $\ell < 1$. Moreover, since $v(\mu_A) = v(1 - \mu_A)$ and $v'(\mu_A) = -v'(1 - \mu_A)$, we have

$$\mu_P^*(\mu_A) + \mu_P^*(1 - \mu_A) = 1.$$

Moreover, when $\mu_P^*(\mu_A) > 0$ and $\mu_A \leq 1/2$, we must have

$$\begin{aligned} \frac{d}{d\mu_A} \mu_P^*(\mu_A) &= \frac{d}{d\mu_A} \mu_A \frac{v(\mu_A) + (1 - \mu_A)v'(\mu_A) + \frac{\delta_P}{\delta_A - \delta_P}v(1)}{v(\mu_A) + \frac{\delta_P}{\delta_A - \delta_P}v(1)} \\ &= \frac{\mu_P^*(\mu_A)}{\mu_A} + \frac{(1 - \mu_A)v''(\mu_A)}{v(\mu_A) + \frac{\delta_P}{\delta_A - \delta_P}v(1)} - \frac{\mu_P^*(\mu_A)v'(\mu_A)}{v(\mu_A) + \frac{\delta_P}{\delta_A - \delta_P}v(1)} \end{aligned}$$

Since $v''(\mu_A) \geq 0$ and $v'(\mu_A) < 0$, for $\mu_A < 1/2$, the above is positive. Moreover, if $\mu_A > 1/2$ and $\mu_P^*(\mu_A) < 1$, then

$$\frac{d}{d\mu_A} \mu_P^*(\mu_A) = -\frac{d}{d\mu_A} \mu_P^*(1 - \mu_A) = \frac{d}{d(1 - \mu_A)} \mu_P^*(1 - \mu_A) > 0$$

Therefore, when $\mu_P^*(\mu_A) \in (0, 1)$, μ_P^* is strictly increasing in μ_A .

Let $\lambda^*(\mu_A)$ be defined by

$$v(1) \frac{\lambda^*(\mu_A)}{\lambda^*(\mu_A) + \delta_A} = v(\mu_A).$$

This is the incentive constraint of the agent when

$$\frac{G_{P,1}(t)}{\mu_P} = \frac{G_{P,0}(t)}{1 - \mu_P} = e^{-\lambda^*(\mu_A)t}$$

In line with the rest of the analysis, we assume that $\mu_P = \mu_P^*(\mu_A) \in (0, 1)$. The Lagrange multiplier $\Lambda(t)$ is defined by

$$\Lambda(t) = \frac{1 - \ell}{\ell} \frac{e^{(\delta_A - \delta_P)t}}{v'(\mu_A)(\delta_A - \delta_P)}$$

Given the definition of $\mu_P^*(\cdot)$ and the property mentioned above, $1 - \ell$ and $v'(\mu_A)$ have the same sign, which implies that $\Lambda(t) > 0$, and $\Lambda'(t) > 0$. Given the above definition, $\Lambda(t)$ satisfies

$$\begin{aligned} & (v(1) - v(\mu_A) + \mu_A v'(\mu_A)) \Lambda'(t) - \delta_A v(1) \Lambda(t) = \\ & [(\delta_A - \delta_P)(\mu_A v'(\mu_A) - v(\mu_A)) - \delta_P v(1)] \frac{1 - \ell}{\ell} \frac{e^{(\delta_A - \delta_P)t}}{v'(\mu_A)(\delta_A - \delta_P)} = \\ & \left[\mu_A v'(\mu_A) - v(\mu_A) - \frac{\delta_P}{\delta_A - \delta_P} v(1) \right] \frac{\mu_P^* - \mu_A}{(1 - \mu_P^*) \mu_A} \frac{e^{(\delta_A - \delta_P)t}}{v'(\mu_A)} = \\ & \quad - \left[v(\mu_A) + \frac{\delta_P}{\delta_A - \delta_P} v(1) - \mu_A v'(\mu_A) \right] \times \\ & \quad \frac{\mu_A (1 - \mu_A) v'(\mu_A)}{\left(v(\mu_A) + \frac{\delta_P}{\delta_A - \delta_P} v(1) - \mu_A v'(\mu_A) \right) \mu_A (1 - \mu_A)} \frac{e^{(\delta_A - \delta_P)t}}{v'(\mu_A)} = -e^{(\delta_A - \delta_P)t} \end{aligned}$$

where in the above we have used the definition of μ_P^* from (6). Similarly, we can show that

$$\ell (v(1) - v(\mu_A) - (1 - \mu_A) v'(\mu_A)) \Lambda'(t) - \ell \delta_A v(1) \Lambda(t) + e^{(\delta_A - \delta_P)t} = 0$$

A simple examination of the above illustrates that they are the First Order Condition of the Lagrangian $L(\hat{x}; x)$ with respect to $G_{P,0}(t)$ and $G_{P,1}(t)$. This establishes the optimality of the steady state mentioned above.

In order to prove the claim of the Proposition, we will show the case where $\mu_P > \mu_P^*(\mu_A)$. The other case is the mirror of this proof. Suppose that $\mu_A(t)$ satisfies the differential equation (7) in phase 1 until it reaches the value μ_A^* which is defined by

$$\ell = \frac{\mu_A^*}{1 - \mu_A^*} \frac{1 - \mu_P^*(\mu_A^*)}{\mu_P^*(\mu_A^*)}.$$

Such a value of μ_A^* is uniquely determined since the above becomes

$$\ell = \frac{\delta_A v(1) - (\delta_A - \delta_P) [v(1) - v(\mu_A^*) + \mu_A^* v'(\mu_A^*)]}{\delta_A v(1) - (\delta_A - \delta_P) [v(1) - v(\mu_A^*) - (1 - \mu_A^*) v'(\mu_A^*)]} \quad (23)$$

whenever $\mu_A^* \in (0, 1)$. Since the numerator is decreasing in μ_A^* and the denominator is increasing, the above have at most one solution. When no such μ_A^* exist, we set $\mu_A^* = 1$. We let t^* be the time by which this value of belief is reached. Additionally, let $G_{P,\omega}$ be the associated engagement

policy given by

$$G_{P,1}(t) = \begin{cases} \mu_P & t \leq t^* \\ \mu_P e^{-\hat{\lambda}(t-t^*)} & t > t^* \end{cases}$$

$$G_{P,0}(t) = \begin{cases} \ell \mu_P \frac{1-\mu_A(t)}{\mu_A(t)} & t \leq t^* \\ \ell \mu_P \frac{1-\mu_A(t^*)}{\mu_A(t^*)} e^{-\hat{\lambda}(t-t^*)} & t > t^* \end{cases}$$

Given this, the rest of the proof is very similar to that of Proposition 2. We construct the Lagrange multiplier using the ODE

$$e^{(\delta_A - \delta_P)t} + \Lambda'(t) \left(\underbrace{v(1) - v(\mu_A(t)) + \mu_A(t) v'(\mu_A(t))}_{\Delta_0(t)} \right) - \delta_A v(1) \Lambda(t) = 0, \forall t \leq t^*$$

for a given value of $\Lambda(0+) \geq 0$. For values of $t \geq t^*$, we set

$$\Lambda(t) = \frac{1-\ell}{\ell} \frac{e^{(\delta_A - \delta_P)t}}{v'(\mu_A^*)(\delta_A - \delta_P)}$$

where $\Lambda(0+)$ has to satisfy

$$\frac{1-\ell}{\ell} \frac{e^{(\delta_A - \delta_P)t^*}}{v'(\mu_A^*)(\delta_A - \delta_P)} = e^{\int_0^{t^*} \frac{\delta_A v(1)}{\Delta_0(t)} dt} \left[\Lambda(0+) - \int_0^{t^*} \frac{e^{-\int_0^t \frac{\delta_A v(1)}{\Delta_0(s)} ds} e^{(\delta_A - \delta_P)t}}{\Delta_0(t)} dt \right].$$

Since LHS is positive, $\Lambda(0+)$ that solves the above is also positive. Finally, the Lagrange multipliers λ_0, λ_1 are constructed according to:

$$\lambda_0 = \Lambda(0+) \Delta_0(0) > 0$$

$$\lambda_1 = \Lambda(0+) \Delta_1(0)$$

$$+ \int_0^{t^*} [e^{-\delta_P t} + \ell e^{-\delta_A t} (\Delta_1(t) \Lambda'(t) - \delta_A v(1) \Lambda(t))] dt$$

where $\Delta_1(t) = v(1) - v(\mu_A(t)) - (1 - \mu_A(t)) v'(\mu_A(t))$. The FOC with respect to $G_{P,\omega}(t)$ is satisfied when $\omega = 0$ or when $\omega = 1, t \geq t^*$, given that $\Lambda(t)$ satisfies the above ODE. To show optimality, it is sufficient to show that

$$\int_t^{t^*} e^{-\delta_P s} [e^{(\delta_A - \delta_P)s} + \ell (\Delta_1(s) \Lambda'(s) - \delta_A v(1) \Lambda(s))] \geq 0$$

We prove the stronger result that for all $t \in (0, t^*)$

$$e^{(\delta_A - \delta_P)t} + \ell (\Delta_1(t) \Lambda'(t) - \delta_{AV}(1) \Lambda(t)) \geq 0 \quad (24)$$

To establish this inequality, we realize that $\mu_P > \mu_P^*(\mu_A)$, and since t^* is the lowest time for which $\mu_P(t^*) = \mu_P^*(\mu_A(t^*))$, it must be that for all values of $t \in [0, t^*)$, $\mu_P(t) > \mu_P^*(\mu_A(t))$.

Recall that $\Lambda(t)$ satisfies the following

$$\Lambda(t) = \Gamma(t) \left[\Lambda(0+) - \int_0^t \frac{e^{(\delta_A - \delta_P)s}}{\Delta_0(s) \Gamma(s)} ds \right], \log \Gamma(s) = \int_0^s \frac{\delta_{AV}(1)}{\Delta_0(s)} ds$$

Let $\Omega(t) = \Lambda(t) / \Gamma(t)$. We can use the ODE governing $\Lambda(t)$ and rewrite the LHS of (24) as

$$\begin{aligned} & (\ell \Delta_1(t) - \Delta_0(t)) \Lambda'(t) + (1 - \ell) \delta_{AV}(1) \Lambda(t) = \\ & (\ell \Delta_1(t) - \Delta_0(t)) \left(\frac{\delta_{AV}(1)}{\Delta_0(t)} \Omega(t) \Gamma(t) - \frac{e^{(\delta_A - \delta_P)t}}{\Delta_0(t)} \right) + (1 - \ell) \delta_{AV}(1) \Omega(t) \Gamma(t) \end{aligned}$$

We can divide the above by $\Gamma(t)$ and rewrite the resulting expression as

$$\frac{\ell \Delta_1(t)}{\Delta_0(t)} \left[\delta_{AV}(1) \Omega(t) - \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} \right] + \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} - \delta_{AV}(1) \ell \Omega(t)$$

In order to show that the above is positive, we will show that it is decreasing and its value at t^* is positive. Its derivative with respect to t is given by

$$\begin{aligned} & \left(\frac{\ell \Delta_1(t)}{\Delta_0(t)} \right)' \left[\delta_{AV}(1) \Omega(t) - \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} \right] + \\ & \frac{\ell \Delta_1(t)}{\Delta_0(t)} \left[-\delta_{AV}(1) \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t) \Delta_0(t)} + \frac{\delta_{AV}(1) e^{(\delta_A - \delta_P)t}}{\Gamma(t) \Delta_0(t)} - \frac{(\delta_A - \delta_P) e^{(\delta_A - \delta_P)t}}{\Gamma(t)} \right] + \\ & \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} \left(\delta_A - \delta_P - \frac{\delta_{AV}(1)}{\Delta_0(t)} \right) + \delta_{AV}(1) \ell \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t) \Delta_0(t)} \\ & = \left(\frac{\ell \Delta_1(t)}{\Delta_0(t)} \right)' \left[\delta_{AV}(1) \Omega(t) - \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} \right] + \\ & \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t) \Delta_0(t)} ((\delta_A - \delta_P) (\Delta_0(t) - \ell \Delta_1(t)) + (\ell - 1) \delta_{AV}(1)) \end{aligned} \quad (25)$$

We show that each of the expressions above are negative. To do so, note that for all values of

$t < t^*$, we must have that $\mu_P^*(\mu_A(t)) < \mu_P(t)$, given the definition of t^* . Therefore

$$\begin{aligned} \ell &= \frac{1 - \mu_P(t)}{\mu_P(t)} \frac{\mu_A(t)}{1 - \mu_A(t)} < \frac{1 - \frac{\mu_A(t)v'(\mu_A(t))}{v(\mu_A(t)) + \frac{\delta_P}{\delta_A - \delta_P}v(1)}}{1 + \frac{(1 - \mu_A(t))v'(\mu_A(t))}{v(\mu_A(t)) + \frac{\delta_P}{\delta_A - \delta_P}v(1)}} = \frac{\delta_A v(1) - (\delta_A - \delta_P) \Delta_0(t)}{\delta_A v(1) - (\delta_A - \delta_P) \Delta_1(t)} \\ &\Rightarrow (\delta_A - \delta_P) (\Delta_0(t) - \ell \Delta_1(t)) + (\ell - 1) \delta_A v(1) < 0 \end{aligned}$$

which is proportional to the last expression in (25). Moreover, since $\mu_A(t)$ and $\Delta_0(t)$ are increasing, and $\Delta_1(t)$ is decreasing, we must have that $\left(\frac{\ell \Delta_1(t)}{\Delta_0(t)}\right)' \leq 0$. Finally,

$$\begin{aligned} \delta_A v(1) \Omega(t) - \frac{e^{(\delta_A - \delta_P)t}}{\Gamma(t)} &= \frac{\delta_A v(1) \Lambda(t) - e^{(\delta_A - \delta_P)t}}{\Gamma(t)} \\ &= \frac{\Delta_0(t) \Lambda'(t)}{\Gamma(t)} \geq 0 \end{aligned}$$

This implies that (25) is negative. Thus, it is sufficient to show that (24) holds at $t = t^*$ and this automatically holds given that at $t = t^*$, $\Lambda'(t^*) = (\delta_A - \delta_P) \Delta(t^*)$ and $\Lambda(t^*) = e^{(\delta_A - \delta_P)t^*} \frac{1 - \ell}{\ell v'(\mu_A^*)(\delta_A - \delta_P)}$ which implies that at $t = t^*$, (24) holds with equality. This concludes the proof. \square

A.7 Proof of Lemma 5

Proof. First, we note that we can consider $\hat{T} < \bar{T}$ and its associated solution

$$\hat{x}_{\hat{T}} = \left(G_{P,0} \left(T; \hat{T} \right), G_{P,1} \left(T; \hat{T} \right) \right)$$

. Let T_ω to be the infimum value such that $G_{P,\omega} \left(T_\omega; \hat{T} \right) = 0$ and if such value does not exist, we set $T_\omega = \hat{T}$. We note that we must have $T_1 = T_0$. Suppose that $T_1 < \hat{T}$, then we must have that $\mu_A(T) = 0$ for all $T > T_1$. Hence, the incentive constraint is given by

$$D(T) G_{P,0}(T) v(1) + \int_T^{\hat{T}} D'(s) G_{P,0}(s) ds \geq D(T) G_{P,0}(T) v(1)$$

Since $G_{P,0}(T) \geq 0$ and $D'(T) < 0$, the only way the above constraint can hold is if $G_{P,0}(T) = 0$. This implies that $T_0 \leq T_1$. Using a similar argument, we have that $T_1 = T_0$. Note that since the marginal benefit of $G_{P,\omega}(T)$ is always 1, we can assume that that $d\Lambda_{\hat{T}}(T) = 0$ for all $T \geq T_1$.

Next, let us consider $x = \left(0, \max \left\{ \hat{G}_{P,1} \left(T; \hat{T} \right) - \varepsilon, 0 \right\} \right) \in \Omega$ where $\varepsilon > 0$. We can then write the relations in (14) as

$$\frac{\partial L_{\hat{T}}(\hat{x}_{\hat{T}}; x) - \partial L_{\hat{T}}(\hat{x}_{\hat{T}}; \hat{x}_{\hat{T}})}{\varepsilon} \leq 0.$$

By sending $\varepsilon \rightarrow 0$, the above becomes

$$\begin{aligned} & T_1+ \tag{26} \\ & \ell \int_0^{T_1} \left[v(1) - v \left(\mu_A \left(T; \hat{T} \right) \right) - v' \left(\mu_A \left(T; \hat{T} \right) \right) \left(1 - \mu_A^T \left(T; \hat{T} \right) \right) \right] D(T) d\Lambda_{\hat{T}}(T) + \\ & \ell v(1) \int_0^{T_1} \int_T^{T_1} D'(s) ds d\Lambda_{\hat{T}}(T) - \lambda_{1,\hat{T}} \leq 0 \end{aligned}$$

Similarly, we can consider $x = \left(0, \max \left\{ \hat{G}_{P,1} \left(T; \hat{T} \right) - \varepsilon, 0 \right\} \right) \in \Omega \in \Omega^T$ and do the same as above. This would then imply that the LHS of (26) is non-negative and thus has to be 0. If we let $\Delta_1(\mu) = v(1) - v(\mu) - v'(\mu)(1 - \mu)$, then we can write the above as

$$\int_0^{T_1} \left(-\Delta_1 \left(\mu_A \left(T; \hat{T} \right) \right) D(T) + v(1) (D(T) - D(T_1)) \right) d\Lambda_{\hat{T}} = \frac{T_1 - \lambda_{1,\hat{T}}}{\ell}. \tag{27}$$

Repeating the above for $\omega = 0$, we arrive at

$$\int_0^{T_1} \left(-\Delta_0 \left(\mu_A \left(T; \hat{T} \right) \right) D(T) + v(1) (D(T) - D(T_1)) \right) d\Lambda_{\hat{T}} = T_1 - \lambda_{0,\hat{T}} \tag{28}$$

A property of the Δ_1, Δ_0 functions are

$$\begin{aligned}\frac{\Delta_1(\mu) + \Delta_0(\mu)}{2} &= v(1) - v(\mu) + \frac{v'(\mu)\mu - v'(\mu)(1-\mu)}{2} \\ &= v(1) - v(\mu) + \frac{v'(\mu)(2\mu-1)}{2}\end{aligned}$$

Convexity of $v(\mu)$ implies that

$$\begin{aligned}\mu \geq 1/2 : \frac{v(\mu) - v(1/2)}{\mu - 1/2} &\leq v'(\mu) \rightarrow v(\mu) - v'(\mu) \frac{2\mu-1}{2} \leq v(1/2) \\ \mu \leq 1/2 : \frac{v(1/2) - v(\mu)}{1/2 - \mu} &\geq v'(\mu) \rightarrow v(\mu) - v'(\mu) \frac{2\mu-1}{2} \leq v(1/2)\end{aligned}$$

Which further implies that

$$\frac{\Delta_1(\mu) + \Delta_0(\mu)}{2} \geq v(1) - v(1/2)$$

Now, we consider (27) and (28) and take equally weighted average and use the above inequality to arrive at

$$\int_0^{T_1} (-(v(1) - v(1/2)) D(T) + v(1) (D(T) - D(T_1))) d\Lambda_{\hat{T}} \geq \frac{T_1 - \lambda_{1,\hat{T}} + \ell (T_1 - \lambda_{0,\hat{T}})}{2\ell}$$

Using integration by parts, we can write

$$\begin{aligned}- \int_0^{T_1} (v(1) - v(1/2)) D(T) d\Lambda_{\hat{T}} - v(1) \int_0^{T_1} \Lambda_{\hat{T}}(T) D'(T) dt &\geq \frac{T_1 - \lambda_{1,\hat{T}} + \ell (T_1 - \lambda_{0,\hat{T}})}{2\ell} \Rightarrow \\ \int_0^{T_1} (v(1) - v(1/2)) D(T) d\Lambda_{\hat{T}} + v(1) \int_0^{T_1} \Lambda_{\hat{T}}(T) D'(T) dt &\leq - \frac{T_1 - \lambda_{1,\hat{T}} + \ell (T_1 - \lambda_{0,\hat{T}})}{2\ell} \Rightarrow \\ \int_0^{T_1} D(T) d\Lambda_{\hat{T}} + \frac{v(1)}{v(1) - v(1/2)} \int_0^{T_1} \Lambda_{\hat{T}}(T) D'(T) dt &\leq - \frac{T_1 - \lambda_{1,\hat{T}} + \ell (T_1 - \lambda_{0,\hat{T}})}{2\ell (v(1) - v(1/2))} \Rightarrow \\ \int_0^{T_1} D(T)^{\frac{-v(1/2)}{v(1)-v(1/2)}} d \left(D(T)^{\frac{v(1)}{v(1)-v(1/2)}} \Lambda_{\hat{T}}(T) \right) &\leq - \frac{T_1 - \lambda_{1,\hat{T}} + \ell (T_1 - \lambda_{0,\hat{T}})}{2\ell (v(1) - v(1/2))} \Rightarrow \\ \int_0^{T_1} d \left(D(T)^{\frac{v(1)}{v(1)-v(1/2)}} \Lambda_{\hat{T}}(T) \right) &\leq \int_0^{T_1} D(T)^{\frac{-v(1/2)}{v(1)-v(1/2)}} d \left(D(T)^{\frac{v(1)}{v(1)-v(1/2)}} \Lambda_{\hat{T}}(T) \right)\end{aligned}$$

where in the above we have used the fact that $D(T)^{\frac{-v(1/2)}{v(1)-v(1/2)}} \geq 1$. We then have that

$$\begin{aligned} D(T)^{\frac{v(1)}{v(1)-v(1/2)}} \Lambda_{\hat{T}}(T) - \Lambda_{\hat{T}}(0) &\leq -\frac{T_1 - \lambda_{1,\hat{T}} + \ell(T_1 - \lambda_{0,\hat{T}})}{2\ell(v(1) - v(1/2))} \\ 0 \leq D(T)^{\frac{v(1)}{v(1)-v(1/2)}} \Lambda_{\hat{T}}(T) &\leq \Lambda_{\hat{T}}(0+) - \frac{T_1 - \lambda_{1,\hat{T}} + \ell(T_1 - \lambda_{0,\hat{T}})}{2\ell(v(1) - v(1/2))} \end{aligned}$$

Hence, in order to prove the claim, it is sufficient to show that the right hand side stays bounded as \hat{T} converges to \bar{T} .

We know that $T_1 \leq \bar{T}$ and thus it is bounded above. Moreover, since $\Delta_1(\mu)(1-\mu) + \Delta_0(\mu)\mu = v(1) - v(\mu)$, we can use complementary slackness to write $\partial L_{\hat{T}}(\hat{x}_{\hat{T}}; \hat{x}_{\hat{T}}) = 0$ as

$$\int_0^{T_1} \hat{G}_P(T; \hat{T}) dT = \lambda_{1,\hat{T}} \mu_P + \lambda_{0,\hat{T}} (1 - \mu_P)$$

The RHS of the above is less than $T_1 \leq \bar{T}$ which then implies that $\lambda_{1,\hat{T}}$ and $\lambda_{0,\hat{T}}$ stay bounded below $\max\left\{\frac{\bar{T}}{\mu_P}, \frac{\bar{T}}{1-\mu_P}\right\}$. Finally, plugging in $x = \left(\hat{G}_{P,0}(T), \varepsilon \mathbf{1}[T=0] + \hat{G}_{P,1}(T)\right)$ into (14) and subtract (15) we arrive at

$$\Delta_1(\mu_A) \Lambda_{\hat{T}}(0+) - \lambda_{1,\hat{T}} \leq 0$$

which implies that $\Lambda_{\hat{T}}(0+)$ is also bounded. This concludes the proof. \square