# The Optimal Level of Consultation

Jiayin Zhai*

November 2025

Click here for the latest version

**Abstract**

This paper studies how an information seller optimally manages information precision to maximize revenue from a consumer who pays to learn before taking an irreversible action. The key insight is that the seller provides the least informative signals when the consumer's belief is close to indifference, thereby prolonging engagement and increasing rents. This behavior generates a U-shaped relationship between optimal precision and the consumer's prior belief, and this result is robust to alternative cost structures. When consumers hold heterogeneous priors, standard screening becomes infeasible, inducing the seller to offer a single pooling contract. The paper also develops an empirical framework to test whether sellers adjust information quality to extract rents from consumers with heterogeneous priors. The results show that information distortions may arise from the seller's strategic need to manage the consumer's attention, introducing a new mechanism to the literature on dynamic information provision and advisory markets.

*University of California, Los Angeles (UCLA), Email: jiayinz06@ucla.edu.
I am deeply indebted to Moritz Meyer-ter-Vehn for his invaluable guidance and support throughout this project. I am also very grateful to Jay Lu, Simon Board, Alex Bloedel, Tomasz Sadzik, Ichiro Obara, Daniel Clark for their advice and helpful suggestions. I would also like to thank Jay Yoo, Minghao Pan, Jingyu Fan and participants at the UCLA Theory Proseminar for their helpful comments. All errors are my own.

# 1 Introduction

Advisors, consultants, and algorithmic recommendation systems often sell information to decision makers who must act under uncertainty. While these information sellers may sometimes generate biased or incomplete information because of misaligned preferences or lack of commitment, in many environments they are largely indifferent to which action the consumer ultimately takes. Instead, their primary objective is to secure and sustain the consumer's attention: a digital search platform does not care if you accept a surgery in the end, but it wants you to browse longer so their advertisement slots are more valuable on the other side; a consultant may also have the same incentives to prolong the consultation session to increase billable hours. In this scenario, revenue depends not only on the accuracy of the information they deliver but also on the duration of engagement, and the information sellers face a trade-off: if the legal advisor who reveals too much too quickly risks ending the consultation prematurely, while revealing too little may cause the client to walk away. A digital platform that offers instant, highly precise recommendations may lose users' attention, but yet flooding the search results with ads may drive them somewhere else. These environments reveal a fundamental tension between controlling information quality and sustaining engagement.

This paper studies how an information seller optimally manages the quality of information in such environments to maximize a consumer's engagement and thus revenue. I develop a continuous-time principal-agent model in which the agent must choose between a safe and a risky action whose payoffs depend on an unknown binary state. Before acting, the two players share the same prior belief on the state, and the agent may purchase information from the principal at a constant flow cost. Information arrives through a drift-diffusion process whose precision is chosen by the principal at time 0. We consider the setting that there is no discount and that generating information is costless to the principal. We investigate how the optimal information precision depends on the agent's prior belief.

My main result is that in the optimal design, the principal provides the least informative signals when the agent is most uncertain, precisely when his prior is closest to indifference between the two actions. This is a surprising result because in an single agent setting, the intuition suggests that the experimentation level should peak when he has moderate belief, because he is most uncertain about which action to take, and the posterior belief variance is greatest for moderate priors, which means he learns very fast in the beginning. This may suggest that the agent with central, unfocused beliefs has the highest underlining demand for information. However, when the signal precision is chosen by a seller who wants to maximize consumer engagement, the agent has the most central prior belief ends up receiving the worst information. A similar result was presented in Moscarini and Smith (2001) where the agent acquiring costly information while controlling the intensity himself. With convex flow payoff that increases in experimentation intensity and discounting, dynamic considerations drive the agent to delay high-intensity experimentation to the end when the beliefs are close to the stopping beliefs. Therefore, high experimentation appears at focused beliefs.

In the strategic setting, the uncertain agent gets the worst information precisely because he has the highest underlying information demand and the principal can take advantage of it. When the agent is very unsure, information is extremely valuable to him, so the principal exploits this feature by slowing down the signal to stretch out engagement. When an agent's prior is more certain, information is less valuable, so to keep them paying attention at all, the principal must offer more precise signals. This generates a U-shaped relationship between optimal precision and the agent's prior.

The paper also develops two extensions that clarify the robustness and institutional implications of the main mechanism. First, I show that the qualitative predictions of the model are robust to alternative flow cost structures, including convex cost. When the agent pays a flow cost that depends on the precision level, the agent's stopping behavior depends only on the ratio of precision to cost, and the principal's revenue can again be expressed purely as a function of the stopping posterior. As a result, the optimal stopping posteriors and the U-shaped relationship between optimal precision and prior beliefs coincide with those in the benchmark model with constant flow cost. This robustness result implies that what matters for behavior is the relative tightness of information per unit of cost, not the particular functional form of the fee schedule.

Second, I consider a screening problem where the agent's prior belief is private information and the principal can offer a menu of contracts. Because the agent's value is strictly increasing in the normalized precision, all types of agents rank contracts in the same way. This destroys the usual scope for screening: any menu of contracts collapses to the single contract with the optimal normalized precision. The optimal mechanism is therefore a pooling contract chosen to maximize expected revenue against the distribution of priors, effectively targeting the agents with intermediate priors between two prior cutoffs. Those agents whose priors fall out of the region will opt out. This extension shows that standard adverse-selection tools cannot be used to segment agents by prior beliefs in this class of dynamic information-selling problems. It also brings out a practical implication: when an information seller is not able to use her knowledge about the consumers' priors, either due to lack of access or legal restrictions, she cannot differentiate across types through menus. Instead, she can offer a single, one-size-fits-all product: a legal advisor will deliver the same level of service to all clients, and a digital platform will mix the same density of advertising across all users' search results.

A natural question that follows is whether information providers extract rents by using their knowledge of consumers' private beliefs. The model delivers a central empirical implication: information quality follows a U-shaped relationship with consumers' priors. Sellers provide the least precise signals to clients whose beliefs lie near indifference and offer sharper information to those with more extreme priors. This pattern arises directly from the engagement motive and lends itself to empirical testing. It can be detected using observable proxies for precision—such as the seniority of the consultant or analyst assigned to a client, the specificity or depth of personalized recommendations, or the density of commercial content mixed into a platform's search results.

The paper bridges dynamic information provision, gaussian learning, and the economics of attention. It shows that information distortions arise not just from bias or misalignment but also from the seller's strategic need to prolong engagement, adding a new mechanism to the literature to that typically focuses on bias, persuasion, or commitment. It also sheds light on real-world advice markets where revenue is tied to time or attention, such as consulting or subscription-based algorithmic services.

The paper proceeds as follows. Section 2 discusses relevant literature. Section 3 introduces the model. Section 4 characterizes the agent's optimal stopping rule. Section 5 solves the principal's revenue-maximization problem and derives the U-shaped optimal precision. Section 7 discusses empirical implications. Section 6 studies robustness and screening. Section 8 concludes.

## 2 Related Literature

This paper relates to several strands of literature. First, this paper contributes to the literature on dynamic information provision and costly experimentation. It connects the classical Wald sequential sampling problem with principal agent models of persuasion and attention (e.g.Orlov et al., 2020; Che et al., 2023), though in Che et al. (2023), the sender seeks to maintain the receiver's attention because persuasion requires time due to flow-cost constraints, whereas in the current paper attention retention arises endogenously from the principal's control of information precision. They derive the belief dynamics under the optimal information structure and show that the sender must leave rents to the receiver to compensate for his listening costs. They also demonstrate that any outcome between Kamenica and Gentzkow (2011)'s sender-optimal persuasion and full revelation is attainable in a Markov Perfect Equilibrium as persuasion costs vanish. In contrast, in my model, persuasion occurs through continuous-time diffusion signals, and the sender obtains only a fraction of the static persuasion payoff when persuasion takes place. Henry and Ottaviani (2019) also study a dynamic persuasion problem with state-dependent Gaussian drift processes. Similar to my motivation, they reinterpret the Wald problem as a strategic communication problem between a sender and a receiver, but they fix the precision of the information flow. They compare belief thresholds and welfare across three institutional arrangements that capture the evolution of the FDA's drug approval process: sender control, receiver control, and ex-ante commitment by the receiver. Eső and Szentes (2007) explore the optimal contracting problem of a consulting firm selling advice, modeled as noise drawn from a fixed distribution regarding a project's value, when the buyer's action is contractible. Lehrer and Wang (2024) investigate an optimal transfer scheme to sell a given Blackwell experiment for a fixed state that may be used for repeated sampling of conditionally independent signals and they find that the revenue-maximizing transfer scheme is an upfront transfer scheme.

This paper also contributes to the literature on markets for information. Eső and Szentes (2007) study a setting in which a monopolistic seller of information provides imperfect signals about the value of a project to a buyer whose ultimate action is contractible. Bergemann et al. (2018) analyze an environment in which a data intermediary sells targeted information to advertisers seeking to reach consumers with heterogeneous beliefs. A broader overview of markets for information can be found in the survey by Bergemann and

Bonatti (2019).

The agent's problem in this paper is fundamentally an optimal stopping problem, which has been extensively studied since Wald (1947). The classical Wald framework has been analyzed under drift–diffusion learning (e.g., Moscarini and Smith (2001); Ke et al. (2016); Fudenberg et al. (2015)), and generalized to environments with endogenous experimentation (Zhong (2022)). In particular, Gonçalves (2024) shows that decision makers facing a speed–accuracy tradeoff choose stopping times generated by a drift–diffusion optimal-stopping rule in laboratory settings, and the resulting decision structure mirrors the one faced by the agent in my model. This paper extends this tradition by embedding the stopping problem in a strategic setting, where the principal controls the experimentation intensity and the agent decides when to stop.

Finally, the paper is related to the literature on costly communication and attention management. Gentzkow and Kamenica (2014) extend Kamenica and Gentzkow (2011)'s concavification approach to costly signals with entropy-reduction costs, preserving tractability through posterior-based optimization. In my model, although the information process is continuous and dynamic, the sender's problem can likewise be represented as a choice over posterior beliefs. Persson (2018) examine persuasion under "information overload," where complexity imposes cognitive costs on the receiver, while Bloedel and Segal (2018) introduce rational inattention as an endogenous attention constraint. Relatedly, Orlov et al. (2020) study a setting where waiting is a payoff-relevant action in a real-option environment, contrasting with my framework where waiting serves as a medium of persuasion rather than an end in itself.

# 3 Model

*Terminal Payoffs.* A principal ("she") sells costly information to an agent ("he") who must decide when to stop learning and take one of two actions $a \in \{A, B\}$. The unknown state is $\theta \in \{H, L\}$, with common prior $p_0 = \Pr(\theta = H)$. Each action yields state-dependent payoffs:

$$u(A,H) = -l < 0, \qquad u(A,L) = l > 0,$$
$$u(B,H) = h > 0, \qquad u(B,L) = -h < 0,$$

The agent's expected payoffs at belief $p$ are:

$$u_A(p) = l(1-2p), \qquad u_B(p) = h(2p-1).$$

The agent is indifferent between $A$ and $B$ at $p = \frac{1}{2}$, where both expected payoffs are zero. The principal obtains zero terminal payoff regardless of the action taken.

*Timing.* Time flows continuously from $t = 0$. There's no discount factor for either player. At time 0, the principal commits to an information structure. At each time instance, the agent can decide whether to observe a piece of information to update his belief, then he can choose to take action $A$ or$B$, or to wait to listen to the principal's next piece of information. There's no cost if the agent chooses not to listen to the principal. Once an action is taken, the game ends.

*Information.* Before choosing an action, the agent can purchase informative signals of the state $\theta$. The signal process $\{x_t\}_{t\geq 0}$ is a Brownian motion with constant uncertain drift chosen by Nature, $\mu^\theta$ in state $\theta$, where $\mu^H = 1 = -\mu^L$. The principal chooses the flow variance $1/n$ at the beginning of the game, so the signal process solve the stochastic differential equation (SDE)

$$dx_t^\theta = \mu^\theta \, dt + \frac{1}{\sqrt{n}} \, dW_t, \tag{1}$$

in state $\theta$, where $W_t \sim N(0,t)$ is a Brownian motion, and $n > 0$ is the the consultation level or signal quality chosen by the principal at the beginning.

*Information Cost and Revenue.* While the agent listens to the principal, he pays flow cost $k > 0$ per unit time to the principal. We maintain the constant flow cost assumption to present the main results of the paper but we discuss the implication and other cost structures in later sections. If the agent stops at time $T$, the transfer is $kT$. Later sections discuss generalizations of this flow-cost assumption.

# 4 Agent's Problem: Optimal Stopping

*Belief Dynamics.* Given $n$, the signal process $\{x_t\}_{t\geq 0}$ induces a diffusion belief process $\{p_t\}_{t\geq 0}$. As in Moscarini and Smith (2001), or originally by Theorem 9.1 of Liptser and Shiryaev (2013), the agent's posterior belief $p_t = \Pr_t(\theta = H)$ satisfies

$$dp_t = 2p_t(1-p_t)\sqrt{n}\, d\bar{W}_t, \tag{2}$$

where $\bar{W}_t$ is the observation-adapted Brownian innovation process.

*Optimal Stopping Problem.* The agent maximizes his expected return less information costs. Denote $U(p_0, n)$ as the supremum value with respect to the stopping time $T$ when the agent has prior $p_0$ and signal quality $n$. Given precision $n$ and prior $p_0$, the agent chooses a stopping time $T$ to maximize:

$$\begin{aligned} U(p_0; n) = \sup_T \mathbb{E}\left[u(p_T) - \int_0^T k\, dt \;\middle|\; p_0\right], \\ dp_t = 2p_t(1-p_t)\sqrt{n}\, d\bar{W}_t. \end{aligned} \tag{3}$$

where $u(p_T)$ denotes the expected payoff of the optimal action given posterior $p_T$. Note that the agent will not postpone consuming information because all the information is ex ante identical, so it is without loss of generality to assume that the agent acquires information from the principal until he decides to take action. The agent selects action $A$ for $p \leq \underline{p}$, $B$ for $p \geq \bar{p}$, and experiments for $p \in (\underline{p}, \bar{p})$. (3) becomes an optimal stopping problem for boundaries $\underline{p}, \bar{p}$.

Since $U(p_0; n)$ is the value function of the receiver's optimal learning, $U(p_0; n)$ is convex. We will always have $U(p; n) \geq u(p)$. And in particular, we have $U(p; n) > u(p)$ when $p \in (\underline{p}, \bar{p})$. Beliefs $\{p_t\}$ are a martingale. For any given experimentation region $(\underline{p}, \bar{p})$, the Hamilton-Jacobi-Bellman (HJB) equation for the optimal stopping problem is

$$U(p; n) = -kdt + \mathbb{E}\left[U(p + dp)\right] \tag{4}$$

yielding the ODE

$$U''(p; n) = \frac{k}{2np^2(1-p)^2}. \tag{5}$$

The general solution is

$$U(p;n) = \frac{k}{2n}L(p) + c_1 p + c_2, \tag{6}$$

$$L(p) = p\ln\frac{p}{1-p} + (1-p)\ln\frac{1-p}{p}. \tag{7}$$

Here $L(p)$ denotes the the expected log-likelihood-ratio (LLR) function, which is non-negative and convex on $(0,1)$ and symmetric around $1/2$. It is an extreme case of the KL divergence for the binary state, where the alternative distribution is the mirrored distribution.

The stopping posteriors and the two constants in the value function $\left\{\underline{p}, \bar{p}, c_1, c_2\right\}$ can be pinned down by two sets of smooth pasting and value matching conditions:

$$U(\underline{p};n) = u_A(\underline{p}) = l(1-2\underline{p}), \tag{8}$$

$$U'(\underline{p};n) = u'_A(\underline{p}) = -2l, \tag{9}$$

$$U(\bar{p};n) = u_B(\bar{p}) = h(2\bar{p}-1), \tag{10}$$

$$U'(\bar{p};n) = u'_B(\bar{p}) = 2h. \tag{11}$$

With the symmetry property of the LLR function $L(p)$ and the indfference belief $\hat{p} = 1/2$, we can show that the two stopping posteriors are symmetric around $1/2$.

**Lemma 1.** *If the agent is indifferent at $p = \frac{1}{2}$, then for any $h, l > 0$,*

$$\bar{p} = 1 - \underline{p}.$$

*Proof.* Consider a transformed value function $\tilde{U}(p;n)$ defined by subtracting the linear asymmetry between the payoffs:

$$\tilde{U}(p;n) = U(p;n) - (h-l)\left(p - \frac{1}{2}\right). \tag{12}$$

Since the subtracted term is linear in $p$, the second derivative remains unchanged: $\tilde{U}''(p) = U''(p)$. Thus, $\tilde{U}(p)$ represents the value function of a problem with the same information cost structure, which is symmetric around $p = 1/2$.

We apply the same transformation to the terminal payoff obstacles:

$$\tilde{u}_B(p) = u_B(p) - (h-l)\left(p - \frac{1}{2}\right) = (h+l)\left(p - \frac{1}{2}\right),$$
$$\tilde{u}_A(p) = u_A(p) - (h-l)\left(p - \frac{1}{2}\right) = -(h+l)\left(p - \frac{1}{2}\right).$$

The transformed problem is perfectly symmetric: the cost of learning depends on $p(1-p)$, and the transformed payoffs are symmetric reflections $\tilde{u}_B(p) = -\tilde{u}_A(1-p)$ with effective stakes $(h+l)$. Consequently, the optimal stopping set must be symmetric around the indifference point 1/2. It follows that $\bar{p} - 1/2 = 1/2 - \underline{p}$, which implies $\bar{p} = 1 - \underline{p}$. □

Using the symmetry result and subtracting derivative conditions, we get

$$U'(\bar{p}; n) - U'(\underline{p}; n) = u'_B(\bar{p}) - u'_A(\underline{p}) = 2h - (-2l) = 2(h+l).$$

Substituting $U'(p; n) = \frac{k}{2n} L'(p) + c_1$ and using $L'(\underline{p}) = -L'(\bar{p})$, we obtain:

$$\frac{k}{n} L'(\bar{p}) = 2(h+l).$$

**Proposition 1** (Characterization of Optimal Stopping Boundaries)**.** *For signal precision $n$, flow cost $k$, and payoff magnitudes $h, l > 0$, the stopping boundaries satisfy:*

$$L'(\bar{p}) = -L'(\underline{p}) = \frac{2n(h+l)}{k}, \tag{13}$$

*with $\bar{p} = 1 - \underline{p}$.*

**Proposition 2** (Comparative Statics)**.** *Let $(\underline{p}, \bar{p})$ be defined above. Then:*

- *Information cost: As $k$ decreases, the experimentation region widens.*
- *Information quality: As $n$ increases, the experimentation region widens.*
- *Stakes: As $h + l$ increases, the experimentation region widens.*

Basically, for a fixed $n$, we can fully characterize the stopping posteriors $\{\underline{p}, \bar{p}\}$, which does not depend on the prior $p_0$. Therefore, for a fixed signal precision $n$, if $p_0 \in (\underline{p}, \bar{p})$,

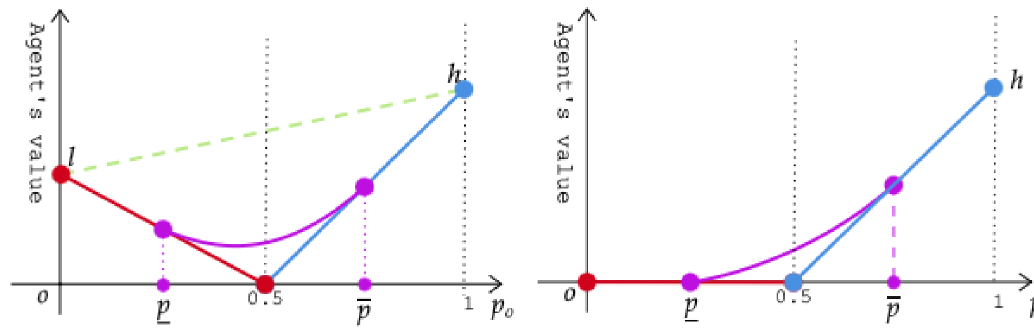


Figure 1: Agent Value Function under Given $n$

Left panel: both actions are risky; right panel: action $A$ is safe and always generates 0 payoff. Notice that the symmetry result holds in both scenarios.

then the agent experiments till belief exceeds the experimentation region; if $p_0 \notin (\underline{p}, \bar{p})$, they will not acquire costly signals and will act based on their prior belief at time 0.

Another noticeable point is that the optimal stopping beliefs only depends on the ratio of precision to flow cost $n/k$, or the normalized precision. This result suggests that the model may be able to accommodate more flexible flow cost structures.

# 5 Principal's Problem and Optimal Consultation Level

Since the principal knows the agent's prior $p_0$, anticipating the agent's stopping behavior, the principal chooses the information precision $n(p_0)$ to maximize expected revenue:

$$\begin{aligned} V(p_0) &= \underset{n}{max}\, \mathbb{E}\left[\int_0^T kdt \Big| p_0\right] = k\mathbb{E}[T|p_0] \\ s.t. \qquad T &= \underset{\tau}{argsup}\, \mathbb{E}\left[u(\tilde{q}_\tau)|p_0\right] - k\mathbb{E}[\tau|p_0] \end{aligned}$$

Morris and Strack(2019) proves the equivalence of cost of Wald's sequential sampling model and the log likelihood cost. In particular,

$$\begin{aligned}\mathbb{E}[T] &= \frac{1}{2n}\mathbb{E}_\pi\left[L(\tilde{q}) - L(p_0)\right] \qquad \text{where } \mathbb{E}_\pi\left[\tilde{q}\right] = p_0 \\ &= \frac{1}{2n}\left[L(\bar{p}(n)) - L(p_0)\right] = \frac{1}{2n}\left[L(\underline{p}(n)) - L(p_0)\right]\end{aligned}$$

This is because by choosing $n$, the agent's optimal stopping will make him stop at $\underline{p}$ or $\bar{p}$. Because of the symmetry result, we have $\bar{p} = 1 - \underline{p}$, and expected LLR function $L(p)$ is symmetric around 1/2, therefore the expected stopping time, i.e. the expected engagement length can be simplified to be a function of only the stopping beliefs. Besides, using the relationship between the stopping beliefs $\{\underline{p}, \bar{p}\}$ and information precision $n$, the principal's problem can be further simplified to

$$\underset{\bar{p}}{sup}\,(h+l)\frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})} \quad \text{or} \quad \underset{\underline{p}}{sup}\, -(h+l)\frac{L(\underline{p}) - L(p_0)}{L'(\underline{p})}$$

This representation shows that the principal's problem of choosing $n$ based on $p_0$ is equivalent to targeting the posterior $\bar{p}(p_0)$ or $\underline{p}(p_0)$ at which the agent stops, balancing informativeness against engagement length. We can trace how optimal $\bar{p}$ or $\underline{p}$ changes in $p_0$, and use the monotonicity of $\bar{p}$ or $\underline{p}$ and $n$, further establishing how the optimal consultation level $n$ changes in $p_0$:

**Theorem 1** (Optimal Consultation Level)**.** *For each prior $p_0 \in (0,1)$, there exists a unique optimal consultation level $n^*(p_0)$ that maximizes the principal's expected revenue. $n^*(p_0)$ is symmetric around $1/2$ it strictly decreases on $(0, 1/2)$, strictly increases on $(1/2, 1)$ and reaches the unique minimum when $p_0 = 1/2$.*

*Proof.* See Appendix A.2. □

We show that when the principal can choose the information quality for agents with different priors $p_0$, there exists a unique optimal information precision level $n^*(p_0)$ for all possible priors $p_0 \in (0,1)$, and the informativeness is minimized when $p_0 = 1/2$, i.e. for the agent with the most unfocused prior belief. The optimal consultation level $n^*(p_0)$ is also symmetric around 1/2.

This result also implies that the experimentation region for the agent is the smallest when $p_0 = 1/2$; as the prior belief moves away from 1/2, either more pessimistic or more negative, the principal will provide better signal quality, and therefore widens the learning region. Further, the optimal information design is such that the principal has a unique $n^*(p_0)$ for all $p_0 \in (0, 1)$, i.e. as long as the agent does not know the state for sure, no matter how confident he is, the principal will provide some information level to him and extracts positive rents.

The proof sketch is as follows. First we focus on the case when $p_0 \in [1/2, 1)$, and let the principal choose the optimal consultation level $n$ to target at the belief upper bound $\bar{p}$. We showed earlier that the principal's problem is essentially choosing $\bar{p}$ to maximize the expected engagement, and we want to know how the optimal $\bar{p}$ changes in $p_0$. We can use either implicit function theorem, or Topkis' theorem as in Milgrom and Shannon (1994) to show that $\bar{p}^*$ strictly increases in $p_0$ for $p_0 \in [1/2, 1)$. Then since $\bar{p}$ strictly increases in $n$, we can show that $n^*(p_0)$ strictly increases in $p_0$ for $p_0 \in [1/2, 1)$. When $p_0 \in (0, 1/2)$, we can focus on the optimization problem in terms of $\underline{p}$. Similarly, we can show that $\underline{p}^*$ strictly increases in $p_0$ for $p_0 \in (0, 1/2)$, and because $\underline{p}$ strictly decreases in $n$, it follows that $n^*(p_0)$ strictly decreases in $p_0$ for $p_0 \in (0, 1)$. Therefore, the optimal $n^*(p_0)$ reaches the lowest value when $p_0 = 1/2$.

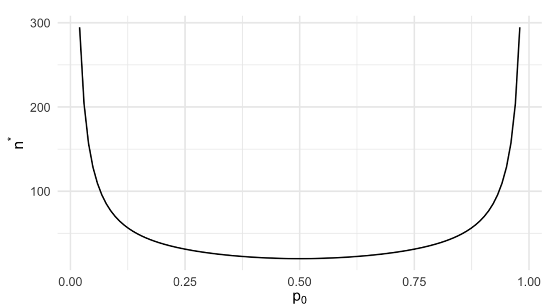


Figure 2: Optimal Consultation Level $n^*$

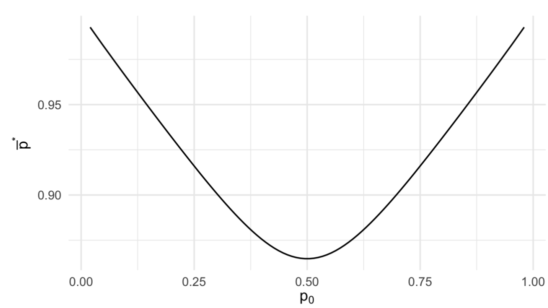


Figure 3: Optimal Stopping boundary $\bar{p}^*$

Here the principal earns revenue equal to the flow fee $k$ multiplied by the expected duration of learning. We know that the agent with more unfocused prior belief gets lower information precision, and the experimentation region is the smallest. Although for the same experimentation region, it takes longer for agents with more unfocused prior to reach the stopping boundaries; with the personalized information quality $n^*(p_0)$ and the resulting experimentation region $(\underline{p}(p_0), \bar{p}(p_0)$, it's unclear how the expected learning time, as well as the principal's expected revenue changes with $p_0$. Using the envelope theorem, we obtain the following result:

**Theorem 2** (Principal's Value Function)**.** *Under the optimal consultation level $n^*(p_0)$, the principal's value function $V(p_0)$ is symmetric around $1/2$ and strictly single-peaked, attaining its unique maximum at $p_0 = 1/2$.*

*Proof.* See Appendix A.4. □

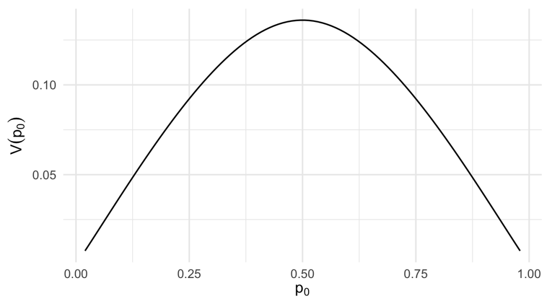


Figure 4: Principal's Value Function

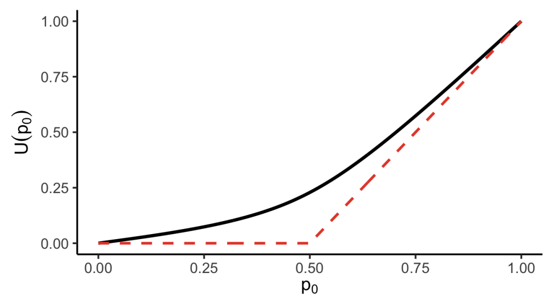


Figure 5: Agent's Value Function

We can see that the principal's value is also maximized when $p_0 = 1/2$. h the information precision is lowest and the experimentation region is smallest at this prior, the expected learning time remains longest for the agent with the most unfocused belief, yielding the highest payoff for the principal.

Notice that there is no real cost in the setting: the flow payment from the agent to the principal is essentially a transfer. If the agent could choose $n$ himself, he would simply select the highest precision and stop only upon learning the state with certainty. Therefore, all premature stopping in equilibrium arises entirely from the principal's strategic incentive to prolong the agent's learning process and extract more revenue.

# 6 Extensions

## 6.1 Non-constant Flow Cost

In this section, we allow the flow cost $k$ to increase with information precision $n$. This assumption is natural: higher-quality information typically commands a higher price: more experienced consultants charge higher fees, and reducing platform advertising often requires purchasing an upgraded service. In the baseline model, the agent pays a constant cost $k$ per unit of time while learning. The analysis, however, extends directly to any cost function $k(n)$ that is increasing in the precision level $n > 0$ chosen by the principal at time 0. The key observation is that the agent's continuation value on the experimentation region satisfies

$$U''(p) = \frac{k(n)}{2n} \frac{1}{p^2(1-p)^2},$$

so the stopping boundaries are again pinned down by

$$L'(\bar{p}) = -L'(\underline{p}) = \frac{2n(h+l)}{k(n)},$$

Thus, for any admissible cost function, the agent's stopping behavior depends only on the *ratio* $n/k(n)$.

Wald's representation of the expected stopping time further implies that the principal's expected revenue can be written purely as a function of the stopping posterior:

$$\mathrm{V}(p_0, n) = (h+l)\frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})}.$$

This is identical to the objective in the benchmark model with constant flow cost. Hence the optimal stopping posteriors and consequently the optimal precision profile coincide exactly with those derived earlier.

**Theorem 3** (Robustness to Non-constant Flow Cost)**.** *Suppose the principal chooses a precision level $n > 0$ at time 0 and the agent pays a flow cost $k(n) > 0$ that is increasing in $n$. Then for every prior $p_0 \in (0,1)$:*

1. *The agent's stopping boundaries depend on $(n, k(n))$ only through the ratio $n/k(n)$.*
2. *The principal's revenue maximization reduces to*

$$\sup_{\bar{p} > 1/2} (h+l)\frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})}.$$

3. *The optimal stopping posteriors $\bar{p}^*(p_0)$ and the corresponding optimal precision levels coincide with those in the constant flow-cost model.*

This result shows that the qualitative prediction that the U-shaped relationship between optimal precision and the prior is robust to arbitrary increasing flow-cost functions, provided the principal commits to a fixed precision level at time 0.

## 6.2 Screening with Heterogeneous Priors

We now consider a setting in which the agent's prior $p_0$ is private information. The principal observes only a known distribution $F$ over types $p_0 \in (0,1)$ and may offer a menu of contracts to screen the agent.

*Players, Types, and Information*:

- **Principal**: chooses a menu of consultation contracts.
- **Agent**: privately knows his prior $p_0$.
- **Type space**: $p_0 \in (0,1)$ with commonly known distribution $F$.

*Contract Space and Actions*: A contract is a pair $(n, k)$ specifying a precision level and a flow cost. The principal offers a menu

$$\mathcal{M} = \{(n(p), k(p))\}_{p \in (0,1)}.$$

A type-$p_0$ agent selects a contract from $\mathcal{M}$ and then chooses the optimal stopping time given the induced ratio $\rho = n/k(n)$.

*Payoffs*: Under precision $n$ with flow cost $k$, a type-$p_0$ agent obtains

$$U(p_0; n) = \sup_T \mathbb{E}\left[u(p_T) - \int_0^T k(n)\, dt\right].$$

The principal's profit from such a type is

$$\Pi(p_0; n) = k(n)\, \mathbb{E}[T(p_0, n)].$$

*Incentives*: Because the agent's continuation value is increasing in the ratio $\rho = n/k(n)$, all types agree on the ranking of contracts. Formally, for any two contracts $(n_i, k_i)$ and

$(n_j, k_j)$:

$$\frac{n_i}{k_i} > \frac{n_j}{k_j} \quad \Longrightarrow \quad U(p_0; n_i) \geq U(p_0; n_j) \qquad \forall\, p_0.$$

Hence no menu can separate types by prior.

**Theorem 4** (Impossibility of Screening by Prior)**.** *Let the principal offer any menu of contracts* $\mathcal{M} = \{(n(p), k(p))\}$ *to agents with private priors* $p_0$*. Suppose the agent's optimal stopping rule depends only on the ratio* $\rho = n/k(n)$*. Then:*

1. *All types* $p_0 \in (0, 1)$ *weakly prefer the contract with the highest ratio* $\rho$*.*
2. *No incentive-compatible menu can separate types by prior beliefs.*
3. *The principal's optimal mechanism is a* single pooling contract*, chosen to maximize expected revenue against the distribution* $F$*.*

Intuitively, the only feature of a contract that affects the agent's continuation value is its normalized precision $\rho$. Since all types agree on the ranking of $\rho$, any attempt at screening collapses: all types select the most informative contract. The principal therefore offers a single pooling contract, and and only agents with moderate priors choose to participate.

# 7 Empirical Implications and Measurement

## 7.1 Testable Predictions and Empirical Identification

A key feature of the model is that belief dynamics alone generate systematic variation in consultation duration across clients, even when the information quality is held fixed. In a drift-diffusion environment, posterior beliefs move most slowly near the center of the belief space. Therefore, even if two clients receive signals of identical precision, the one whose prior is closer to 1/2 will, in expectation, require a longer time to reach a stopping threshold. Formally, fixing the diffusion precision $n$,

$$E[T(p_0) \mid n] \text{ is maximized at } p_0 = 1/2.$$

This intrinsic property of Bayesian learning implies that clients with intermediate priors should naturally exhibit longer consultation durations, independent of the firm's strategic behavior.

Building on this baseline, the model predicts an additional layer of discrimination created by the information seller. The principal strategically chooses *lower* precision for clients with more uncertain priors, precisely because their slow belief dynamics already keep them engaged. This means that, relative to the baseline of purely mechanical Bayesian updating, we should observe:

1. an *inverse U-shaped* relationship between prior $p_0$ and consultation time, even under constant information quality; and
2. a *U-shaped* relationship between prior $p_0$ and information quality supplied, reflecting the principal's active manipulation.

Empirically, these two forces can be separated. The first arises mechanically from learning dynamics; the second identifies the firm's strategic response. Observing systematically lower information quality for clients with intermediate priors would be evidence of intentional slowing of learning.

## 7.2 Proxies for Information Quality

Direct measures of diffusion precision are rarely observable. However, many professional services markets offer natural proxies. A leading example is the *seniority* or *rank* of the consultant assigned to the client:

- Senior partners, attorneys, or analysts typically produce more precise, sharper, and faster-to-interpret assessments.
- Junior staff, trainees, or algorithmic assistants often produce noisier, more generic, or less refined recommendations.

Thus, consultant seniority serves as a reasonable empirical proxy for the precision parameter $n$. Assigning a junior consultant is analogous to providing a low-precision signal; assigning a senior expert is analogous to high precision.

Under the model, clients with intermediate priors should be disproportionately matched to lower-seniority consultants (i.e., lower precision), above and beyond what would be expected from belief dynamics alone.

## 7.3 Empirical Framework

Let $i$ index consultation episodes. Define:

- $p_{0i}$: the client's prior confidence (e.g., historical forecast accuracy, initial diagnostic guess, baseline volatility);
- $T_i$: consultation duration (session length, time-to-decision, number of exchanges);
- $s_i$: a measure of consultant seniority or expertise (inverse proxy for signal quality).

Two complementary regressions can be implemented.

### 1. Consultation duration.

$$T_i = \alpha + \delta_1 p_{0i} + \delta_2 p_{0i}^2 + X_i' \eta + \nu_i.$$

The mechanical drift–diffusion model predicts $\delta_1 > 0$ and $\delta_2 < 0$, giving an inverse-U even when $s_i$ is excluded.

### 2. Information quality (seniority assignment).

$$s_i = \alpha' + \beta_1 p_{0i} + \beta_2 p_{0i}^2 + X_i' \gamma + \varepsilon_i.$$

The strategic model predicts $\beta_1 < 0$ and $\beta_2 > 0$, indicating that clients with intermediate priors receive systematically lower-quality information—evidence consistent with discriminatory allocation of precision.

Because the baseline drift–diffusion effect operates through $T_i$ but not through $s_i$, any U-shaped pattern in seniority assignment cannot be explained mechanically and instead reflects the firm's strategic manipulation of information precision.

### 7.4 Interpretation and Welfare Implications

The model highlights a welfare-relevant divergence between socially optimal and profit-maximizing information provision. When clients are highly uncertain, their marginal value of information is highest. Yet precisely in this region, the principal supplies the *lowest* precision to prolong engagement. This generates two distortions:

- *Delay distortion*: clients with the greatest need for clarity receive the slowest effective learning.
- *Quality distortion*: information is deliberately degraded relative to the welfare-maximizing benchmark.

These distortions are particularly salient in legal advice, medical consulting, or financial planning, where delayed or degraded information directly affects decision quality. Policy interventions such as transparency in consultant assignment, client-controlled pacing of information, or minimum-quality standards may reduce the incentive to exploit client uncertainty.

Overall, the empirical framework distinguishes between the intrinsic learning dynamics of Bayesian updating and the firm's strategic manipulation of information quality. By examining seniority assignments and consultation durations jointly, the data can directly test whether firms exploit the slower learning of uncertain clients by selectively reducing the precision of information they receive.

## 8 Conclusion

This paper analyzes how an information seller optimally chooses the precision of a continuous-time signal when serving a client who can end the interaction at any time. The joint determination of the client's stopping thresholds and the principal's precision choice reveals a simple but powerful pattern: information quality is lowest when the client is most uncertain, and increases as the client becomes more confident. This creates a U-shaped relationship between optimal precision and the prior belief, reflecting the principal's incentive to slow learning when engagement is most valuable. A second prediction arises directly from drift–diffusion learning: even holding information quality fixed, clients with intermediate priors exhibit the longest consultation durations because belief evolution is slowest near the center of the belief space. This mechanical inverse-U pattern in consultation time provides a natural benchmark. Deviations from this benchmark, specifically the provision of systematically lower precision to uncertain clients, identify the principal's discriminatory manipulation of information speed.

The analysis shows that these mechanisms are robust to richer cost structures. Allowing the flow cost to depend on precision leaves the core logic intact: the agent's stopping boundaries depend only on the ratio $n/k(n)$, and the principal's expected revenue can still be written as a function of the stopping posterior. Consequently, the optimal stopping posteriors and the U-shaped relationship between optimal precision and prior beliefs are unchanged. What matters for behavior is the effective "precision per unit flow cost," not the specific form of the pricing rule. This robustness result suggests that the model's qualitative predictions should carry over to a wide variety of fee structures used in practice, from hourly billing to tiered retainers and performance-contingent fees. When the agent's prior is private information, the model further implies that standard screening via menus of contracts is infeasible. Because all types agree on the ranking of contracts by their normalized precision, any attempt to design a separating menu collapses: every type selects the contract with the highest precision–cost ratio. The optimal mechanism is therefore a single pooling contract, targeted toward the marginal client whose prior is closest to indifference. This impossibility of screening by priors underscores a key tension in information markets: sellers can finely tailor the path of information revelation for a given client, but they cannot use static menus to sort clients by their initial beliefs in this environment.

Taken together, these results generate a tractable empirical framework. By examining both the shape of consultation durations and the allocation of information quality, researchers can disentangle the inherent dynamics of Bayesian belief updating from the principal's strategic response. An inverse-U in consultation time combined with a U-shape in measured information quality for different prior levels would be consistent with the model's central mechanism: firms exploit slower learning among uncertain clients by selectively reducing the precision of the advice they receive. The analysis also highlights a welfare-relevant distortion. The clients who value information most, i.e. those with intermediate priors, receive the lowest precision and thus the slowest learning under the firm's optimal policy. This delay is not a mistake but a predictable feature of profit maximization. In settings where decision quality has large social or economic consequences, such as legal advice, medical diagnostics, or financial counseling, there may be a role for policy interventions that promote transparency in information quality, client-directed pacing of information, or minimum precision standards. The robustness and screening results suggest that such interventions should focus on the effective precision-cost tradeoff and on the allocation of high precision resources across clients, rather than on the particular billing scheme.

The analysis demonstrates that an information seller may optimally distort the learning process by reducing precision when the buyer is most uncertain. This mechanism provides a novel explanation for inefficiencies in markets for information: distortions can stem not only from bias or persuasion motives, but also from the seller's incentive to prolong engagement. The mechanism helps rationalize patterns observed in consulting relationships and algorithmic platform services, where compensation structures tied to time or ongoing attention may encourage slower information revelation. More broadly, the results underscore how the institutional design of advice markets or searching markets, including pricing, commitment and data privacy, can shape the efficiency of learning and decision-making.

## References


**Bergemann, Dirk, Alessandro Bonatti, and Alex Smolin**, "The design and price of information," *American economic review*, 2018, *108* (1), 1–48.

— **and** — , "Markets for information: An introduction," *Annual Review of Economics*, 2019, *11* (1), 85–107.

**Bloedel, Alexander W and Ilya R Segal**, "Persuasion with rational inattention," *Available at SSRN 3164033*, 2018.

**Che, Yeon-Koo, Kyungmin Kim, and Konrad Mierendorff**, "Keeping the listener engaged: a dynamic model of bayesian persuasion," *Journal of Political Economy*, 2023, *131* (7), 000–000.

**Eső, Péter and Balázs Szentes**, "The price of advice," *The Rand Journal of Economics*, 2007, *38* (4), 863–880.

**Fudenberg, Drew, Philipp Strack, and Tomasz Strzalecki**, "Stochastic choice and optimal sequential sampling," *arXiv preprint arXiv:1505.03342*, 2015.

**Gentzkow, Matthew and Emir Kamenica**, "Costly persuasion," *American Economic Review*, 2014, *104* (5), 457–462.

**Gonçalves, Duarte**, "Speed, Accuracy, and Complexity," *arXiv preprint arXiv:2403.11240*, 2024.

**Henry, Emeric and Marco Ottaviani**, "Research and the approval process: The organization of persuasion," *American Economic Review*, 2019, *109* (3), 911–955.

**Kamenica, Emir and Matthew Gentzkow**, "Bayesian persuasion," *American Economic Review*, 2011, *101* (6), 2590–2615.

**Ke, T Tony, Zuo-Jun Max Shen, and J Miguel Villas-Boas**, "Search for information on multiple products," *Management Science*, 2016, *62* (12), 3576–3603.

**Lehrer, Ehud and Tao Wang**, "The value of information in stopping problems," *Economic Theory*, 2024, *78* (2), 619–648.

**Liptser, Robert S and Albert N Shiryaev**, *Statistics of random processes: I. General theory*, Vol. 5, Springer Science & Business Media, 2013.

**Milgrom, Paul and Chris Shannon**, "Monotone comparative statics," *Econometrica: Journal of the Econometric Society*, 1994, pp. 157–180.

**Moscarini, Giuseppe and Lones Smith**, "The optimal level of experimentation," *Econometrica*, 2001, *69* (6), 1629–1644.

**Orlov, Dmitry, Andrzej Skrzypacz, and Pavel Zryumov**, "Persuading the principal to wait," *Journal of Political Economy*, 2020, *128* (7), 2542–2578.

**Persson, Petra**, "Attention manipulation and information overload," *Behavioural Public Policy*, 2018, *2* (1), 78–106.

**Topkis, Donald M**, "Minimizing a submodular function on a lattice," *Operations research*, 1978, *26* (2), 305–321.

**Wald, Abraham**, "Foundations of a general theory of sequential decision functions," *Econometrica, Journal of the Econometric Society*, 1947, pp. 279–313.

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

# A Proofs

## A.1 Derivation of the HJB equation of the agent

The agent's value function under precision level $n$ is

$$U(p_0; n) = \sup_T \mathbb{E}\left[u(p_T) - \int_0^T k\, dt \;\Big|\; p_0\right], \qquad dp_t = 2p_t(1-p_t)\sqrt{n}\, d\widetilde{W}_t,$$

where $u(p_T)$ is the expected payoff of the optimal static action at belief $p_T$.

Over a small interval $[0, h]$, the agent either stops immediately or continues for duration $h$:

$$U(p; n) = \max\left\{u(p),\ \mathbb{E}\big[-kh + U(p_h; n) \mid p_0 = p\big]\right\}.$$

By Itô's lemma, the generator of $p_t$ satisfies

$$\mathcal{L}f(p) = \frac{1}{2}\sigma^2(p)f''(p) = 2np^2(1-p)^2 f''(p),$$

so that

$$\mathbb{E}[U(p_h) \mid p_0 = p] = U(p) + 2np^2(1-p)^2 U''(p)h + o(h).$$

Substituting and simplifying,

$$U(p) = \max\Big\{u(p),\, U(p) + \big(2np^2(1-p)^2U''(p) - k\big)h + o(h)\Big\}.$$

Taking the limit as $h \to 0$, we obtain the HJB equation:

$$\boxed{0 = \max\left\{u(p) - U(p),\ 2np^2(1-p)^2U''(p) - k\right\}.} \tag{HJB–U}$$

In the *stopping region*, $U(p) = u(p)$ and $2np^2(1-p)^2U''(p) - k \leq 0$; in the *continuation region*, $U(p) > u(p)$ and the equality holds.

Defining the value of experimentation $V(p) := U(p) - u(p) \geq 0$, we obtain

$$0 = \max\left\{-V(p),\ 2np^2(1-p)^2\big(V''(p) + u''(p)\big) - k\right\}.$$

Since $u(p)$ is piecewise linear and $u''(p) = 0$ except at indifference points, the HJB simplifies to

$$\boxed{0 = \max\left\{-V(p),\ 2np^2(1-p)^2V''(p) - k\right\}.} \tag{HJB–V}$$

This characterizes the agent's optimal stopping rule in the absence of discounting.

## A.2 Monotonicity of the optimal upper boundary

**Theorem 5** (Monotonicity of the optimal upper boundary)**.** *Fix a prior $p_0 \in (1/2, 1)$ and consider the problem*

$$\sup_{\bar{p}\in[1/2,1)} f(\bar{p}, p_0), \qquad f(\bar{p}, p_0) := \frac{\pi}{4L'(\bar{p})}\big(L(\bar{p}) - L(p_0)\big),$$

*where*

$$L(p) := p \ln \frac{p}{1-p} + (1-p)\ln\frac{1-p}{p}, \qquad p \in (0,1).$$

*Let $\Gamma(p_0) \subset [1/2, 1)$ denote the set of maximizers.*

*Suppose that for each $p_0 \in (1/2, 1)$ the set $\Gamma(p_0)$ is nonempty. Then:*

1. *The function $f$ has (strictly) increasing differences in $(\bar{p}, p_0)$ on $(1/2, 1) \times (1/2, 1)$.*
2. *The argmax correspondence $\Gamma$ is nondecreasing in $p_0$ in the sense of set inclusion: if $1/2 < p_0 < p_0' < 1$, then*
$$\sup \Gamma(p_0) \ \le\ \inf \Gamma(p_0').$$
*In particular, there exists a selection $\bar{p}^* : (1/2, 1) \to [1/2, 1)$ with $\bar{p}^*(p_0) \in \Gamma(p_0)$ for all $p_0$ such that $\bar{p}^*$ is weakly increasing in $p_0$.*

*Proof.* Since the positive multiplicative constant $\pi/4$ does not affect the maximizer, it is convenient to work with

$$g(\bar{p}, p_0) := \frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})},$$

which has the same argmax as $f$ in $\bar{p}$ for each $p_0$.

On $(0, 1)$, the function $L$ is strictly convex, so $L''(p) > 0$ for all $p \in (0, 1)$, and $L'(p)$ is strictly increasing with $L'(1/2) = 0$. In particular, for $\bar{p} > 1/2$ we have $L'(\bar{p}) > 0$.

First compute the partial derivative of $g$ with respect to $p_0$:

$$\frac{\partial g}{\partial p_0}(\bar{p}, p_0) = \frac{-L'(p_0)}{L'(\bar{p})}.$$

Differentiating this once more with respect to $\bar{p}$ yields the cross–partial:

$$\frac{\partial^2 g}{\partial \bar{p}\, \partial p_0} = \frac{\partial}{\partial \bar{p}}\left(\frac{-L'(p_0)}{L'(\bar{p})}\right) = -L'(p_0) \cdot \frac{\partial}{\partial \bar{p}}\left(\frac{1}{L'(\bar{p})}\right) = -L'(p_0)\left(-\frac{L''(\bar{p})}{[L'(\bar{p})]^2}\right) = \frac{L'(p_0)\, L''(\bar{p})}{[L'(\bar{p})]^2}.$$

On the domain of interest, $p_0 > 1/2$ and $\bar{p} > 1/2$, we have $L'(p_0) > 0$, $L''(\bar{p}) > 0$, and $[L'(\bar{p})]^2 > 0$, so

$$\frac{\partial^2 g}{\partial \bar{p}\, \partial p_0}(\bar{p}, p_0) > 0 \quad \text{for all } (\bar{p}, p_0) \in (1/2, 1) \times (1/2, 1).$$

Thus $g$ (and hence $f$) has strictly increasing differences in $(\bar{p}, p_0)$ on this domain.

Let $X := [1/2, 1)$ with the usual order, which is a complete lattice, and let $T := (1/2, 1)$ be the parameter space for $p_0$. By the monotone maximum theorem of Topkis (Topkis (1978); see also Milgrom and Shannon (1994)), if a function has increasing differences in $(x, t) \in X \times T$, then its argmax correspondence is nondecreasing in $t$ with respect to the usual order on $X$. Applying this result to $g$, we obtain that $\Gamma(p_0) = \arg\max_{\bar{p} \in X} f(\bar{p}, p_0)$ is nondecreasing in $p_0$: if $1/2 < p_0 < p_0' < 1$, then

$$\sup \Gamma(p_0) \ \leq \ \inf \Gamma(p_0').$$

The last statement about the existence of a weakly increasing selection follows from standard properties of increasing set–valued maps on a totally ordered set.

For $p_0 < 1/2$, we can alternatively apply Topkis Theorem for $\underline{p}$, and since $\bar{p}$ strictly decreases in $\underline{p}$, we can derive our main result that $\bar{p}$ is minimized when $p_0 = 1/2$. As a result, the optimal consultation level is also minimized when $p_0 = 1/2$.

□

## A.3 Alternative Proof of the Comparative Statics Using Implicit Function Theorem

**Lemma 2** (Comparative statics via the implicit function theorem)**.** *Fix* $p_0 \in (1/2, 1)$ *and consider*

$$\sup_{\bar{p}>1/2} f(\bar{p}, p_0), \qquad f(\bar{p}, p_0) := \frac{\pi}{4} \frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})},$$

*where*

$$L(p) := p \ln \frac{p}{1-p} + (1-p) \ln \frac{1-p}{p}, \qquad p \in (0, 1).$$

*Suppose that for each* $p_0 \in (1/2, 1)$ *there exists a unique interior maximizer* $\bar{p}^*(p_0) \in (1/2, 1)$ *such that:*

*[(i)]*

1. $f$ *is twice continuously differentiable in* $(\bar{p}, p_0)$ *on* $(1/2, 1) \times (1/2, 1)$*;*
2. *the first-order condition* $f_{\bar{p}}(\bar{p}^*(p_0), p_0) = 0$ *holds; and*
3. *the second-order condition* $f_{\bar{p}\bar{p}}(\bar{p}^*(p_0), p_0) < 0$ *holds.*

*Then* $\bar{p}^*(p_0)$ *is strictly increasing in* $p_0$ *on* $(1/2, 1)$*.*

*Proof.* Since the positive constant $\pi/4$ does not affect the argmax or the signs of derivatives, it is convenient to work with

$$g(\bar{p}, p_0) := \frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})},$$

so that $f(\bar{p}, p_0) = (\pi/4)\, g(\bar{p}, p_0)$ and, in particular,

$$f_{\bar{p}p_0} = \frac{\pi}{4} g_{\bar{p}p_0}, \qquad f_{\bar{p}\bar{p}} = \frac{\pi}{4} g_{\bar{p}\bar{p}}.$$

For each $p_0$, the interior maximizer $\bar{p}^*(p_0)$ satisfies the first-order condition

$$f_{\bar{p}}(\bar{p}^*(p_0), p_0) = 0 \quad \Longleftrightarrow \quad g_{\bar{p}}(\bar{p}^*(p_0), p_0) = 0.$$

Define

$$F(\bar{p}, p_0) := g_{\bar{p}}(\bar{p}, p_0).$$

By assumption (ii), $F(\bar{p}^*(p_0), p_0) = 0$ for all $p_0 \in (1/2, 1)$; by (iii),

$$F_{\bar{p}}(\bar{p}^*(p_0), p_0) = g_{\bar{p}\bar{p}}(\bar{p}^*(p_0), p_0) < 0.$$

Hence the implicit function theorem applies and yields a continuously differentiable function $\bar{p}^*(p_0)$ solving $F(\bar{p}^*(p_0), p_0) = 0$ with derivative

$$\frac{d\bar{p}^*}{dp_0}(p_0) = -\frac{F_{p_0}(\bar{p}^*(p_0), p_0)}{F_{\bar{p}}(\bar{p}^*(p_0), p_0)} = -\frac{g_{\bar{p}p_0}(\bar{p}^*(p_0), p_0)}{g_{\bar{p}\bar{p}}(\bar{p}^*(p_0), p_0)}.$$

It remains to sign $g_{\bar{p}p_0}$. First compute

$$g(\bar{p}, p_0) = \frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})}.$$

Differentiating with respect to $p_0$ gives

$$g_{p_0}(\bar{p}, p_0) = \frac{-L'(p_0)}{L'(\bar{p})}.$$

Differentiating this with respect to $\bar{p}$ yields the mixed partial:

$$g_{\bar{p}p_0}(\bar{p}, p_0) = \frac{\partial}{\partial \bar{p}}\left(\frac{-L'(p_0)}{L'(\bar{p})}\right) = -L'(p_0)\,\frac{\partial}{\partial \bar{p}}\left(\frac{1}{L'(\bar{p})}\right) = -L'(p_0)\left(-\frac{L''(\bar{p})}{[L'(\bar{p})]^2}\right) = \frac{L'(p_0)\,L''(\bar{p})}{[L'(\bar{p})]^2}.$$

On $(0, 1)$, the function $L$ is strictly convex, so $L''(p) > 0$ for all $p \in (0, 1)$, and $L'$ is strictly increasing with $L'(1/2) = 0$. In particular, for $p_0 > 1/2$ and $\bar{p} > 1/2$ we have $L'(p_0) > 0$, $L''(\bar{p}) > 0$, and $[L'(\bar{p})]^2 > 0$, so

$$g_{\bar{p}p_0}(\bar{p}, p_0) > 0 \quad \text{for all } (\bar{p}, p_0) \in (1/2, 1) \times (1/2, 1).$$

Evaluating at $(\bar{p}^*(p_0), p_0)$ and using $g_{\bar{p}\bar{p}}(\bar{p}^*(p_0), p_0) < 0$ from the second-order condition, we obtain

$$\frac{d\bar{p}^*}{dp_0}(p_0) = -\frac{g_{\bar{p}p_0}(\bar{p}^*(p_0), p_0)}{g_{\bar{p}\bar{p}}(\bar{p}^*(p_0), p_0)} > 0.$$

Thus $\bar{p}^*(p_0)$ is strictly increasing in $p_0$ on $(1/2, 1)$. □

## A.4 Proof of Theorem 2

Fix any prior $p_0 \in (0, 1)$ and consider

$$V(p_0) := \sup_{\bar{p}>1/2} f(\bar{p}, p_0), \qquad f(\bar{p}, p_0) := (h + l)\frac{L(\bar{p}) - L(p_0)}{L'(\bar{p})}.$$

Suppose that for each $p_0$ in an interval $I \subset (1/2, 1)$:

1. there exists a unique maximizer $\bar{p}^*(p_0) \in (1/2, 1)$;
2. the maximizer is interior and satisfies the first-order condition

$$f_{\bar{p}}(\bar{p}^*(p_0), p_0) = 0$$

   and the second-order condition

$$f_{\bar{p}\bar{p}}(\bar{p}^*(p_0), p_0) < 0;$$

3. $f$ is continuously differentiable in $p_0$ and twice continuously differentiable in $\bar{p}$ in a neighborhood of $(\bar{p}^*(p_0), p_0)$.

By definition,

$$V(p_0) = f(\bar{p}^*(p_0), p_0).$$

The envelope theorem applies to the parametric maximization problem $\sup_{\bar{p}>1/2} f(\bar{p}, p_0)$, giving

$$V'(p_0) = \frac{\partial f}{\partial p_0}(\bar{p}^*(p_0), p_0).$$

Holding $\bar{p}$ fixed, differentiation yields

$$\frac{\partial f}{\partial p_0}(\bar{p}, p_0) = -\frac{\pi}{4}\frac{L'(p_0)}{L'(\bar{p})}.$$

Thus

$$V'(p_0) = -\frac{\pi}{4} \frac{L'(p_0)}{L'(\bar{p}^*(p_0))}.$$

Since $L$ is strictly convex and symmetric around $1/2$, we have:

$$L'(1/2) = 0, \qquad L'(p) < 0 \text{ for } p < 1/2, \qquad L'(p) > 0 \text{ for } p > 1/2.$$

On $I \subset (1/2, 1)$, both $p_0$ and $\bar{p}^*(p_0)$ exceed $1/2$, so $L'(p_0) > 0$ and $L'(\bar{p}^*(p_0)) > 0$. Therefore,

$$V'(p_0) < 0 \qquad \text{for all } p_0 \in I.$$

An identical argument applied to $(0, 1/2)$ (symmetry of $f$ and of the optimizers) implies $V'(p_0) > 0$ on $(0, 1/2)$.

Combining the two inequalities and the symmetry $V(p_0) = V(1 - p_0)$ yields the stated result: $p_0 = 1/2$ is the global minimizer of $V$ on $(0, 1)$. □