

Trust, but Verify: Information Acquisition in Credence-Good Markets*

David Kim[†]

Job Market Paper

[Click here for the latest version](#)

October 19, 2025

Abstract

This paper studies consumer information acquisition in credence-good markets. An expert first sets a price menu, observes the consumer’s problem, and makes a treatment recommendation. The consumer can then acquire costly information to verify the expert’s recommendation. We find a non-monotonic relationship between verification costs and consumer surplus. At prohibitively high or zero cost, consumer surplus is zero, sustained by wasteful rejection or full surplus extraction, respectively. For an intermediate range of costs, a mutually beneficial outcome emerges where the expert builds credibility more efficiently and the consumer secures a positive surplus. In contrast to the existing credence-good literature, we show that expert cheating arises as an intrinsic feature of the market, without consumer heterogeneity. This finding also informs the debate on the impact of new information technologies, suggesting that the proliferation of low-cost verification tools (e.g., AI-powered diagnostics) is beneficial to consumers by shifting the market into this Pareto-improving range.

*I am deeply indebted to my advisors, Fei Li and Peter Norman, for their guidance and support. I am also grateful to my committee members—Gary Biglaiser, Ilwoo Hwang, and Curtis Taylor—and to Ting Liu for many valuable discussions and insightful suggestions. I also thank Jaden Chen, Can Urgan, and participants at the UNC theory workshop for helpful comments. All remaining errors are my own.

[†]Department of Economics, University of North Carolina at Chapel Hill. Email: dskim10@email.unc.edu

Contents

1	Introduction	1
1.1	Related Literature	3
2	Model	5
3	The Recommendation Subgame	12
3.1	Equilibrium Refinements and Simplifications	12
3.2	The Core Strategic Problem	13
4	Benchmark Analysis	17
4.1	Free Learning	18
4.2	No Learning	19
5	Main Result	20
5.1	Intermediate Learning Cost	20
5.2	Numerical Illustration	23
6	Conclusion	26
A	Appendix: Proofs	29
A.1	Omitted Proofs in section 4	29
A.2	Proof of Proposition 1	32
A.3	Proof of Proposition 2	34
A.4	Proof of Theorem 1	35
A.5	Expanding \mathcal{P}	42

1 Introduction

When hiring a repair technician, consulting a medical doctor, or seeking legal advice, consumers must rely on the specialized knowledge of experts. These markets are classic examples of credence goods, defined by a fundamental information asymmetry: the expert is privately informed about the consumer’s true problem, while the consumer is not. This creates a core agency problem, as the expert has a clear financial incentive to exploit this advantage. Specifically, he can misrepresent a minor problem as a major one in order to recommend an unnecessarily expensive treatment, a practice often referred to as overtreatment. This paper investigates how the consumer’s ability to acquire information to verify an expert’s claim shapes this strategic landscape.

This question has taken on new urgency as information acquisition has become cheaper and more flexible. Historically, consumer verification was a high-cost, rigid process, such as seeking a second opinion from another expert or sifting through fragmented, often unreliable information online. Recent AI-enabled tools—from assistive diagnostics to conversational search—lower verification costs and provide an on-demand “digital second opinion.” The intuitive policy view is that reducing high information costs should empower consumers.

However, the strategic relationship between expert credibility and consumer learning creates a delicate balance, locking the players in a feedback loop. High expert credibility can breed consumer complacency, discouraging costly verification. This lack of scrutiny, however, can create a strong incentive for the expert to exploit that trust and cheat. Conversely, a high perceived risk of being cheated motivates consumers to engage in verification, which in turn disciplines the expert, restoring credibility and completing the strategic cycle. This dynamic suggests that the cost of information is a critical determinant of the market outcome, as it governs the incentives that can either sustain or break this cycle. This tension motivates the central question of our paper: how does the cost of consumer verification shape market outcomes, and who ultimately captures the gains from trade?

To formalize this puzzle, we develop a model of a monopolistic credence-good market. A consumer faces one of two problems—major or minor—which the expert observes perfectly. We make standard assumptions that it is socially efficient to treat both problems, but the cost of the major treatment exceeds the consumer’s ex-ante expected benefit, creating a “lemon market” problem that can lead to inefficient trade. In our model, the expert first commits to a price menu for the two treatments. After the expert diagnoses the problem and makes a recommendation,

the consumer can flexibly and endogenously choose an information structure—a “test”—to verify the expert’s claim. The cost of this test is increasing in its informational content, with a unit cost of information, k , serving as the key parameter of our analysis. This modeling choice distinguishes our work by focusing not on the optimal design of a given signal, but on the consumer’s strategic decision of how much and what kind of information to acquire in response to the expert’s recommendation.

Our central finding reveals a “paradox of friction”: the consumer is strictly better off in the presence of moderate information cost compared to when it is perfectly free. We show that for any cost of information in the intermediate range, the equilibrium outcome is a strict Pareto improvement over the benchmark where learning is prohibitively costly ($k = \infty$). In this outcome, the consumer captures a strictly positive surplus. In sharp contrast, when learning is frictionless ($k = 0$), the consumer’s surplus falls back to zero, and all gains from the efficient market accrue to the expert.

The intuition for this result is driven by what we term the “price of credibility,” which inverts the standard logic of information rents. Instead of the uninformed consumer paying the informed expert to reveal information, here the informed expert effectively subsidizes the uninformed consumer’s learning by offering a price cut on the expensive treatment. A positive cost of learning ($k > 0$) grants the consumer a powerful strategic tool: the credible commitment to *not* learn if the price is too high. To avoid losing the sale, the expert must offer a price cut large enough to make learning a worthwhile investment, and this subsidy becomes the consumer’s surplus. When learning is free, this commitment power vanishes, as the consumer cannot credibly threaten not to verify.

This mutually beneficial outcome, however, requires the expert’s incentive to cheat to persist. We show that the expert’s potential for opportunism is a “necessary evil.” If the expert were known to be perfectly honest, the consumer would have no reason to pay for verification. The threat of consumer learning would disappear, dissolving the expert’s more efficient commitment device and causing the market to unravel back to an inefficient benchmark. The very possibility of being cheated is what underpins the value of verification and allows for the surplus-generating equilibrium to exist.

This paper makes two primary contributions. First, to the credence-good literature, we offer a new rationale for the persistence of expert fraud, showing it can be an intrinsic feature of the market even with a representative, homogeneous consumer. While previous work often relies

on consumer heterogeneity to explain fraud, our model generates opportunism endogenously: to sustain the consumer’s incentive to verify, the expert must cheat with positive probability. This finding suggests that policy and empirical work should consider drivers of fraud beyond observable consumer characteristics. Second, our findings inform the recent debate on the impact of AI and other low-cost diagnostic tools. Our model suggests these technologies are beneficial for consumers precisely because they shift the market from a high-cost benchmark to a Pareto-improving intermediate range where both experts and consumers are strictly better off. This highlights a subtle but important policy takeaway: while reducing prohibitively high information costs is valuable, the ultimate goal should not be to make verification perfectly frictionless. As our analysis shows, driving costs to zero would inadvertently disempower consumers by removing the very friction that grants them bargaining power.

1.1 Related Literature

This paper lies at the intersection of two strands of literature: credence-goods and consumer information acquisition. By introducing a consumer with a flexible, costly learning technology into a classic credence-good model, we contribute to both. First, we provide a new rationale for the persistence of expert fraud, a central puzzle in the credence-good literature. Second, our finding—that a moderate cost of information can generate a mutually beneficial outcome—adds to the active debate in the information acquisition literature on how frictions shape market outcomes.

Credence-Goods. We study the role of consumer information in the credence-good market, starting with [Darby and Karni \(1973\)](#), [Pitchik and Schotter \(1987\)](#), [Wolinsky \(1993\)](#), and [Taylor \(1995\)](#). A key puzzle in this literature has been to generate equilibrium outcomes that reflect the real-world phenomenon of expert cheating. To generate cheating in equilibrium, the literature has typically relied on introducing consumer heterogeneity. [Fong \(2005\)](#) shows that if consumers have observable differences (e.g., genders), the expert will selectively cheat certain types. [Hyndman and Ozerturk \(2011\)](#) provide a micro-foundation for this by showing how a public signal can create such observable types. [Kim et al. \(2025\)](#), takes this a step further by showing that even unobservable consumer heterogeneity, in the form of an exogenous private signal, can lead to expert cheating. Similarly, others examine the expert’s response when he can observe the consumer’s information directly ([Jost et al. \(2021\)](#)). In sharp contrast, our

work differs by focusing on the representative, homogeneous consumer’s ability to endogenously acquire information after a recommendation.

Information Acquisition. This paper also contributes to the growing literature on information acquisition—often referred to as the rational inattention—pioneered by [Sims \(2003\)](#). Our result resonates with findings in recent works in that consumers may benefit from inattentive learning. However, the underlying economic mechanisms are distinct and depend on the specific friction being modeled. [Ravid \(2020\)](#) studies a bargaining environment where a rationally inattentive buyer pays costs to interpret a complex contract. [Fabbri \(2025\)](#) identifies an “attention holdup” problem where a seller can exploit the consumer’s sunk learning costs. Our work provides a new rationale rooted in the agency problem of a credence-good market: the consumer’s surplus arises from a “price of credibility”, where an informed expert strategically subsidizes verification to build trust more efficiently.

Methodologically, our analysis also departs from these papers by establishing the result for a general class of information costs. Whereas [Ravid \(2020\)](#) and [Fabbri \(2025\)](#) study environments with entropic cost functions, our analysis accommodates any cost function satisfying standard properties such as Blackwell monotonicity. This shows that our mechanism is not tied to a particular functional form but reflects a more general feature of the strategic environment. Our assumptions cover a wide class of costs, including the log-likelihood ratio (LLR) costs axiomatically characterized by [Pomatto et al. \(2023\)](#).

Our work also relates to [Chatterjee et al. \(2025\)](#), whose model shares our feature of post-offer learning. The primary distinction is the learning technology: [Chatterjee et al. \(2025\)](#) assumes the buyer chooses from a binary signal that can conclusively confirm the high value but provides no information otherwise (“*H*-focused” in their language), while our consumer can choose from a fully flexible information structure. Our focus on an informed expert also distinguishes our work from models with uninformed sellers, which either solve for optimal screening mechanisms ([Mensch \(2022\)](#)) or find that pre-offer learning can lead to Pareto-worst outcomes ([Ravid et al. \(2022\)](#)).

Organization The rest of the paper is organized as follows. Section 2 presents the formal model. Section 3 defines and simplifies the recommendation subgame. Section 4 analyzes the benchmark cases. Section 5 presents our main result and illustrates it with a numerical simulation. Section 6 concludes. All proofs are in the Appendix.

2 Model

This section formally develops the model. We begin by describing the economic environment, including the players, payoffs, and key parametric assumptions. We then detail the timing of the game, and define the strategies and equilibrium concept.

Environment. There are two risk-neutral players: an *expert* (he) and a *consumer* (she). The consumer has a problem of unknown severity, $\theta \in \{H, L\}$. The problem is major ($\theta = H$) with a commonly known prior probability $\mu_0 \in (0, 1)$ and minor ($\theta = L$) otherwise. If the problem is treated, the consumer has state-dependent payoff of v_θ . An equivalent interpretation is, if the problem is left untreated, the consumer incurs a state-dependent loss v_θ . We assume $v_L < v_H$. This captures settings where failing to address a real issue (high state) is costly—for example, not fixing an actual engine defect risks severe failure (e.g., seizure on the road)—whereas ignoring a benign issue (low state), such as a faulty sensor reading, entails much smaller downside.

The expert can treat problem θ by incurring a state-dependent cost c_θ . The expert can do one of three options: either recommend treatment $\hat{\theta} \in \{H, L\}$ at the price $p_{\hat{\theta}}$ or refuse treatment $\hat{\theta} = \emptyset$. The expert is liable to fix the consumer’s problem after the consumer accepts the recommendations. However, the consumer do not observe which type of treatment was performed by the expert.

Parameters. Treating the problem in each state is efficient: $c_\theta < v_\theta, \forall \theta \in \{L, H\}$. To focus on interesting cases where the consumer has learning incentive after being recommended an expensive treatment from the expert, we assume the lemon condition:

$$E[v] = \mu_0 v_H + (1 - \mu_0) v_L < c_H.$$

Therefore, a price for H -treatment the expert sets above his marginal cost is too expensive for the consumer to accept without learning.

Timing.

Stage 0: The expert posts $\mathbf{p} = (p_L, p_H) \in \mathbb{R}_+^2$.

Stage 1: Nature draws $\theta \in \{H, L\}$.

recommendation game

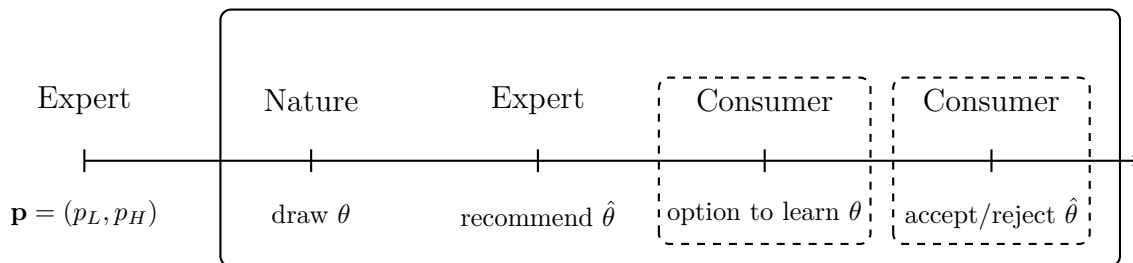


Figure 1: Timeline

Stage 2: The expert observes θ and recommends $\hat{\theta} \in \{L, H, \text{no treatment}\}$.

Stage 3: After observing the recommendation $\hat{\theta}$, the consumer chooses

- (i) whether and how to acquire costly information about θ ,
- (ii) decides whether to accept or reject the recommended treatment $\hat{\theta}$.

A key feature of our model is the expert has commitment to a price schedule in the beginning. Importantly, this price commitment does not confer extra leverage in the recommendation game. Because the expert can always choose “no treatment” at Stage 2, any price that would make treatment unattractive relative to not treating will simply not be implemented in a sequentially rational continuation.

This structure is a deliberate choice to reflect the institutional realities of many credence-good markets. For instance, in the auto repair industry, shops often use standardized software like ALLDATA or Mitchell 1—comprehensive databases that provide official estimates for both parts and average labor hours across different services—to generate a fixed price schedule. Similarly, in healthcare, medical procedures are tied to CPT codes with pre-negotiated reimbursement rates. This institutional framework forces the expert to commit to a service-contingent price.

Our model is therefore designed to isolate this crucial dimension of trust and strategic information acquisition. The central strategic tension does not arise from price bargaining—consumers rarely bargain with a doctor over the price of a specific surgery. Instead, the core agency problem, which we model in the recommendation subgame, is one of credibility. The consumer’s primary concern is not the price of the recommended treatment, but the truthfulness

of the recommendation itself.

Information Acquisition. After observing the expert’s recommendation, the consumer can choose to investigate the state θ by designing and deploying an experiment, also known as *signal structure* (or simply, *signals*). Formally, a signal structure denoted by

$$\phi : \{H, L\} \rightarrow \Delta(S)$$

is a statistical test that maps the true state of the problem, $\theta \in \{H, L\}$, to a distribution over the set of signal realizations S . This framework encompasses a spectrum of familiar information structures:

- *An Uninformative Signal:* The consumer could choose a signal that is completely uncorrelated with the true problem, like flipping a coin. For any signal realization $s \in S$, the probability of it occurring is the same in either state. This means the likelihood ratio $\frac{\phi_H(s)}{\phi_L(s)} = 1$ for all s , and observing the signal does not update the consumer’s belief.
- *A Fully Revealing Signal:* The consumer could seek a perfectly accurate test. This signal structure perfectly separates the states. The simplest example uses a signal space with two realizations, $S = \{s_H, s_L\}$, where s_H is conclusive evidence of problem H , and s_L is conclusive evidence of problem L . Formally, this means $\phi_H(s_H) = 1$, $\phi_L(s_L) = 1$.
- *A Partially Informative Signal:* Most realistically, the consumer designs a signal where each realization $s_i \in S$ provides some evidence. The informativeness of each piece of evidence s_i is captured by its likelihood ratio, $\gamma_i = \frac{\phi_H(s_i)}{\phi_L(s_i)}$. Observing a signal s_i with a high likelihood ratio ($\gamma_i \gg 1$) is strong evidence for the major problem, while observing one with a low ratio ($\gamma_i \ll 1$) is strong evidence for the minor problem. Our analysis focuses on the consumer’s optimal choice of a signal structure, which is equivalent to choosing an optimal distribution of these likelihood ratios.

Information Cost. Acquiring an information structure ϕ entails a cost

$$k \cdot C(\phi),$$

with two components:

- (a) A unit cost $k \geq 0$ capturing how expensive information is for the consumer. This is the key comparative-statics parameter in our analysis.
- (b) A prior-independent cost function $C(\phi)$ that is Blackwell monotone (weakly increasing in informativeness of the experiment ϕ).¹

The Revelation Principle. Following a revelation-principle-style argument (e.g., [Woodford \(2009\)](#)), it is without loss to restrict attention to binary experiments in the form of action recommendations.

In principle, the consumer may choose any experiment ϕ , which is a rich object. What ultimately matters, however, is the *action* taken after the signal is observed. If two distinct signals always induce the same action, then distinguishing between them is decision-irrelevant. Merging such signals into a single signal leaves induced behavior unchanged and only coarsens the experiment, which by Blackwell monotonicity cannot increase cost. Iterating this merging argument yields an experiment in which each signal is itself a *action recommendation*.

Hence, without loss of generality, we can represent the consumer’s choice by a binary recommendation experiment

$$\phi = (\phi_H, \phi_L) \in [0, 1]^2.$$

The object $\phi = (\phi_H, \phi_L)$ fully describes a binary recommendation experiment: ϕ_θ is the probability the consumer receives the recommendation “accept” in state $\theta \in \{H, L\}$. Equivalently, letting A denote the “accept” recommendation and R the “reject” recommendation, Bayes’ rule gives the posterior beliefs

$$\mu(H | A) = \frac{\mu_0 \phi_H}{\mu_0 \phi_H + (1 - \mu_0) \phi_L}, \quad \mu(H | R) = \frac{\mu_0 (1 - \phi_H)}{\mu_0 (1 - \phi_H) + (1 - \mu_0) (1 - \phi_L)}.$$

Two benchmarks are useful. (i) *Uninformative*: $\phi_H = \phi_L$, in which case the recommendation is independent of the state and posteriors equal the prior ($\mu(H | A) = \mu(H | R) = \mu_0$). (ii) *Perfect*: $(\phi_H, \phi_L) = (1, 0)$, which yields correct recommendations with probability one.

Since we will work with binary, action-recommendation experiments, we impose the following mild regularity on C as a function of (ϕ_H, ϕ_L) :

Assumption 1. *The cost $C : [0, 1]^2 \rightarrow \mathbb{R}_+ \cup \{+\infty\}$ satisfies:*

¹Formally, $C(\cdot)$ satisfies Blackwell monotonicity if for any pair of experiments ϕ', ϕ for which ϕ' is no less informative than ϕ , it must be that $C(\phi') \geq C(\phi)$.

state	accept	reject
$\theta = H$	ϕ_H	$1 - \phi_H$
$\theta = L$	ϕ_L	$1 - \phi_L$

Figure 2: Binary recommendation experiment.

(i) Continuity: C is continuous on $(0, 1)^2$ (it may take the value $+\infty$ on the boundary).

(ii) Strict convexity: C is strictly convex on the set $\{(\phi_L, \phi_H) \in (0, 1)^2 : \phi_L \neq \phi_H\}$ (we normalize $C(\phi) = 0$ whenever $\phi_L = \phi_H$, i.e., when the experiment is uninformative).

(iii) Decreasing differences (submodularity): for any $\phi_H < \phi'_H$ and $\phi_L < \phi'_L$,

$$C(\phi'_L, \phi'_H) - C(\phi'_L, \phi_H) \leq C(\phi_L, \phi'_H) - C(\phi_L, \phi_H).$$

Conditions (ii) and (iii) capture the standard property of diminishing marginal returns. Condition (ii) implies that the marginal cost of increasing each signal realization's informativeness is increasing. Condition (iii) implies that the marginal cost of improving one signal realization is higher when the other is already more informative.

These properties are satisfied by a broad class of established specifications, including the log-likelihood-ratio (LLR) cost of Pomatto et al. (2023).

Example 1 (Log-likelihood-ratio (LLR) cost). *A prominent specification is the (possibly weighted) sum of Kullback–Leibler divergences between the state-conditional signal:*

$$C(\phi) = \sum_{\theta, \theta' \in \{H, L\}} \beta_{\theta\theta'} D_{KL}(\phi_\theta || \phi_{\theta'}),$$

where $\beta_{\theta\theta'} \geq 0$ capture the intrinsic difficulty of distinguishing between the two states and

$$D_{KL}(\phi_\theta || \phi_{\theta'}) = \phi_\theta \ln \frac{\phi_\theta}{\phi_{\theta'}} + (1 - \phi_\theta) \ln \frac{1 - \phi_\theta}{1 - \phi_{\theta'}},$$

is the Kullback–Leibler divergence between $\phi_\theta, \phi_{\theta'}$. If $\phi_H = \phi_L$ (uninformative), then every likelihood ratio equals 1 and $C(\phi) = 0$. The LLR cost is Blackwell monotone, strictly convex on informative experiments, and exhibits decreasing differences.

Another prior-independent example is the squared Hellinger distance. It measures how similar the two state-conditional distributions are—the more similar, the smaller the distance (and it is zero when they are identical).²

Example 2 (Squared Hellinger cost). *A second prior-independent specification is based on the squared Hellinger distance between the state-conditional distribution. For binary signals it can be written as*

$$C(\phi) = \sum_{\theta, \theta' \in \{H, L\}} \beta_{\theta\theta'} H^2(\phi_\theta, \phi_{\theta'}),$$

where $\beta_{\theta\theta'} \geq 0$, and for $\phi_\theta, \phi_{\theta'} \in (0, 1)$,

$$H^2(\phi_\theta, \phi_{\theta'}) = (\sqrt{\phi_\theta} - \sqrt{\phi_{\theta'}})^2 + (\sqrt{1 - \phi_\theta} - \sqrt{1 - \phi_{\theta'}})^2.$$

Like LLR cost in the previous example, the Hellinger cost satisfies Assumption 1.

Payoffs. The payoffs for the expert and consumer depend on the recommendation, the consumer's learning and acceptance decisions, and the prices. If the consumer acquires an experiment ϕ at cost $k \cdot C(\phi)$ and accepts the expert's recommendation $\hat{\theta} \in \{H, L\}$ at price $p_{\hat{\theta}}$, her net payoff is $v_{\hat{\theta}} - p_{\hat{\theta}} - kC(\phi)$. The expert's corresponding profit is $p_{\hat{\theta}} - c_{\hat{\theta}}$.

If the consumer rejects the recommendation, her payoff is solely the cost of information, $-kC(\phi)$, and the expert's profit is zero. If the expert chooses not to treat ($\hat{\theta} = \emptyset$), both get zero.

Strategies and Belief. The expert's strategy is a pair (\mathbf{p}, β) . First, he chooses a price menu $\mathbf{p} = (p_L, p_H) \in \mathbb{R}_+^2$. Second, after observing the true state of the problem $\theta \in \{H, L\}$, the expert chooses a recommendation probability

$$\beta : \{H, L\} \times \mathbb{R}_+^2 \rightarrow \Delta(\{L, H, \emptyset\})$$

where $\beta(\hat{\theta}|\theta, \mathbf{p})$ is the probability of recommending $\hat{\theta}$ in state θ and given price \mathbf{p} .

Given the price \mathbf{p} and following the expert's strategy, the consumer belief that the true problem is $\theta = H$ following $\hat{\theta}$ -treatment recommendation is updated using Bayes-rule:

$$\mu(H|\hat{\theta}, \mathbf{p}) = \frac{\mu_0 \beta(\hat{\theta}|H, \mathbf{p})}{\mu_0 \beta(\hat{\theta}|H, \mathbf{p}) + (1 - \mu_0) \beta(\hat{\theta}|L, \mathbf{p})}.$$

²This falls under the state-wise divergence costs analyzed by Cheng and Kim (2025, §4.4), which are Blackwell monotone.

From our discussion of the revelation-principle-style argument, the consumer’s strategy is that of choosing recommendation probabilities

$$\phi = (\phi_H, \phi_L) : \{H, L\} \times \mathbb{R}_+^2 \rightarrow [0, 1]^2$$

where $\phi_\theta(\hat{\theta}, \mathbf{p})$ is the probability of receiving the “accept” signal realization (to accept the expert’s $\hat{\theta}$ recommendation) conditional on the true state θ of the problem and given price \mathbf{p} .

Equilibrium Notion. In our environment each price \mathbf{p} set at the beginning of the game defines a well-specified subgame. We impose subgame perfection by adopting the *almost perfect Bayesian equilibrium* (*almost PBE*) concept of Mailath (2018). An assessment $((\mathbf{p}, \beta), \phi; \mu)$ is an *almost PBE* if

- (i) for each price \mathbf{p}' , $(\beta(\cdot, \mathbf{p}'), \phi(\cdot, \mathbf{p}'), \mu(\cdot, \mathbf{p}'))$ is *almost PBE* in the subgame following \mathbf{p}' , where

$$\mu(H|\hat{\theta}, \mathbf{p}') = \frac{\mu_0 \beta(\hat{\theta} | H, \mathbf{p}')}{\mu_0 \beta(\hat{\theta} | H, \mathbf{p}') + (1 - \mu_0) \beta(\hat{\theta} | L, \mathbf{p}')} .^3 \quad (\star)$$

- (ii) \mathbf{p} is optimal for the expert given the continuation play across all prices:

$$\mathbf{p} \in \arg \max_{\mathbf{p}' \in \mathbb{R}_+^2} \pi(\mathbf{p}'),$$

where $\pi(\mathbf{p}')$ is the expert’s profit in the subgame equilibrium $(\beta(\cdot, \mathbf{p}'), \phi(\cdot, \mathbf{p}'), \mu(\cdot, \mathbf{p}'))$.

A key feature of *almost PBE* is that the Bayes requirement in (\star) is *conditional on the subgame root* (here, any price \mathbf{p}'): it does not require \mathbf{p}' to occur on the equilibrium path. Thus, even if \mathbf{p}' is reached with zero unconditional probability, (\star) applies whenever the downstream recommendation node is reached with positive probability conditional on \mathbf{p}' .

We further refine by eliminating weakly dominated strategies for both players. Any *almost PBE* satisfying this refinement is referred to simply as an *equilibrium*.

³In our finite two-period extensive-form game with perfect recall, the standard notions coincide: *almost PBE*, *weak PBE*, and sequential equilibrium yield the same set of outcomes. See Mailath (2018, Chs. 11–12) on *almost PBE* and Kreps and Wilson (1982) on sequential equilibrium.

3 The Recommendation Subgame

3.1 Equilibrium Refinements and Simplifications

We begin with studying a recommendation subgame for a fixed price profile to understand how the expert's incentive to make recommendations depends on the price and his anticipation of the consumer's acquisition strategy, and conversely, how the consumer's acquisition strategy depends on her belief about the expert's honesty. We establish some preliminary results that allow us to simplify the problem.

Relevant Price Range. Without loss of generality, we restrict the expert's price domain to a compact, payoff-relevant set:

$$\mathbf{p} \in \mathcal{P} \equiv [c_L, v_L] \times [c_H, v_H].$$

Formally, for any price $\mathbf{p} \in \mathbb{R}_+^2 \setminus \mathcal{P}$, there exists $\mathbf{p}' \in \mathcal{P}$ such that the expert's continuation payoff under any equilibrium following \mathbf{p}' is weakly greater than his continuation payoff under any equilibrium following \mathbf{p} . If the two payoffs coincide, then the consumer's equilibrium surplus following \mathbf{p}' is weakly greater than under \mathbf{p} .

This type of price-range restriction is standard in the credence-good literature, following a similar logic to that in Fong (2005) and Kim et al. (2025). The key intuition is any treatment with price $p \notin [c_L, v_L] \cup [c_H, v_H]$ won't be traded in equilibrium. Therefore, expanding the expert's choice set outside the range benefits neither the expert nor the consumer. We prove this argument in the Appendix.

Weakly Dominated Strategies. As specified in our equilibrium concept, we focus on strategies that are not weakly dominated. Applying this refinement allows us to narrow the set of possible equilibrium actions and focus on the central agency problem. Lemma 1 formalizes this by identifying two key dominated strategies.

Lemma 1. *Suppose that $\mathbf{p} \in \mathcal{P}$. In the recommendation subgame $\Gamma(\mathbf{p})$:*

1. *Recommending an L-treatment for an H-problem is weakly dominated for the expert.*
2. *Rejecting an L-treatment is weakly dominated for the consumer.*

Proof. (1) If the expert recommends L when $\theta = H$ and the consumer accepts, his payoff is $p_L - c_H$. Under $\mathbf{p} \in \mathcal{P}$ we have $v_L < c_H$, hence $p_L - c_H \leq v_L - c_H < 0$. If the consumer rejects, the payoff is 0. Therefore recommending L when $\theta = H$ yields a payoff ≤ 0 in every contingency and is strictly worse than 0 when accepted; it is weakly dominated by recommending no treatment.

(2) If the consumer accepts an L -treatment, her payoff is $v_H - p_L$ when $\theta = H$ and $v_L - p_L$ when $\theta = L$. Since $\mathbf{p} \in \mathcal{P}$ implies $p_L \leq v_L$, we have $v_L - p_L \geq 0$, and trivially $v_H - p_L > 0$. Rejecting yields 0. Thus acceptance weakly dominates rejection. \square

The preliminary results in Lemma 1 allow us to simplify the analysis and focus on the central strategic interaction of the model. This interaction involves two key decisions: (i) the expert's choice of whether to misrepresent an L problem by recommending an H -treatment, and (ii) the consumer's response to such a recommendation, which includes her choice of an optimal information structure and her subsequent acceptance decision. Henceforth, we refer to the act of recommending an H -treatment for an L problem as *cheating*, while an expert who recommends a θ -treatment for a θ problem is referred to as *honest*.

3.2 The Core Strategic Problem

We now recast each player's problem within the subgame.

The Consumer's Problem. By Lemma 1, accepting any recommendation of the L treatment is weakly dominant for the consumer; consequently, she never acquires information in that case. Upon receiving a recommendation of the H treatment, she updates her belief about the problem's severity. Given the expert's strategy β , her posterior that $\theta = H$ is:

$$\mu(\beta) = \frac{\mu_0}{\mu_0 + (1 - \mu_0)\beta}.$$

With this updated belief $\mu = \mu(\beta)$, she then chooses an information structure $\phi = (\phi_H, \phi_L)$ to maximize her expected utility, balancing the benefit of making an informed decision against the cost of acquiring information.

The expected gross benefit of choosing a signal ϕ , denoted $W_0(\phi; \mu, \mathbf{p})$,

$$W_0(\phi; \mu, \mathbf{p}) = \mu\phi_H(v_H - p_H) + (1 - \mu)\phi_L(v_L - p_H).$$

The first term is the expected payoff when the state is truly H : with posterior probability μ , she accepts with probability ϕ_H and gains $v_H - p_H$. The second term is the expected payoff

when the state is L (which, given the H -recommendation, implies the expert cheated): with posterior probability $1 - \mu$, she accepts with probability ϕ_L and receives $v_L - p_H$.

The consumer's problem is therefore to choose the acceptance probabilities that maximize this gross payoff minus the cost of information:

$$\max_{\phi \in [0,1]^2} \{W_0(\phi; \mu, \mathbf{p}) - k C(\phi)\}.$$

The solution to this problem, $\phi(\mu, \mathbf{p}; k)$, is the consumer's best-response learning strategy.

When the expert recommends the high treatment, the consumer faces a simple trade-off: pay for information now to avoid a costly mistake, or skip learning and either accept or reject based on her current belief. Intuitively, If the price looks like a great deal given that belief, she doesn't bother with costly learning—she simply accepts. If the price is too high, she reaches the opposite conclusion and rejects without checking. Only when the price sits in a middle range—close enough to be tempting but high enough to be risky—does information have value. The next lemma formalizes this intuition by showing that there are two price thresholds: below the lower one, accept without learning; above the upper one, reject without learning; and in between, it is optimal to acquire information.

Lemma 2 (Consumer's Learning Thresholds). *Let μ be the consumer's belief that $\theta = H$. There exist unique price thresholds $\underline{p}(\mu) \leq \bar{p}(\mu)$ such that consumer learns if and only if*

$$p_H \in (\underline{p}(\mu), \bar{p}(\mu))$$

1. If $p_H \geq \bar{p}(\mu)$, the consumer acquires no information and rejecting is optimal
2. If $p_H \leq \underline{p}(\mu)$, the consumer acquires no information and accepting is optimal.
3. Moreover, $E[v | \mu] \in (\underline{p}(\mu), \bar{p}(\mu))$.
4. If $\underline{p}(\mu) = \bar{p}(\mu) = E[v|\mu]$ the learning interval is empty.

Figure 3 provides a graphical representation of the consumer's best-response regions which depends on the consumer's belief (μ , on the horizontal axis) and the price of the high-cost treatment (p_H , on the vertical axis). The diagonal dashed line represents the consumer's posterior expected value

$$E[v | \mu] = \mu v_H + (1 - \mu)v_L$$

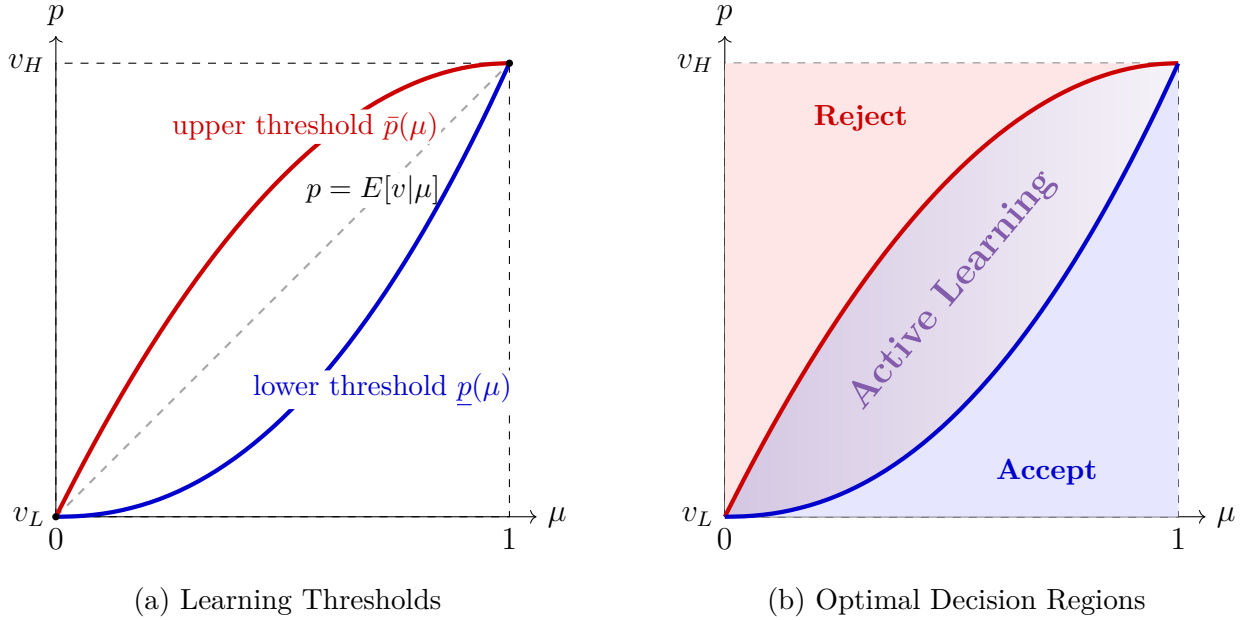


Figure 3: Consumer's Best Response

before acquiring information. The two curves above and below the dashed line, respectively, delineates three distinct regions of consumer behavior.

Above the red curve (the one above the dashed line) the price is too high relative to the consumer's belief to justify the treatment, and the consumer rejects without investigation. Below the blue curve (the other below the dashed line), the price is low enough relative to the consumer's belief that the consumer accepts without investigation. It is only in the intermediate region enclosed by these two curves that the consumer finds it optimal to acquire costly information. The intensity of the blue heatmap in this region represents the amount of information the consumer acquires, showing that learning is most intense when the price p_H is close to the consumer's no-information expected value (the dashed line).

Figure 4 illustrates the consumer's best response under the two extreme cases for the cost of information. As information becomes cheaper ($k \downarrow$), the intermediate (learning) region expands; as it becomes more costly ($k \uparrow$), the region shrinks. In the limit $k = 0$ (costless verification), the two threshold curves move apart so that the intermediate region covers the entire (μ, p_H) panel: the consumer learns for every pair (μ, p_H) and then accepts or rejects perfectly based on the realized state. In the opposite limit $k = \infty$ (no learning), the two thresholds coincide at the curve $p_H = E[v | \mu]$ (the dashed line), so the intermediate region collapses to that line and

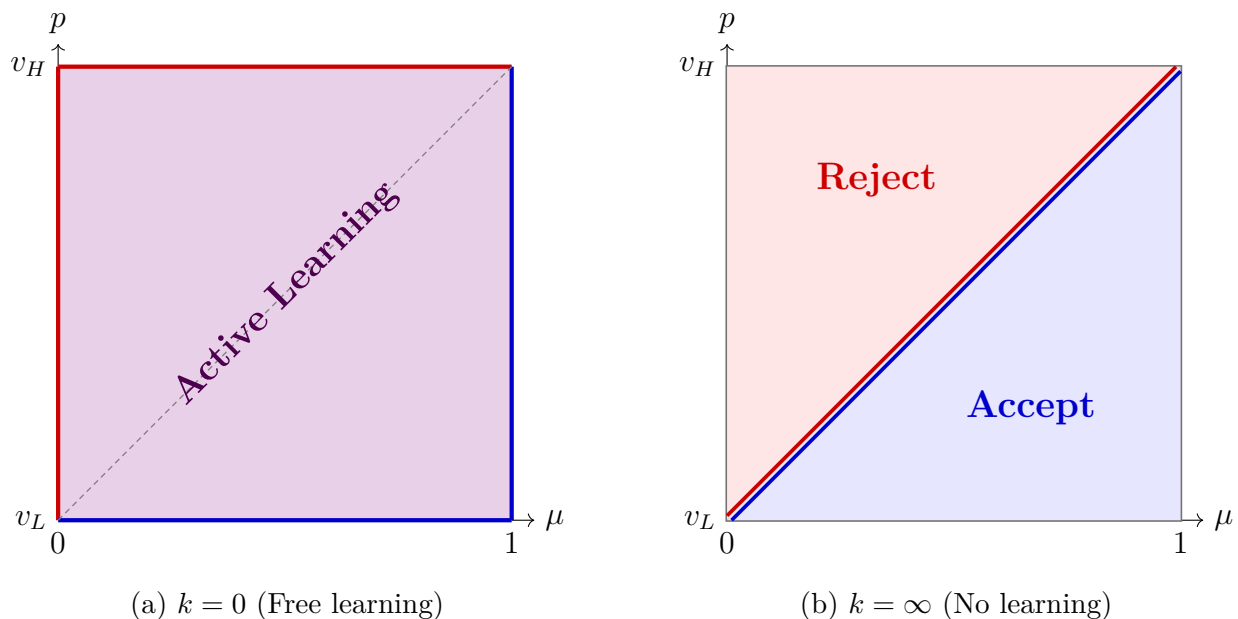


Figure 4: A comparison of consumer behavior under different learning cost scenarios.

the consumer never acquires information.

The Expert’s Problem. When the true state is an L problem, the expert chooses a cheating probability $\beta \in [0, 1]$ to maximize his expected profit, anticipating the consumer’s learning strategy ϕ . He faces a trade-off. With probability $1 - \beta$, he can be *honest*, recommend the L -treatment, and earn a certain profit of $p_L - c_L$. Alternatively, with probability β , he can *cheat* by recommending the H -treatment. In this case, his profit is risky: he earns $p_H - c_L$ only if the consumer’s signal yields an “accept” recommendation, which occurs with probability ϕ_L . The expert’s problem is to choose β to maximize a weighted average of these two outcomes. Formally, his maximization problem is:

$$\max_{\beta \in [0, 1]} (1 - \beta)(p_L - c_L) + \beta\phi_L(p_H - c_L)$$

The solution is a best-response correspondence, $\beta(\phi, \mathbf{p})$, which depends on the relative profitability of being honest versus cheating. The expert’s optimal strategy is a corner solution unless he is exactly indifferent:

- If $p_L - c_L > \phi_L(p_H - c_L)$, the profit from honesty is strictly greater. The expert is always honest, so $\beta(\phi, \mathbf{p}) = 0$.

- If $p_L - c_L < \phi_L(p_H - c_L)$, the expected profit from cheating is strictly greater. The expert always cheats, so $\beta(\phi, \mathbf{p}) = 1$.
- If $p_L - c_L = \phi_L(p_H - c_L)$, the expert is indifferent between being honest and cheating. He is willing to randomize, so $\beta(\phi, \mathbf{p}) \in [0, 1]$.

This characterization highlights how the expert's incentive to cheat is increasing in the price of the major treatment (p_H) and the consumer's propensity to accept a recommendation for it when her problem is minor (ϕ_L), and decreasing in the price of the minor treatment (p_L).

Equilibrium of Recommendation Subgame. For any given price menu \mathbf{p} and information cost k , an equilibrium in the recommendation subgame consists of a strategy profile (β^*, ϕ^*) and a posterior belief μ^* that satisfy two conditions. First, each player's strategy must be a best response to the other: β^* solves the expert's problem given ϕ^* , and ϕ^* solves the consumer's problem given her belief μ^* . Second, the consumer's belief must be consistent with the expert's strategy via Bayes' rule: $\mu^* = \frac{\mu_0}{\mu_0 + (1 - \mu_0)\beta^*}$.

The structure of our model, particularly our focus on the relevant price range and weakly dominant strategies, yields two convenient features that greatly simplify this equilibrium analysis. First, since the expert always recommends H -treatment when the state is H , an H -recommendation occurs with positive probability on the equilibrium path. This ensures the posterior μ^* is always well-defined by Bayes' rule. Second, accepting an L -recommendation is a weakly dominant strategy for the consumer. This allows us to remain agnostic about her off-equilibrium beliefs following such a recommendation.

In the next section, we solve for this equilibrium at the extreme values of k .

4 Benchmark Analysis

Before presenting the main result, we characterize the equilibrium outcomes in two benchmark cases corresponding to extreme values of the information cost parameter, k . These benchmarks establish the baseline against which my central finding—that a small amount of friction can be welfare-improving—is compared.

4.1 Free Learning

First, consider the idealized case where the consumer can acquire perfect information at no cost, i.e., $k = 0$.

Proposition 1. *When information is costless ($k = 0$), there exists a unique separating equilibrium outcome. In this equilibrium, the expert charges $p_L^* = v_L$ and $p_H^* = v_H$ and makes truthful recommendations. The consumer acquires precise information enough to incentive the expert's honest recommendation; and accepts all on-path recommendations. Trade is efficient, and the consumer's surplus is zero.*

The intuition behind this result is straightforward: costless verification eliminates the expert's credibility problem. Since the consumer can perfectly and freely check any recommendation, an attempt by the expert to cheat—by recommending an H -treatment for an L -type problem—would be futile, as it would be immediately detected and rejected. Anticipating this, the expert has no incentive to be dishonest, ensuring that recommendations are truthful in equilibrium.

At the price $p = \mathbf{v}$, the recommendation subgame admits multiple continuations. In all of them the expert is truthful and the consumer accepts any L recommendation, but after an H recommendation she can choose any acceptance probability up to a cutoff that just makes the expert indifferent. At one extreme, she never acquires information and rejects all H recommendations and the market for H effectively collapses; at the other, she acquires perfect information accepts all H recommendations, granting expert the highest payoff equal to the efficient gain from trade— $E[v_\theta - c_\theta]$. In an equilibrium of the full game, the latter continuation is uniquely sustainable.

By setting prices just below the consumer's valuation, such as $\mathbf{p}^\varepsilon := \mathbf{v} - \varepsilon$, where $\varepsilon = (\varepsilon, \varepsilon)$, the expert gives the consumer a strict incentive to verify any recommendation. This is because successfully verifying a truthful recommendation yields a positive surplus of ε , which is better than the zero surplus from rejecting the offer. This off-path price strategy serves as a powerful credible threat that stabilizes the full-surplus-extraction outcome. If the consumer were to follow an equilibrium strategy at the price \mathbf{v} that did not yield the full surplus to the expert, the expert could profitably deviate to the price $\mathbf{p}^\varepsilon := \mathbf{v} - \varepsilon$. At that price \mathbf{p}^ε , the consumer's best response is to verify, which in turn forces the expert to be truthful. This deviation would guarantee the expert a profit of $E[v_\theta - c_\theta] - \varepsilon$. Since this profit can be made arbitrarily close

to the full surplus (as $\varepsilon \downarrow 0$), any equilibrium other than the expert's most preferred one is unsustainable. A formal proof is provided in the Appendix.

Remark 1. *The consumer's ability to costlessly and perfectly verify the expert's claim completely resolves the expert's credibility problem. This allows the expert to use the consumer's verification threat as a commitment device to his own truthfulness. Because the consumer cannot commit not to acquire information (and thereby rejecting high price) the expert is able to set the highest price that extract all the gains from trade, leaving the consumer with zero surplus. While this outcome is efficient, the consumer derives no benefit from her access to free information.*

4.2 No Learning

We next consider the regime in which information cost is so large that the consumer does not learn—more precisely, she does not learn enough to affect the expert's incentives in the recommendation subgame. In this no-learning benchmark, external verification is effectively unavailable and recommendations reduce to cheap talk.⁴ This benchmark establishes the outcome in a standard credence-good setting studied by Fong (2005) where external verification is not an option and recommendations are cheap talk.

Proposition 2. *When information is prohibitively costly ($k = \infty$), there exists a unique separating equilibrium outcome. In this outcome, the expert makes truthful recommendations at prices $p_L^* = v_L$ and $p_H^* = v_H$. To deter cheating, the consumer rejects the H -recommendation with a positive probability. The outcome is inefficient, and the consumer's surplus is zero.*

The intuition here is that without an external verification mechanism, the expert's recommendations are mere cheap talk. To make an expensive H -recommendation credible, the consumer must hold a credible threat of punishment. By randomly rejecting the H -recommendation with a specific probability, the consumer makes the expert indifferent between honestly recommending H for an H -problem and cheating by recommending H for an L -problem. This threat of an inefficient market where H -type problem is undertreated is what disciplines the expert and sustains truthfulness.

This enforcement directly leads to zero consumer surplus. The expert, knowing the consumer's rejection strategy, sets the highest possible prices $\mathbf{p}^* = \mathbf{v} = (v_L, v_H)$. The consumer's

⁴Formally, for any fixed primitives and cost function C , there is a (model-dependent) threshold above which the consumer's best response is uninformative (or yields the same equilibrium behavior as without learning).

rejection probability is precisely calibrated to deter the expert, but this leaves the consumer with an expected payoff of zero. A formal proof is provided in Appendix.

Remark 2. *This result shows the inefficiency of establishing credibility without external information. The market relies on a costly and socially wasteful commitment device: the consumer’s threat to reject a recommendation that is, in equilibrium, always honest. Both the expert (who risks rejection) and the consumer (who gets zero surplus) would be better off if a more efficient credibility mechanism were available. This sets the stage for our main result, which shows that a small, positive cost of information provides exactly such a mechanism.*

5 Main Result

5.1 Intermediate Learning Cost

The benchmark analysis reveals a stark outcome for the consumer: whether information is free or prohibitively costly, her option to learn does not help the consumer at all. In the free-learning case, her ability to verify claims serves only to enhance the expert’s credibility, allowing him to extract all gains from trade. This raises a natural question: what happens in the intermediate case, where information is neither free nor prohibitively costly?

Our main result, presented in Theorem 1, shows that under the presence of moderate information cost, both the expert and the consumer strictly benefits over the no-learning benchmark. In particular, the consumer strictly benefits over both benchmarks.

Theorem 1. *There exists $\bar{k} > 0$ such that for all information costs $k \in (0, \bar{k})$, every equilibrium outcome $(\Pi(k), U(k))$ satisfies the following properties:*

- (i) *The consumer learns and benefits from learning: $U(k) > 0$.*
- (ii) *The expert benefits from the consumer’s learning: $\Pi(k) > \Pi(\infty)$.*
- (iii) *The expert cheats with positive probability and gets rejected with positive probability.*

The intuition behind Theorem 1 centers on the expert’s use of a price cut ($p_H^* < v_H$) as a subsidy to incentivize consumer learning. This act of subsidized learning and cheating is noteworthy as it inverts the standard logic of information rents. Whereas a classic information rent involves an uninformed party paying an informed party to reveal information, here the

Table 1: Comparison of Equilibrium Outcomes by Information Cost

Cost of information	Credibility built on	Equilibrium payoff (expert, consumer)
$k = 0$ (Free Learning)	(almost-) Perfect Verification	$(\Pi(0), 0)$
$k = \infty$ (No Learning)	Random Rejection	$(\Pi(\infty), 0)$
$k \in (0, \bar{k})$	Subsidized Learning	$(\Pi(k) > \Pi(\infty), U(k) > 0)$

Notes: $\Pi(0) = E[v - c]$ is the efficient social surplus, $\Pi(\infty)$ the no-learning benchmark with $\Pi(0) > \Pi(\infty)$.

informed expert pays the uninformed consumer to *learn*. The following paragraphs unpack how this benefits both players. As we show, the expert’s incentive to cheat is a necessary evil, as it is precisely what makes the consumer’s costly verification a valuable action.

Why the Expert Benefits. The expert’s gain arises from finding a more efficient method to establish credibility. Without consumer learning ($k = \infty$), the expert’s profits are eroded by a strategy of random rejection, which discards even honest recommendations. However, when learning is possible ($k \in (0, \bar{k})$), the expert can strategically lower the price to incentivize the consumer to verify his advice—an action the consumer would forgo at the higher price of $\mathbf{p}^* = \mathbf{v}$. While this price cut reduces the expert’s margin per sale, the trade-off is profitable: the consumer’s ability to verify builds credibility, which in turn drives higher demand and boosts the expert’s overall profit.

Why the Consumer Benefits. The consumer’s surplus arises from a powerful form of two-sided commitment that does not exist in the benchmarks. Her commitment power comes from her ability to refuse to learn—more precisely, not learn precise information—if the price is too high, a threat that is credible only because learning is costly ($k > 0$). This contrasts sharply with the free-learning ($k = 0$) case, where she cannot commit *not* to do so and thus has no leverage. Faced with this credible threat, the expert is forced to offer a price cut that not only covers her verification costs but also leaves her with a positive surplus. The flexibility of her learning technology is crucial; it prevents the expert from shaving this surplus down to zero, as she can always find a profitable deviation by choosing a different information structure if he tries to make her indifferent. This creates the central paradox of information cost: the consumer is strictly better off with a small cost of information than with none at all, because

the friction is what grants her the commitment power to secure a share of the surplus.

Cheating as Intrinsic Feature. The entire surplus-generating economics we have so far discussed is sustained by the expert’s potential for opportunism: the consumer is only willing to engage in costly verification because she fears being sold an unnecessary treatment. If the expert were known to be perfectly honest, her incentive to learn would disappear. This, in turn, would dissolve the very foundation of the expert’s more efficient commitment device. Therefore, the expert’s potential for cheating is a necessary evil that underpins the value of verification and allows for the mutually beneficial outcome to exist.

Having established the intuition, we now sketch the proof of Theorem 1. The formal proof, detailed in the appendix, establishes the theorem’s three properties by proving two main propositions. First, for an intermediate cost of information (k), the expert is strictly better off than in the no-learning benchmark, and his equilibrium strategy must involve cheating (Proposition 3). Second, this necessarily implies the consumer captures a strictly positive surplus (Proposition 4).

The argument for the expert’s improved profit relies on continuity. We establish that the expert’s worst-case equilibrium profit function, $\Pi(k)$, is lower semi-continuous at $k = 0$ (i.e., it cannot jump down). A key step shows that as the cost k vanishes, the consumer’s optimal strategy converges to perfect verification when the state is $\theta = H$ ($\phi_H^* \uparrow 1$ as $k \downarrow 0$). This ensures that for any k below some threshold $\bar{k} > 0$, the expert’s profit remains close to the efficient, free-learning benchmark, which is strictly greater than his profit in the no-learning case.

We then show that positive cheating is a necessary feature of any such equilibrium. The logic is straightforward: if the expert were always honest, a consumer with $k > 0$ would never pay to verify his claims. This lack of oversight would give the expert a profitable incentive to deviate and cheat, unraveling any supposedly honest equilibrium.

The argument for the consumer’s positive surplus proceeds by contradiction. We assume the consumer’s surplus is zero and show this leads to one of two impossible outcomes. The consumer’s strategy cannot be uninformative (e.g., random rejection). If it were, the outcome would revert to the no-learning benchmark, yielding a low profit for the expert. This contradicts our first result that the expert must be strictly better off. The strategy cannot be informative. If the strategy were informative but yielded zero surplus, the consumer would have a profitable

deviation. This step relies on the strict quasi-concavity of the consumer’s objective function. This property guarantees that a consumer held to zero surplus can always choose a slightly less informative (and thus cheaper) signal to secure a positive payoff, contradicting the optimality of her strategy. Since a zero-surplus outcome is impossible, the consumer’s equilibrium payoff must be strictly positive. Together, these steps formalize the intuition that a small information cost can create a mutually beneficial outcome.

5.2 Numerical Illustration

To build intuition for the equilibrium trade-offs, this section numerically illustrates the model’s predictions for a baseline set of parameters.⁵ We first analyze the equilibrium mechanism for a fixed information cost before turning to the comparative statics.

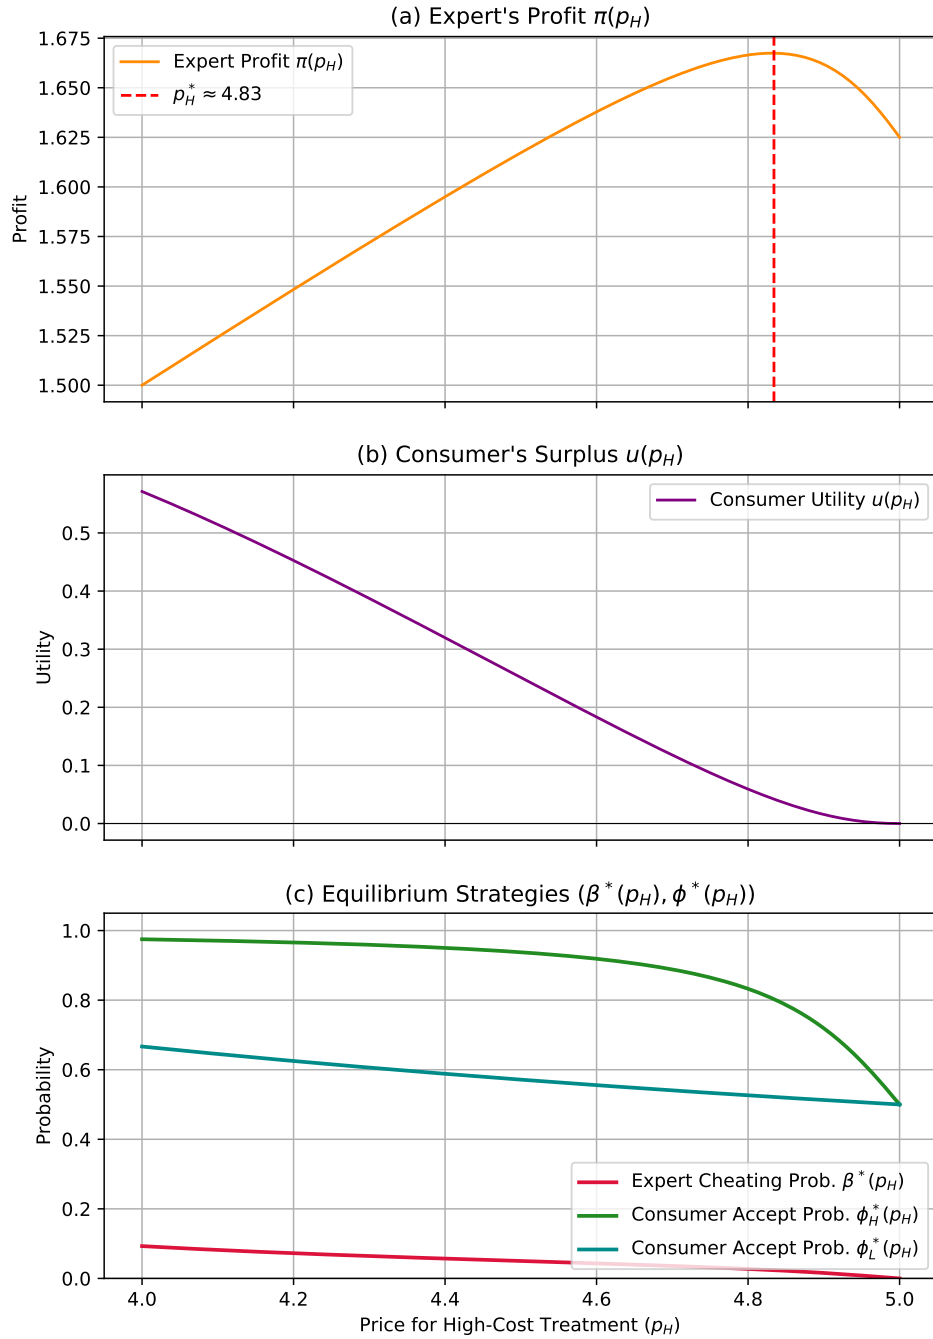
Equilibrium at fixed k . Figure 5 illustrates the equilibrium of the recommendation subgame for a fixed information cost ($k = 0.1$) as the expert’s price for the H -treatment p_H varies (with $p_L = v_L$ fixed).

The bottom panel shows the players’ equilibrium strategies. For any price below the consumer’s valuation ($p_H < v_H = 5.0$), the expert employs a semi-separating strategy—cheating with a positive probability ($\beta^*(p_H) > 0$). To counter this, the consumer engages in costly learning, which is reflected in the separation between her acceptance probabilities conditional on the state ($\phi_H^* > \phi_L^*$). As the price approaches the consumer’s full valuation ($p_H \rightarrow v_H$), the expert’s incentive to cheat diminishes, and the consumer’s costly verification ceases as the gap $\phi_H^* - \phi_L^*$ converges to zero. The top two panels show the resulting payoffs. The expert’s profit is maximized at an interior price of $p_H^* \approx 4.82$, where he balances the benefit of a higher price against the cost of inducing more precise learning from the consumer.

Comparative Statics. Figure 6 presents the central comparative statics of the model, plotting the equilibrium outcomes as information cost becomes cheaper (as k decreases from 1 to 0). The results are strikingly consistent with our theoretical findings.


⁵Parameters are set to: $v_H = 5.0, v_L = 3.0, c_H = 4.0, c_L = 1.0, \mu_0 = 0.25$. The cost function is the LLR cost of Pomatto et al. (2023) in Example 1 with $\beta_{HL} = \beta_{LH} = 0.5$. These values satisfy the core assumptions of the model, including the lemon condition.

Figure 5: Numerical Illustration of Equilibrium at $k = 0.1$



Notes: The figure shows the equilibrium strategies and payoffs as a function of p_H , holding other parameters fixed. Panel (a) shows the expert's profit, Panel (b) shows the consumer's surplus, and Panel (c) shows the equilibrium choice of ϕ_H (acceptance probability) and β (cheating probability).

Figure 6: Comparative Statics for $k \in (0, 1]$



Notes: The figure plots equilibrium outcomes as a function of the consumer's unit cost of information, k , changes from $k = 1$ to $k = 0.001$. Panel (a) shows the expert's optimal price, Panel (b) shows the expert's equilibrium profit, and Panel (c) shows the consumer's equilibrium surplus as a function of k .

- *Optimal Price (Top Panel (a))*: The expert’s optimal price, p_H^* , is non-monotone in the cost of information. For high costs (e.g., $k = 1.0$), the expert sets the price at the consumer’s valuation, $p_H^* = v_H = 5.0$. As k falls, the expert strategically lowers the price to subsidize the consumer’s learning, before raising it again as information becomes nearly free ($k \rightarrow 0$).
- *Expert’s Profit (Middle Panel (b))*: The expert’s profit is monotonically increasing as information gets cheaper. This highlights his benefit from shifting from the inefficient random-rejection benchmark (high k) to the more efficient subsidized-learning regime (low k), where credibility is established at a lower cost.
- *Consumer’s Surplus (Bottom Panel (c))*: Most notably, the consumer’s surplus is non-monotonic, confirming the “paradox of friction” established in Theorem 1. The consumer is worst off when information is either prohibitively expensive (high k) or free ($k = 0$), capturing zero surplus in both extremes. She achieves a strictly positive surplus only for an intermediate range of information costs.

Remark 3. *A feature of the consumer’s surplus is that it is not single-peaked in general. This complex pattern arises because a change in k affects the consumer’s welfare through two distinct channels: a **price effect** and a **learning effect**. As k increases, the expert may lower the price (a positive price effect for the consumer), but the direct cost of learning also increases (a learning effect). The final shape of the surplus curve depends on which of these two effects dominates.*

6 Conclusion

This paper analyzes how a consumer’s ability to investigate an expert’s recommendation shapes outcomes in a credence-good market. Our analysis reveals that the consumer is left with no net benefit when verification is either free—allowing the expert to capture all gains from trade—or prohibitively expensive, where credibility is maintained only through inefficient random rejection. In contrast, a modest cost of inspection empowers the consumer. It creates a credible threat to forgo investigation and reject if the price is too high, compelling the expert to offer a lower price and thereby allowing the consumer to capture a positive payoff. This mutually beneficial arrangement hinges on the expert’s potential for opportunism; without a

risk of being misled, the consumer would have no incentive to conduct a costly verification, which in turn would unravel the surplus-generating equilibrium.

A key policy implication of our model is that the proliferation of AI-powered tools, such as Large Language Models, can generate mutually beneficial outcomes for both consumers and experts. Our model provides a new lens to interpret these technologies: they represent a low-cost regime where our central “paradox of friction” operates. AI-powered tool has their own limitations—perfect verification is not truly costless. Nevertheless, these tools enable a more efficient equilibrium where the expert establishes credibility by subsidizing consumer learning, a path that is preferable for both parties over the inefficient benchmark.

The results are driven by a few key assumptions that offer avenues for future research. Our model features a monopolistic expert, homogeneous consumers, and a specific timing protocol. One promising extension is to introduce competition, exploring how the strategic subsidization of consumer inquiry changes when consumers can seek a second opinion from a rival expert. Another avenue is to relax the assumption of a perfectly informed expert by introducing diagnostic costs. If the expert must incur an expense for an accurate diagnosis, a new strategic tension emerges where an expert might recommend a default treatment to avoid diagnostic fees. Analyzing how the consumer’s own investigation interacts with the expert’s diagnostic effort would provide a more comprehensive understanding of these markets.

References

- Chatterjee, K., M. Dong, and T. Hoshino (2025). Bargaining and information acquisition. *American Economic Journal: Microeconomics* 17(3), 75–102.
- Cheng, X. and Y. Kim (2025). On the monotonicity of information costs.
- Darby, M. R. and E. Karni (1973). Free competition and the optimal amount of fraud. *The Journal of law and economics* 16(1), 67–88.
- Fabbri, F. (2025). Attention holdup. Technical report, Working Paper.
- Fong, Y.-f. (2005). When do experts cheat and whom do they target? *RAND Journal of Economics*, 113–130.

- Hyndman, K. and S. Ozerturk (2011). Consumer information in a market for expert services. *Journal of Economic Behavior & Organization* 80(3), 628–640.
- Jost, P.-J., S. Reik, and A. Ressi (2021). The information paradox in a monopolist’s credence goods market. *International Journal of Industrial Organization* 75, 102694.
- Kim, D., F. Li, and T. Liu (2025). Consumer information in credence-good market. *Available at SSRN 5161507*.
- Kreps, D. M. and R. Wilson (1982). Sequential equilibria. *Econometrica: Journal of the Econometric Society*, 863–894.
- Mailath, G. J. (2018). *Modeling strategic behavior: a graduate introduction to game theory and mechanism design*, Volume 6. World Scientific.
- Mensch, J. (2022). Screening inattentive buyers. *American Economic Review* 112(6), 1949–1984.
- Milgrom, P. and I. Segal (2002). Envelope theorems for arbitrary choice sets. *Econometrica* 70(2), 583–601.
- Milgrom, P. and C. Shannon (1994). Monotone comparative statics. *Econometrica: Journal of the Econometric Society*, 157–180.
- Pitchik, C. and A. Schotter (1987). Honesty in a model of strategic information transmission. *The American Economic Review* 77(5), 1032–1036.
- Pomatto, L., P. Strack, and O. Tamuz (2023). The cost of information: The case of constant marginal costs. *American Economic Review* 113(5), 1360–1393.
- Ravid, D. (2020). Ultimatum bargaining with rational inattention. *American Economic Review* 110(9), 2948–2963.
- Ravid, D., A.-K. Roesler, and B. Szentes (2022). Learning before trading: on the inefficiency of ignoring free information. *Journal of Political Economy* 130(2), 346–387.
- Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics* 50(3), 665–690.

Taylor, C. R. (1995). The economics of breakdowns, checkups, and cures. *Journal of Political Economy* 103(1), 53–74.

Wolinsky, A. (1993). Competition in a market for informed experts' services. *The RAND Journal of Economics*, 380–398.

Woodford, M. (2009). Information-constrained state-dependent pricing. *Journal of Monetary Economics* 56, S100–S124.

A Appendix: Proofs

A.1 Omitted Proofs in section 4

Lemma 3 (Optimal Learning Thresholds). *Fix $\beta \in [0, 1]$ and let $\mu \equiv \mu(\beta) = \frac{\mu_0}{\mu_0 + (1 - \mu_0)\beta}$ be the consumer's belief that $\theta = H$ after an H -recommendation, and write $v(\mu) \equiv E[v \mid \mu] = \mu v_H + (1 - \mu)v_L$. There exist unique price thresholds $\bar{p}(\beta) \geq \underline{p}(\beta)$ such that:*

1. *If $p_H \geq \bar{p}(\beta)$, the consumer acquires no information and rejecting is optimal, so $V = 0$.*
2. *If $p_H \leq \underline{p}(\beta)$, the consumer acquires no information and accepting is optimal, so $V = v(\mu) - p_H$.*
3. *If $p_H \in (\underline{p}(\beta), \bar{p}(\beta))$, the consumer acquires informative signal ϕ^* .*

Moreover,

$$\underline{p}(\beta) \leq v(\mu(\beta)) \leq \bar{p}(\beta).$$

The interval $(\underline{p}(\beta), \bar{p}(\beta))$ may be empty.

Proof. Fix $\beta \in [0, 1]$ so that $\mu(\beta) = \frac{\mu_0}{\mu_0 + (1 - \mu_0)\beta}$. Write $p \equiv p_H$ and $v(\beta) \equiv E[v \mid \mu(\beta)]$. Following H -recommendation, the consumer chooses one of three actions: *Reject* (R), *Accept* (A), or *Learn* (L). These yield

$$R(p) = 0, \quad A(p) = v(\beta) - p, \quad L(p) = \max_{\phi \in \Phi} \{W_0(\phi, \mu, p) - kC(\phi)\},$$

where $W_0(\phi, \mu, p) = \mu\phi_H(v_H - p) + (1 - \mu)\phi_L(v_L - p)$. Clearly,

$$L(p) \geq \max\{A(p), 0\}. \tag{1}$$

For each p , $W_0(\phi, \mu, p) - kC(\phi)$ is linear and nonincreasing in p . Hence, $L(\cdot)$ is continuous, convex, and nonincreasing in p . We also have boundary values

$$L(v_H) = 0 \quad (\text{since } W_0(\phi, \mu, v_H) \leq 0), \quad L(v_L) = A(v_L) = v(\beta) - v_L > 0.$$

Define

$$\bar{p}(\beta) \equiv \inf\{p : L(p) = 0\} \leq v_H, \quad \underline{p}(\beta) \equiv \sup\{p : L(p) = A(p)\} \geq v_L.$$

By construction and monotonicity, $L(p) = 0$ for all $p \geq \bar{p}(\beta)$, while $L(p) = A(p) \geq 0$ for all $p \leq \underline{p}(\beta)$. Since $A(\cdot)$ is affine with slope -1 and $R \equiv 0$,

$$L(p) = 0 = R(p) \quad \text{for } p \geq \bar{p}(\beta), \quad L(p) = A(p) \quad \text{for } p \leq \underline{p}(\beta),$$

establishing (1)–(2). For (3), on $p \in (\underline{p}(\beta), \bar{p}(\beta))$ we have $L(p) > 0$ and $L(p) > A(p)$, so the consumer strictly prefers learning and achieves $W_0(\phi^*; \mu(\beta), \mathbf{p}) - kC(\phi^*)$ for some informative ϕ^* .

To see $\underline{p}(\beta) \leq \bar{p}(\beta)$, suppose not. If $\underline{p}(\beta) > \bar{p}(\beta)$, then for any $p \in (\bar{p}(\beta), \underline{p}(\beta))$, $V(p) = 0 = A(p)$, so $A(p)$ would be constant over an interval, contradicting its slope -1 .

Finally, prove $\underline{p}(\beta) \leq v(\beta) \leq \bar{p}(\beta)$. If $p < v(\beta)$, then $A(p) > 0$, and by (1) $L(p) \geq A(p) > 0$, hence $p < \bar{p}(\beta)$. Therefore, $v(\beta) \leq \bar{p}(\beta)$. If $p > v(\beta)$, then $A(p) < 0$, so $L(p) = A(p)$ cannot hold because $L(p) \geq 0$ by (1); thus $p > \underline{p}(\beta)$. Therefore, $v(\beta) \geq \underline{p}(\beta)$. This yields the stated sandwich inequality. \square

Lemma 4 (Properties of the Value Function). *The consumer's value function $V(\mu, \mathbf{p}; k)$ is non-decreasing in the belief μ , and non-increasing in the price p_H .*

Proof of Lemma 4. Fix a price vector $p = (p_L, p_H) \in \mathcal{P}$. Write

$$V(\mu, \mathbf{p}; k) = \max_{\phi \in \Phi} \{ W_0(\phi; \mu, \mathbf{p}) - kC(\phi) \}, \quad W_0(\phi; \mu, \mathbf{p}) = \mu \phi_H (v_H - p_H) + (1 - \mu) \phi_L (v_L - p_H),$$

where by revelation principle we take $\Phi = [0, 1]^2$ and $\phi = (\phi_H, \phi_L)$ are accept probabilities in states H, L after an H -recommendation. The cost C satisfies Assumption 1.

For each ϕ , $W_0(\phi; \mu, \mathbf{p})$ is affine in (μ, p_H) . Let $\phi^* \in \arg \max_{\phi} \{ W_0(\phi; \mu, \mathbf{p}) - kC(\phi) \}$. $C(\phi^*)$ is finite at ϕ^* . Hence, $V(\mu, \mathbf{p}) = W_0(\phi^*; \mu, \mathbf{p}) - kC(\phi^*)$ has well-defined partial with respect to both μ and p_H . Thus Envelope Theorem (see [Milgrom and Segal \(2002\)](#)) applies to V with respect to both μ and p_H . Then

$$\partial_{\mu} W_0(\phi^*; \mu, \mathbf{p}) = \phi_H^* (v_H - p_H) - \phi_L^* (v_L - p_H), \quad \partial_{p_H} W_0(\phi^*; \mu, \mathbf{p}) = -\mu \phi_H^* - (1 - \mu) \phi_L^*,$$

and the $-kC(\phi^*)$ term drops out since it does not depend on μ or p_H . On the price domain we have $v_H - p_H \geq 0$ and $v_L - p_H < 0$, so

$$\partial_\mu W_0(\phi^*; \mu, \mathbf{p}) = \underbrace{\phi_H^* (v_H - p_H)}_{\geq 0} + \underbrace{\phi_L^* (p_H - v_L)}_{\geq 0} \geq 0.$$

By the envelope theorem, every one-sided derivative and every sub/supergradient of V in μ is a convex combination of such nonnegative numbers. Hence $V(\mu, \mathbf{p})$ is nondecreasing in μ . Likewise, for any maximizer ϕ^* ,

$$\partial_{p_H} W_0(\phi^*; \mu, \mathbf{p}) = -\mu \phi_H^* - (1 - \mu) \phi_L^* \leq 0,$$

with equality only if $\phi_H^* = \phi_L^* = 0$ (outright rejection). So $V(\mu, \mathbf{p})$ is nonincreasing in p_H . \square

Corollary 1 (Shape of V in p_H). *Fix $\beta \in (0, 1]$ and write $\mu \equiv \mu(\beta)$. Let $\underline{p}(\beta), \bar{p}(\beta)$ be the thresholds from Lemma 3. Then*

$$V(\mu(\beta), \mathbf{p}) = \begin{cases} v(\mu(\beta)) - p_H, & \text{if } p_H \leq \underline{p}(\beta), \\ W_0(\phi^*; \mu(\beta), \mathbf{p}) - kC(\phi^*), & \text{if } p_H \in (\underline{p}(\beta), \bar{p}(\beta)), \\ 0, & \text{if } p_H \geq \bar{p}(\beta), \end{cases}$$

where, in the intermediate price region, $\phi^* \in \arg \max_{\phi \in \Phi} \{W_0(\phi; \mu(\beta), \mathbf{p}) - kC(\phi)\}$ (by Lemma 3 this optimizer is informative there).

Moreover:

(a) Monotonicity and convexity. *As a function of p_H , $V(\mu(\beta), \mathbf{p})$ is nonincreasing and convex. Its slope lies in $[-1, 0]$ everywhere.*

(b) Envelope inside the learning region. *For any $p_H \in (\underline{p}(\beta), \bar{p}(\beta))$ at which the maximizer is unique,*

$$\frac{\partial V}{\partial p_H}(\mu(\beta), \mathbf{p}) = -[\mu(\beta) \phi_H^* + (1 - \mu(\beta)) \phi_L^*] \in (-1, 0),$$

and V is (locally) affine in p_H with that slope.

(c) Kinks at thresholds.

$$\partial_{p_H}^- V(\underline{p}(\beta)) = -1, \quad \partial_{p_H}^+ V(\bar{p}(\beta)) = 0,$$

and

$$\begin{aligned}\partial_{p_H}^+ V(\underline{p}(\beta)) &= -[\mu(\beta)\phi_H^* + (1 - \mu(\beta))\phi_L^*] \Big|_{p_H=\underline{p}(\beta)^+} \in [-1, 0), \\ \partial_{p_H}^- V(\bar{p}(\beta)) &= -[\mu(\beta)\phi_H^* + (1 - \mu(\beta))\phi_L^*] \Big|_{p_H=\bar{p}(\beta)^-} \in (-1, 0].\end{aligned}$$

Lemma 5 (1-Lipschitz in p_H). *Fix $\beta \in (0, 1]$ and write $p \equiv p_H$. Let*

$$F(\phi, p) = W_0(\phi; \mu(\beta), \mathbf{p}) - kC(\phi) = K(\phi) - q(\phi)p, \quad q(\phi) := \mu(\beta)\phi_H + (1 - \mu(\beta))\phi_L \in [0, 1],$$

and $V(\beta, p) = \sup_{\phi \in \Phi} F(\phi, p)$. Then for all p_1, p_2 ,

$$|V(\beta, p_2) - V(\beta, p_1)| \leq |p_2 - p_1|.$$

Proof. Assume $p_2 > p_1$ (the other case is symmetric). Pick $\phi_1 \in \arg \max_{\phi} F(\phi, p_1)$. Then

$$V(\beta, p_2) \geq F(\phi_1, p_2) = F(\phi_1, p_1) - q(\phi_1)(p_2 - p_1) = V(\beta, p_1) - q(\phi_1)(p_2 - p_1) \geq V(\beta, p_1) - (p_2 - p_1),$$

since $q(\phi_1) \in [0, 1]$. Moreover, for every ϕ , $F(\phi, \cdot)$ is nonincreasing in p , so $V(\beta, p_2) = \sup_{\phi} F(\phi, p_2) \leq \sup_{\phi} F(\phi, p_1) = V(\beta, p_1)$. Hence

$$-(p_2 - p_1) \leq V(\beta, p_2) - V(\beta, p_1) \leq 0,$$

which implies $|V(\beta, p_2) - V(\beta, p_1)| \leq p_2 - p_1$. □

A.2 Proof of Proposition 1

Proof of Proposition 1. The proof proceeds in two steps. At the first step, we characterize the subgame equilibrium for each price $\mathbf{p} \in \mathcal{P} \setminus \{\mathbf{v}\}$ and show that the expert's profit in the subgame is uniquely $E[p_\theta - c_\theta]$ for any interior price $p \in \mathcal{P}^\circ$ and for all other prices $\mathcal{P} \setminus \{\mathbf{v}\}$, the expert's profit is strictly smaller than $E[v_\theta - c_\theta]$ —the efficient social surplus. In the next step, we show that following price \mathbf{v} there are multiple subgame equilibria. However, the subgame equilibrium play which yields the expert his highest profit, $E[v_\theta - c_\theta]$, is a unique equilibrium outcome consistent with equilibrium play of the entire game.

First, we claim the unique subgame equilibrium is separating. Suppose for contradiction that a non-separating equilibrium exists ($\beta > 0$). A consumer receiving an H -recommendation faces a risk of a negative payoff ($v_L - p_H < 0$). Since information is costless ($k = 0$), her unique best response is to acquire a perfectly informative signal. Anticipating this, an expert observing

$\theta = L$ strictly prefers to be truthful (profit $p_L - c_L > 0$) over cheating and being rejected (profit 0). This contradicts the premise. Thus, for any $\mathbf{p} \in \mathcal{P}^\circ$, the unique subgame equilibrium is separating ($\beta^* = 0$), and on the equilibrium path, the consumer accepts all recommendations. The expert's profit is $E[p_\theta - c_\theta] := \mu_0(p_H - c_H) + (1 - \mu_0)(p_L - c_L)$.

Next, we show that other boundary prices $\mathbf{p} \in \partial\mathcal{P} \setminus \{\mathbf{v}\}$ must yield the expert's profit strictly less than $E[v_\theta - c_\theta]$. We consider all possible subgame equilibria:

Case 1: The subgame equilibrium is separating ($\beta^ = 0$).* In this case, the expert is always truthful. His maximum possible profit is $E[p_\theta - c_\theta]$ when the consumer accepts his recommendation with probability one. Since $\mathbf{p} \neq \mathbf{v}$, at least one price must be strictly less than its corresponding valuation. It follows that $E[p_\theta - c_\theta] < E[v_\theta - c_\theta]$.

Case 2: The subgame equilibrium is non-separating ($\beta^ > 0$).* If the expert cheats with positive probability, the consumer's unique best response (since $k = 0$) is to acquire a perfectly informative signal and reject any detected cheating. The expert's profit is $\mu_0(p_H - c_H) + (1 - \mu_0)(1 - \beta^*)(p_L - c_L)$. Since $\beta^* > 0$, this profit is strictly less than the profit from a separating equilibrium at the same price, $E[p_\theta - c_\theta]$, which is strictly less than $E[v_\theta - c_\theta]$.

Lastly, at the price of $\mathbf{p} = \mathbf{v}$, the consumer is indifferent between accepting and rejecting a truthful H -recommendation, which gives rise to multiple separating subgame equilibria (e.g., the expert recommends truthfully and the consumer always rejects H -recommendation without learning). We show that a subgame play consistent with the entire game involves (i) the expert making truthful recommendations; (ii) the consumer acquire information—precise enough to incentives the expert's honest recommendation—and thereby accepting all on-path recommendations. The resulting outcome is efficient, the expert's profit is $E[v_\theta - c_\theta]$, and the consumer's payoff is zero.

To see this a unique outcome of the full game, suppose for contradiction that the anticipated outcome at \mathbf{v} was a different subgame equilibrium, yielding the expert a profit $\pi^* < E[v_\theta - c_\theta]$. The expert could profitably deviate to a price $\mathbf{p}^\varepsilon := \mathbf{v} - \boldsymbol{\varepsilon}$, where $\boldsymbol{\varepsilon} = (\varepsilon, \varepsilon)$, for a sufficiently small $\varepsilon > 0$. By our previous finding, the unique subgame equilibrium at this interior price \mathbf{p}^ε involves full acceptance, yielding a profit of $E[v_\theta - c_\theta] - \varepsilon$. For a small enough ε , we have $E[v_\theta - c_\theta] - \varepsilon > \pi^*$. This profitable deviation rules out any outcome at $\mathbf{p} = \mathbf{v}$ other than the

expert-best one, a contradiction. Since any other price gives the expert's payoff strictly less than $E[v_\theta - c_\theta]$, this is a unique equilibrium up to consumer's learning strategy. □

A.3 Proof of Proposition 2

Remark 4. *Our no-learning benchmark Proposition 2 is analogous to Proposition 1 in Fong (2005), who study a model without endogenous information acquisition. We prove the result here in our notation for completeness.*

Proof of Proposition 2. The proof proceeds in two steps. In the first step, we characterize the subgame equilibrium for prices $\mathbf{p} \in \mathcal{P}$ with $p_H < v_H$. In the next step, we show that pricing at \mathbf{v} constitutes an unique equilibrium of the full game.

First, we show that any equilibrium in the subgame following price $\mathbf{p} \in \mathcal{P}$ with $p_H < v_H$ must be in mixed strategies. A pure-strategy pooling equilibrium (where the expert always recommends H) is not sustainable due to the lemon condition ($E[v] < c_H \leq p_H$), which would cause the consumer to always reject H recommendation. Likewise, a fully separating (truthful) equilibrium is not sustainable. Given the expert's honest H recommendation the consumer is prompted to always accept ($v_H - p_H > 0$), but this leads the expert to deviate to cheating on L problem ($p_H - c_L > p_L - c_L$).

Therefore, any equilibrium in the subgame following \mathbf{p} with $p_H < v_H$ must involve mixing by both players to render them indifferent. The consumer must accept H -recommendation with probability $\alpha(\mathbf{p}) \equiv \frac{p_L - c_L}{p_H - c_L}$ to make the expert indifferent between recommending H and L at $\theta = L$. Simultaneously, the expert must recommend H at $\theta = L$ with probability $\beta(\mathbf{p}) \equiv \frac{\mu_0(v_H - p_H)}{(1 - \mu_0)(p_H - v_L)}$ to make the consumer indifferent between accepting and rejecting H -treatment recommendation. These two indifference conditions characterize a unique subgame equilibrium for each price \mathbf{p} . The expert's profit in each subgame is

$$\pi(\mathbf{p}) = \mu_0 \alpha(\mathbf{p})(p_H - c_H) + (1 - \mu_0)(p_L - c_L),$$

which is strictly increasing in both p_L, p_H . Hence, an equilibrium price of the entire game cannot be other than \mathbf{v} .

As price \mathbf{p} approach \mathbf{v} , the subgame equilibrium strategies converge as follows:

$$\alpha^* = \lim_{\mathbf{p} \uparrow \mathbf{v}} \alpha(\mathbf{p}) = \frac{v_L - c_L}{v_H - c_L} \in (0, 1)$$

$$\beta^* = \lim_{\mathbf{p} \uparrow \mathbf{v}} \beta(\mathbf{p}) = \lim_{p_H \uparrow v_H} \frac{\mu_0(v_H - p_H)}{(1 - \mu_0)(p_H - v_L)} = 0$$

The rest of the reasoning is similar to the proof of Proposition 1. There are multiple subgame equilibria at the price $\mathbf{p} = \mathbf{v}$ but the strategies (α^*, β^*) constitute the expert's best subgame equilibrium play. This yields the expert a profit of $\pi(\mathbf{v}) = \mu_0\alpha^*(v_H - c_H) + (1 - \mu_0)(v_L - c_L)$. Since the expert's profit function $\pi(\mathbf{p})$ is continuous and strictly increasing over \mathcal{P} , any other subgame play involving the consumer rejecting with higher probability cannot be a part of an equilibrium of the entire game. If the consumer were to do so, the resulting expert's profit would be $\pi^* < \pi(\mathbf{v})$, prompting the expert to deviate to a price $\mathbf{p}' = \mathbf{v} - (\epsilon, \epsilon) \in \mathcal{P}$ for a sufficiently small $\epsilon > 0$. For a small enough ϵ , we have $\pi(\mathbf{p}') > \pi^*$. This profitable deviation rules out any outcome at $\mathbf{p} = \mathbf{v}$ other than the expert-best one, a contradiction. \square

A.4 Proof of Theorem 1

Proof of Theorem 1. Follows immediately from Proposition 3 and Proposition 4. \square

The remainder of this appendix builds the formal proof of Theorem 1. The argument rests on two key results, Proposition 3 and Proposition 4, which are proven in sequence. Proposition 3, which establishes the expert's improved profit and the necessity of cheating, involves a couple of steps. Its proof requires auxiliary lemmas 6 and 8 that characterize the solution to the consumer's problem. Specifically, Lemma 6 establishes a continuity property as information costs vanish, while subsequent lemmas establish monotonicity properties based on strictly decreasing differences. Proposition 4 then uses these results to show that the consumer must earn a positive surplus. We begin by stating and proving the necessary lemmas.

Lemma 6. *Fix any price $\mathbf{p} \in \mathcal{P}^\circ$ and belief $\mu > 0$. For any sequence of costs $\{k_n\}_{n=1}^\infty$ such that $k_n \downarrow 0$, let $\{\phi^*(k_n)\}_{n=1}^\infty$ be a corresponding sequence of solutions to the consumer's problem:*

$$\max_{\phi \in [0,1]^2} \left\{ -k_n C(\phi) + W_0(\phi; \mu, \mathbf{p}) \right\},$$

where $W_0(\phi; \mu, \mathbf{p})$ is the consumer's informational benefit. Then, any accumulation point $\hat{\phi} = (\hat{\phi}_H, \hat{\phi}_L)$ of the sequence of solutions $\{\phi^*(k_n)\}_{n=1}^\infty$ must satisfy $\hat{\phi}_H = 1$.

Remark 5 (Why the Maximum Theorem does not resolve Lemma 6). *The classic Berge Maximum Theorem (e.g., Berge, 1963) requires that (i) the objective $f(\phi, k)$ is continuous in (ϕ, k) ,*

and (ii) the feasible correspondence is nonempty, compact-valued, and continuous (upper and lower hemicontinuous). It then concludes that the value function is continuous and the argmax correspondence is nonempty, compact-valued, and upper hemicontinuous in k .

In our setting $f(\phi, k) = W_0(\phi) - kC(\phi)$ is not continuous on $[0, 1]^2 \times [0, \bar{k}]$ because we allow for $C(\phi) = +\infty$ for experiments involving fully revealing signal realizations, so for some ϕ it may be that $f(\phi, k) = -\infty$ for $k > 0$. Hence the hypotheses of the Maximum Theorem fail and it cannot be invoked directly.

Even under the stronger assumption that C were finite and continuous everywhere (so Berge applies), the conclusion would only give upper hemicontinuity at $k = 0$: any cluster point of maximizers of $W_0(\cdot) - kC(\cdot)$ belongs to $\arg \max W_0$. By itself this does not pin down the limit coordinate ϕ_H . The identification $\hat{\phi}_H = 1$ additionally uses the strict monotonicity of W_0 in ϕ_H (when $\mu > 0$ and $p_H < v_H$).

Proof of Lemma 6. Fix any price $\mathbf{p} \in \mathcal{P}^\circ$ and belief $\mu > 0$. For any sequence of costs $\{k_n\}_{n=1}^\infty$ such that $k_n \downarrow 0$, let $\{\phi^*(k_n)\}_{n=1}^\infty$ be a corresponding sequence of solutions to the consumer's problem:

$$\max_{\phi \in [0, 1]^2} \left\{ -k_n C(\phi) + W_0(\phi; \mu, \mathbf{p}) \right\} =: W_{k_n}(\phi; \mu, \mathbf{p}),$$

where $W_0(\phi; \mu, \mathbf{p})$ is the consumer's informational benefit,

$$W_0(\phi; \mu, \mathbf{p}) = \mu \phi_H (v_H - p_H) + (1 - \mu) \phi_L (v_L - p_H).$$

For brevity, we suppress the dependency of W_0, W_{k_n} on μ, \mathbf{p} and write $W_0(\phi), W_{k_n}(\phi)$.

We first note that $W_0(\phi)$ is strictly increasing in ϕ_H , which implies that any maximizer ϕ^* of the limit problem (where $k_n = 0$) must satisfy $\phi_H^* = 1$. We will now show that any limit point of the sequence of optimal choices $\{\phi^*(k_n)\}$ must also satisfy this property. The proof proceeds by contradiction.

Suppose for contradiction that there exists a sequence $k_n \downarrow 0$ and some associated sequence of solutions $\{\phi^*(k_n)\}_{n=1}^\infty$ such that $\phi_H^*(k_n)$ do not converge to 1. Because the sequence of optimal choices $\{\phi^*(k_n)\}$ lies in the compact set $[0, 1]^2$, the Bolzano-Weierstrass Theorem guarantees the existence of a convergent subsequence. Let this subsequence be denoted by $\{\phi^*(k_{n_j})\}_{j=1}^\infty$, and let its limit point be $\hat{\phi}$. Our contradiction hypothesis implies that we can choose this subsequence such that its limit point satisfies $\hat{\phi}_H < 1$.

Since $W_0(\phi)$ is strictly increasing in ϕ_H and our assumed limit point satisfies $\hat{\phi}_H < 1$, we know that the value of the objective function at the limit point, $W_0(\hat{\phi})$, is strictly less than its

value at $(1, \hat{\phi}_L)$. We can therefore define the gap in payoffs as $2\epsilon > 0$:

$$W_0(1, \hat{\phi}_L) = W_0(\underbrace{\hat{\phi}_H}_{<1}, \hat{\phi}_L) + 2\epsilon. \quad (2)$$

We show this leads us to contradiction. By continuity of W_0 in ϕ , for this $\epsilon > 0$, there exists $\delta > 0$ such that if $|\phi' - (1, \hat{\phi}_L)| < \delta$ then

$$|W_0(\phi') - W_0(1, \hat{\phi}_L)| < \epsilon.$$

Taking $\phi' = (1 - \frac{\delta}{\sqrt{2}}, \phi'_L)$ with $\phi'_L \in (\hat{\phi}_L - \frac{\delta}{\sqrt{2}}, \hat{\phi}_L + \frac{\delta}{\sqrt{2}}) \cap (0, 1)$ ensures $|\phi' - (1, \hat{\phi}_L)| < \delta$. This implies

$$\begin{aligned} W_0(\phi') &> W_0(1, \hat{\phi}_L) - \epsilon \\ &= W_0(\hat{\phi}) + \epsilon, \end{aligned} \quad (3)$$

where the last equality follows from substituting $W_0(1, \hat{\phi}_L)$ using (2). By definition of $\phi^*(k_n)$ being an optimum of $W_{k_n}(\cdot)$, for every j in the subsequence,

$$W_{k_{n_j}}(\phi^*(k_{n_j})) \geq W_{k_{n_j}}(\phi'). \quad (4)$$

Substituting the definition of $W_{k_{n_j}}$ into inequality (4) gives

$$-k_{n_j}C(\phi^*(k_{n_j})) + W_0(\phi^*(k_{n_j})) \geq -k_{n_j}C(\phi') + W_0(\phi'),$$

where we note $C(\phi') < \infty$ and C is continuous in a neighborhood around ϕ' (we choose ϕ' strictly inside $(0, 1)^2$). Rearranging this inequality to isolate the cost terms on the RHS gives

$$W_0(\phi') - W_0(\phi^*(k_{n_j})) \leq k_{n_j}[C(\phi') - C(\phi^*(k_{n_j}))]. \quad (5)$$

We now investigate the limit of both sides of (5) as $j \rightarrow \infty$. We show that in the limit the LHS is bounded away from 0 while the RHS is 0, leading to a contradiction.

- The LHS is $W_0(\phi') - W_0(\phi^*(k_{n_j}))$. Since $\phi^*(k_{n_j}) \rightarrow \hat{\phi}$ and W_0 is continuous, the LHS converges to $W_0(\phi') - W_0(\hat{\phi})$. By inequality (3), we know this limit is bounded below by a strictly positive number:

$$\lim_{j \rightarrow \infty} \left(W_0(\phi') - W_0(\phi^*(k_{n_j})) \right) = W_0(\phi') - W_0(\hat{\phi}) > \epsilon > 0.$$

- The RHS is $k_{n_j}[C(\phi') - C(\phi^*(k_{n_j}))]$. We analyze its two components. First, since $C(\phi') < \infty$ and $k_{n_j} \downarrow 0$, the first term $k_{n_j}C(\phi') \rightarrow 0$. For the second term, we note that the consumer's utility at her optimum must be at least that of an outside option yielding zero payoff (e.g., rejecting without learning). Thus, for all j :

$$W_{k_{n_j}}(\phi^*(k_{n_j})) = -k_{n_j}C(\phi^*(k_{n_j})) + W_0(\phi^*(k_{n_j})) \geq 0,$$

which implies $k_{n_j}C(\phi^*(k_{n_j})) \leq W_0(\phi^*(k_{n_j}))$. Since W_0 is continuous on a compact set, it is bounded above, which means the sequence $\{k_{n_j}C(\phi^*(k_{n_j}))\}$ is bounded. As the term $-k_{n_j}C(\phi^*(k_{n_j}))$ is also non-positive, the limit of the RHS must be less than or equal to zero:

$$\limsup_{j \rightarrow \infty} k_{n_j}[C(\phi') - C(\phi^*(k_{n_j}))] = \lim_{j \rightarrow \infty} k_{n_j}C(\phi') - \liminf_{j \rightarrow \infty} k_{n_j}C(\phi^*(k_{n_j})) \leq 0 - 0 = 0.$$

We have shown that the LHS converges to a value $> \epsilon > 0$, while the RHS converges to a value ≤ 0 . This implies $\epsilon \leq 0$, which contradicts our finding that ϵ must be strictly positive.

Therefore, our initial premise that a subsequence of solutions converges to a limit with $\hat{\phi}_H < 1$ must be false. Since this holds for any convergent subsequence, all limit points $\hat{\phi}$ of $\{\phi^*(k_n)\}$ must have $\hat{\phi}_H = 1$.

□

Lemma 7 (Strictly Increasing Differences). *Fix $k > 0$. Fix $\mathbf{p} \in \mathcal{P}$ with $p_H < v_H$. The consumer's objective function $W_k(\phi; \mu, \mathbf{p}) = W_0(\phi; \mu, \mathbf{p}) - kC(\phi)$ has strictly increasing differences in (ϕ_H, μ) and in (ϕ_L, μ) .*

Proof. Since $W_k(\phi; \mu, \mathbf{p}) = W_0(\phi; \mu, \mathbf{p}) - kC(\phi)$ and C does not depend on μ , $\partial_\mu W_k = \partial_\mu W_0$. With $W_0(\phi; \mu, \mathbf{p}) = \mu \phi_H (v_H - p_H) + (1 - \mu) \phi_L (v_L - p_H)$,

$$\begin{aligned} \frac{\partial^2 W_k}{\partial \phi_H \partial \mu} &= \frac{\partial}{\partial \phi_H} [\partial_\mu W_0] = v_H - p_H > 0 \quad (\text{since } p_H < v_H), \\ \frac{\partial^2 W_k}{\partial \phi_L \partial \mu} &= \frac{\partial}{\partial \phi_L} [\partial_\mu W_0] = p_H - v_L > 0 \quad (\text{since } p_H > v_L). \end{aligned}$$

Thus W_k has strictly increasing differences in (ϕ_H, μ) and (ϕ_L, μ) . □

Lemma 8 (Monotone Selection). *Fix k . Fix $\mathbf{p} \in \mathcal{P}$ with $p_H < v_H$. Denote by $\Phi(\mu) \subset [0, 1]^2$ the set of solutions to the consumer's problem:*

$$\max_{\phi} \{W_0(\phi; \mu, \mathbf{p}) - kC(\phi)\}.$$

If $\mu' \geq \mu$ then $\phi'_H(\mu, \mathbf{p}) \geq \phi_H(\mu, \mathbf{p})$, $\phi_L(\mu, \mathbf{p}) \geq \phi'_L(\mu, \mathbf{p})$ for any $\phi' \in \Phi(\mu)$, $\phi \in \Phi(\mu)$.

Proof. The choice set $[0, 1]^2$ is a complete lattice and is independent of μ . By Lemma 7, W_k has strictly increasing differences in (ϕ_H, μ) and (ϕ_L, μ) , hence (strict) single crossing in (ϕ, μ) . Since W_0 is separable in (ϕ_H, ϕ_L) and $C(\phi)$ is submodular in ϕ (our assumption on C), W_k is supermodular in ϕ . Therefore, by Milgrom and Shannon (1994), every selection from $\arg \max_{\phi \in [0, 1]^2} W_k(\phi; \mu, \mathbf{p})$ is nondecreasing in μ . Hence if $\mu' \geq \mu$ and $\phi' \in \Phi(\mu')$, $\phi \in \Phi(\mu)$, then $\phi'_\theta \geq \phi_\theta$ for $\theta \in \{H, L\}$. \square

Proposition 3. *There exists $\bar{k} > 0$ such that for any $k \in (0, \bar{k})$, the expert's equilibrium profit is strictly greater than his profit at no-learning benchmark ($k = \infty$). Moreover, the expert cheats with positive probability in any equilibrium.*

Proof. First, we prove the first part of the statement. Let $\Pi(k)$ denote the expert's worst equilibrium profit in the full game at information cost k , i.e.,

$$\Pi(k) := \sup_{\mathbf{p} \in \mathcal{P}} \underline{\pi}(\mathbf{p}; k), \quad \underline{\pi}(\mathbf{p}; k) := \inf_{(\beta, \mu, \phi) \in \mathcal{E}(\mathbf{p}, k)} \pi(\mathbf{p}; \beta, \mu, \phi; k). \quad (6)$$

Here $\pi(\mathbf{p}; \beta, \mu, \phi; k)$ is the expert's profit in a subgame equilibrium (β, μ, ϕ) following price \mathbf{p} at cost $k > 0$, and $\mathcal{E}(\mathbf{p}, k)$ denotes the set of all such equilibria, which is non-empty. By Propositions 1 and 2,

$$\Pi(0) = \mathbb{E}[v_\theta - c_\theta] > \Pi(\infty).$$

We claim that $\Pi(\cdot)$ is *lower semicontinuous at 0*: for every $\varepsilon > 0$ there exists $\bar{k} > 0$ such that for all $k \in (0, \bar{k})$,

$$\Pi(k) \geq \Pi(0) - \varepsilon.$$

Fix $\varepsilon > 0$ and consider the price $\mathbf{p}^\varepsilon := \mathbf{v} - \frac{1}{2}\boldsymbol{\varepsilon} \in \mathcal{P}$, where $\mathbf{v} = (v_L, v_H)$ and $\boldsymbol{\varepsilon} = (\varepsilon, \varepsilon)$. Consider any equilibrium of the subgame at price \mathbf{p}^ε and denote its objects by $(\beta^*(\mathbf{p}^\varepsilon; k), \mu^*(\mathbf{p}^\varepsilon; k), \phi^*(\mathbf{p}^\varepsilon; k))$. The expert's profit in this subgame is

$$\pi(\mathbf{p}^\varepsilon; k) = \mu_0 \phi_H^*(k)(v_H - c_H - \frac{\varepsilon}{2}) + (1 - \mu_0) \max\{\phi_L^*(k)(v_H - c_L - \frac{\varepsilon}{2}), v_L - c_L - \frac{\varepsilon}{2}\}.$$

Since the expert recommends H with probability 1 in state H and (weakly) less in state L , Bayes' rule implies the posterior on H along the path is at least the prior: $\mu^*(\mathbf{p}^\varepsilon; k) \geq \mu_0$. By Lemma 8, $\phi_\theta^*(k) \geq \phi_\theta(\mu_0; k)$ for each $\theta \in \{L, H\}$. Therefore,

$$\begin{aligned} \pi(\mathbf{p}^\varepsilon; k) &\geq \mu_0 \phi_H(\mu_0; k)(v_H - c_H - \frac{\varepsilon}{2}) + (1 - \mu_0)(v_L - c_L - \frac{\varepsilon}{2}) \\ &= \mu_0 \phi_H(\mu_0; k)(v_H - c_H) + (1 - \mu_0)(v_L - c_L) - \frac{\varepsilon}{2}. \end{aligned}$$

By Lemma 6, $\phi_H(\mu_0; k) \rightarrow 1$ as $k \downarrow 0$, so there exists $k' > 0$ such that for all $k < k'$,

$$\pi(\mathbf{p}^\varepsilon; k) \geq \mathbb{E}[v_\theta - c_\theta] - \varepsilon = \Pi(0) - \varepsilon.$$

Since this inequality holds for *every equilibrium of the subgame at price \mathbf{p}^ε* , the *worst* equilibrium profit at that price is at least $\Pi(0) - \varepsilon$. Because the expert can always post \mathbf{p}^ε , we obtain for all $k < k'$:

$$\Pi(k) \geq \underline{\pi}(\mathbf{p}^\varepsilon; k) \geq \Pi(0) - \varepsilon.$$

Choosing $\varepsilon > 0$ small that $\Pi(0) - \varepsilon > \Pi(\infty)$ yields $\bar{k} > 0$ with $\Pi(k) > \Pi(\infty)$ for all $k \in (0, \bar{k})$, which proves the first part of the statement.

Second, we show the expert cheats with positive probability in any equilibrium when $k \in (0, \bar{k})$. Suppose, to the contrary, that some equilibrium is separating at a price with $p_H < v_H$. Because information is costly and on-path recommendations are truthful (fully revealing), the consumer strictly prefers not to acquire information and accepts H when recommended. At $\theta = L$ the expert then gains by deviating to recommend H (a profitable deviation), contradicting separation. Thus separation can only occur, if at all, at $p_H = v_H$. At $p_H = v_H$ the consumer's informed benefit is zero, she does not learn, and the expert's profit cannot exceed the no-learning benchmark $\Pi(\infty)$ (which is the expert's best equilibrium profit following the price $\mathbf{p} = \mathbf{v}$ at cost $k > 0$), contradicting $\Pi(k) > \Pi(\infty)$ for $k \in (0, \bar{k})$. Hence in any equilibrium for $k \in (0, \bar{k})$, the expert recommends H in state L with positive probability. □

Remark 6. We use sup and inf in the definition (6) for $\Pi(k)$ rather than max and min. If one establishes that $\mathcal{E}(\mathbf{p}, k)$ is nonempty and compact and that $\pi(\mathbf{p}; \beta, \mu, \phi; k)$ is continuous in the equilibrium objects, then the inf is in fact a min; similarly, if $\mathbf{p} \mapsto \underline{\pi}(\mathbf{p}; k)$ attains its supremum on \mathcal{P} , then the sup is a max. None of these stronger attainment properties are required for the argument here.

Proposition 4. There exists $\bar{k} > 0$ such that for any $k < \bar{k}$, the consumer's payoff is strictly positive in any equilibrium. Moreover, the consumer acquires informative signal ϕ^* .

Remark 7. The consumer always acquiring some informative signal ϕ^* combined with the expert cheating with positive probability implies that the outcome is always inefficient as some H -treatment recommended at $\theta = L$ would be rejected and the consumer's cost spent on acquiring information is a social loss.

Proof. Fix arbitrary $k \in (0, \bar{k})$ where \bar{k} is the threshold in Proposition 3 for which $\Pi(k) > \Pi(\infty)$ for any $k < \bar{k}$. Consider *any equilibrium* $(\mathbf{p}^*, \beta^*, \mu^*, \phi^*)$ when the information cost is k .

First, we show the first part of the statement. We prove by contradiction. Suppose that the consumer's payoff is zero. We will show that ϕ^* is neither informative nor uninformative, leading to a contradiction.

We first show that ϕ^* cannot be uninformative. If ϕ^* were uninformative, then the consumer's payoff being zero means

$$0 = W_k(\phi^*, \mu^*, \mathbf{p}^*; k) = -kC(\phi^*) + W_0(\phi^*, \mu^*, \mathbf{p}^*) = E[v \mid \mu^*] - p_H^*, \quad (7)$$

where $E[v \mid \mu]$ is the expected valuation given μ . By Proposition 3, the expert cheats with positive probability in any outcome for $k \in (0, \bar{k})$, so $\mu^* < 1$. If $\mu^* = \mu_0$ then our assumption on parameters and $\mathbf{p} \in \mathcal{P}$ imply $E[v \mid \mu_0] < c_H \leq p_H^*$, which is a contradiction to (7). Therefore, it must be $\mu^* \in (\mu_0, 1)$. This implies that the expert must be semi-separating, i.e., $\beta^* \in (0, 1)$ and that the consumer must be randomizing between accepting and rejecting, without learning, in the manner which keeps the expert indifferent between cheating and recommending honestly for L problem. This is precisely one of subgame equilibria in the no-learning benchmark ($k = \infty$). Because this profit cannot be greater than $\Pi(\infty)$, the expert's equilibrium profit $\Pi(k)$ cannot be greater than $\Pi(\infty)$, a contradiction.

It remains to show that ϕ^* cannot be informative. Assume that ϕ^* were informative. Observe that quasi-convexity of $C(\phi)$ implies that $W_k(\phi)$ is quasi-concave in ϕ , where we suppress dependency on $(\mu^*, \mathbf{p}^*; k)$ for brevity. Then for any $\lambda \in [0, 1]$,

$$W_k(\lambda\phi^* + (1 - \lambda)\phi_{rej}) \geq \min\{W_k(\phi^*), W_k(\phi_{rej})\} = 0, \quad (8)$$

where ϕ_{rej} is the uninformative “reject” and we suppressed the argument β for W_k . The equality in (8) follows from $W_k(\phi_{rej}) = 0$ (rejecting gives zero payoff) and $W_k(\phi^*) = 0$ (our hypothesis that the consumer surplus is zero). The optimality of ϕ^* implies that

$$0 = W_k(\phi^*) \geq W_k(\lambda\phi^* + (1 - \lambda)\phi_{rej})$$

Hence, we must have $0 = W_k(\phi^*) = W_k(\lambda\phi^* + (1 - \lambda)\phi_{rej})$ for all $\lambda \in [0, 1]$. However, for any $\lambda \in (0, 1)$, the convex combination $\lambda\phi^* + (1 - \lambda)\phi_{rej} \equiv \phi^\lambda$ is an informative experiment. Since $C(\phi)$ is strictly quasi-convex over informative experiments, $W_k(\phi, \beta)$ is strictly quasi-concave over informative experiments, which implies that for any $\lambda \in (0, 1)$,

$$W_k(\lambda\phi^\lambda + (1 - \lambda)\phi^*) > \max\{W_k(\phi^*), W_k(\phi^\lambda)\},$$

because ϕ^*, ϕ^λ are both informative experiments. This contradicts the optimality of ϕ^* .

We conclude the proof by showing that the consumer must be acquiring some informative ϕ^* in any equilibrium. By Lemma 3, given the expert's strategy β^* , the consumer acquires some informative signal for any price p_H in the interval

$$(\underline{p}(\beta^*), \bar{p}(\beta^*))$$

where the thresholds $\underline{p}(\beta^*), \bar{p}(\beta^*)$ are defined in Lemma 3. Therefore, it is enough to show that $p_H^* \in (\underline{p}(\beta^*), \bar{p}(\beta^*))$.

First, the interval $(\underline{p}(\beta^*), \bar{p}(\beta^*))$ is not empty. If it were, then $\underline{p}(\beta^*) = \bar{p}(\beta^*) = E[v|\mu(\beta^*)]$, and the consumer never learns given the price p_H^* and the expert's strategy β^* . Similar to the argument above, this is precisely one of subgame equilibria in the no-learning benchmark ($k = \infty$). The expert's equilibrium profit $\Pi(k)$ cannot be greater than $\Pi(\infty)$, a contradiction.

Next, suppose for contradiction that either $p_H \geq \bar{p}(\beta^*)$ or $p_H \leq \underline{p}(\beta^*)$. If $p_H \geq \bar{p}(\beta^*)$ then Lemma 3 implies that the consumer's payoff is zero, a contradiction. If $p_H \leq \underline{p}(\beta^*)$, then, again by Lemma 3, accepting is (uniquely) optimal.⁶ This implies that the expert must be pooling ($\beta^* = 1$) and the consumer's payoff must be

$$V = E[v|\mu_0] - p_H^* > 0$$

which is a contradiction since $E[v | \mu_0] < c_H \leq p_H^*$ from the lemon condition. Therefore, the consumer must be acquiring some informative ϕ^* in equilibrium. □

A.5 Expanding \mathcal{P}

In this section, we show that it is without loss of generality to restrict the expert's choice set for prices to \mathcal{P} , where

$$\mathcal{P} := [c_L, v_L] \times [c_H, v_H].$$

Specifically, we prove the following statement.

Proposition 5. *For every $\mathbf{p} \in \mathbb{R}_+^2 \setminus \mathcal{P}$, there exists $\mathbf{p}' \in \mathcal{P}$ such that the expert's profit in the subgame following \mathbf{p}' is weakly greater than his profit in the subgame following \mathbf{p} for any*

⁶Because the consumer enjoys strictly positive payoff, $V = E[v|\mu(\beta^*)] - p_H^* > 0$, accepting is strictly better than rejecting.

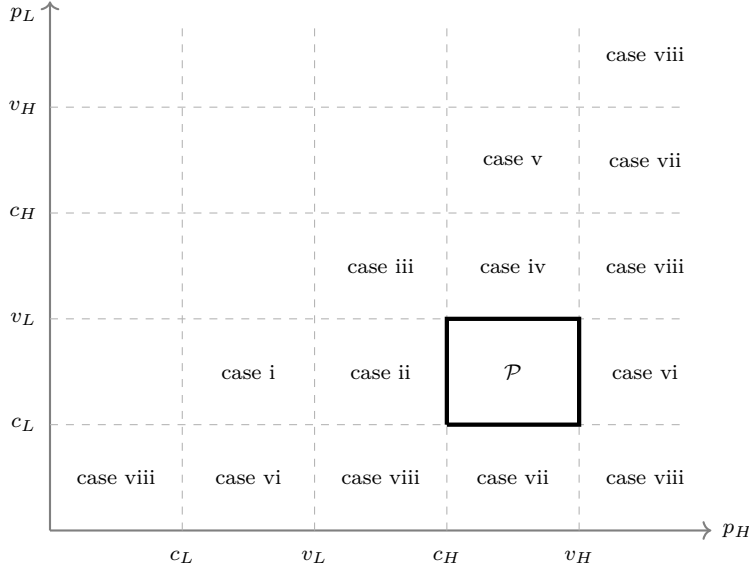


Figure 7: Illustration for Proof of Proposition 5.

equilibrium. Moreover, if these profits coincide, then the consumer's equilibrium surplus is weakly higher under \mathbf{p}' than under \mathbf{p} .

The proof strategy parallels that of Kim et al. (2025), which show that restricting the price domain is without loss when consumers can receive noisy signals (there, a signal is part of the game and not a choice). Here, we exploit the power of monotone comparative statics to establish the similar restriction under our model. But the high level intuition remains the same: due to the expert's inability to commit at a recommendation stage a treatment with price outside the range will not be recommended by the expert in any equilibrium.

Proof. For each $\mathbf{p} \in \mathbb{R}_+^2$, denote by $\pi(\mathbf{p})$ the expert's equilibrium payoffs in the recommendation subgame $\Gamma(\mathbf{p})$.

Case i. Suppose $\mathbf{p} \in [c_L, v_L]^2$. Recommending either treatment at $\theta = H$ gives a negative profit $p_\theta - c_H < 0$ if accepted and zero profit otherwise. Hence, it is weakly dominated by no-treatment which always results in zero profit which we rule out as an equilibrium play. The expert do not recommend any treatment at $\theta = H$, hence no trade occurs at $\theta = H$. The expert's profit $\pi(\mathbf{p})$ at \mathbf{p} is bounded above by $(1 - \mu_0)(\max\{p_L, p_H\} - c_L)$. Choose $\mathbf{p}' = (v_L, c_H) \in \mathcal{P}$. Following this price \mathbf{p}' , the expert guarantees himself a profit at least $(1 - \mu_0)(v_L - c_L)$, i.e.,

$$\pi(\mathbf{p}) \leq (1 - \mu_0)(\max\{p_L, p_H\} - c_L) \leq (1 - \mu_0)(v_L - c_L) \leq \pi(\mathbf{p}').$$

If $\max\{p_L, p_H\} < v_L$, price \mathbf{p} is never chosen in equilibrium since \mathbf{p}' yields strictly higher profit. If $\max\{p_L, p_H\} = v_L$, then the consumer surplus is zero at \mathbf{p} so it cannot be worse under \mathbf{p}' .

Case ii. Suppose $\mathbf{p} \in [c_L, v_L] \times (v_L, c_H)$. Following our discussion, the H -treatment with price $p_H \in (v_L, c_H)$ will not be accepted in any equilibrium. This puts us back to Case i.

Case iii. Suppose $\mathbf{p} \in (v_L, c_H)^2$. No treatment will be accepted by the consumer and both players get zero payoffs. The rest of the reasoning is similar to Case i.

Case iv. Suppose $\mathbf{p} \in (v_L, c_H) \times [c_H, v_H]$. If $p_H = c_H$, then this case is similar to Case iii because the sales for H -treatment do not generate profit. Therefore, assume $p_H > c_H$. Clearly, recommending H -treatment at $\theta = H$ weakly dominates recommending both L - and no-treatment. If there is no trade at H -treatment recommendation, both players get zero payoffs and back to Case iii. If trade happens with positive probability at H -treatment recommendation, then the expert must be recommending H -treatment with probability one at $\theta = L$. Hence, the consumer holds belief μ_0 upon recommended H -treatment (no belief updates). The expert's profit following \mathbf{p} is given by

$$\pi(\mathbf{p}) = \mu_0 \phi_H(\mu_0, \mathbf{p})(p_H - c_H) + (1 - \mu_0) \phi_L(\mu_0, \mathbf{p})(p_H - c_L), \quad (9)$$

where $\phi(\mu_0, \mathbf{p})$ is the consumer's equilibrium learning strategy. We claim that $\phi(\mu_0, \mathbf{p})$ must be informative. To see this, notice that

- $\phi(\mu_0, \mathbf{p})$ cannot be both uninformative and involving positive probability of acceptance. If it were, then the consumer surplus would be negative from the lemon condition:

$$E[v \mid \mu_0] < c_H < p_H.$$

Hence, there is a profitable deviation to rejection with probability one (ϕ_{rej}), which yields zero payoff, a contradiction.

- If $\phi(\mu_0, \mathbf{p}) = \phi_{rej}$, then both the expert and the consumer get zero. Therefore, the expert never chooses such \mathbf{p} in equilibrium, similar to Case i.
- If $\phi(\mu_0, \mathbf{p})$ were informative, then it is a uniquely optimal solution for the consumer from strict quasi-convexity of $C(\phi)$ over informative signals.

Pick $\mathbf{p}' = (v_L, p_H)$ (where p_H is the H -treatment price in profile \mathbf{p}). We show that the expert's profit following \mathbf{p}' is weakly higher than (9) in any equilibrium. The expert's profit following

\mathbf{p}' is given by

$$\pi(\mathbf{p}') = \mu_0 \phi_H^*(p_H - c_H) + (1 - \mu_0) \max\{\phi_L^*(p_H - c_L), v_L - c_L\}. \quad (10)$$

Since $\mathbf{p}' \in \mathcal{P}$, the consumer holds some belief $\mu \geq \mu_0$ weakly above prior upon recommended H -treatment. Since ϕ^* is a best response to μ , $\phi^* = \phi(\mathbf{p}, \mu)$. By Lemma 8, $\mu \geq \mu_0$ implies that $\phi_\theta(\mu(\beta^*), \mathbf{p}') \geq \phi_\theta(\mu_0, \mathbf{p})$, $\forall \theta \in \{L, H\}$. Therefore,

$$\begin{aligned} \pi(\mathbf{p}') &= \mu_0 \phi_H^*(p_H - c_H) + (1 - \mu_0) \max\{\phi_L^*(p_H - c_L), v_L - c_L\} \\ &\geq \mu_0 \phi_H(\mu_0, \mathbf{p})(p_H - c_H) + (1 - \mu_0) \phi_L(\mu_0, \mathbf{p})(p_H - c_L) \\ &= \pi(\mathbf{p}). \end{aligned} \quad (11)$$

In other words, if the price for the treatment being recommended (p_H) stays unchanged, then the consumer's state-contingent acceptance probabilities ϕ_L, ϕ_H are non-decreasing in the consumer's belief (equivalently, non-increasing in the expert's cheating probability β).

By generalized Envelope Theorem, the consumer's value $V(\mathbf{p}, \mu)$ is also non-decreasing in μ for fixed price p_H . Hence, it is either that the expert never chooses such \mathbf{p} in equilibrium ($\pi(\mathbf{p}) < \pi(\mathbf{p}')$), or the expert's profit stays unchanged ($\pi(\mathbf{p}) = \pi(\mathbf{p}')$) but the consumer's surplus cannot be greater at \mathbf{p}' than at \mathbf{p} in any equilibrium.

Case v. Suppose $\mathbf{p} \in [c_H, v_H]^2$. Observe that

- if neither of the prices are traded in equilibrium, we are back to Case iii.
- if only p_θ is traded in equilibrium, we are back to Case iv.

We consider the case where both prices p_H, p_L are traded in equilibrium. This, in turn, implies that both treatments H, L must be recommended at $\theta = H$ with positive probabilities. Now that the consumer has distinct learning decision following both H, L recommendations, we denote by $\phi^{\hat{\theta}}$ by the consumer's acquisition strategy following $\hat{\theta}$ -treatment recommendation, so that $\phi_\theta^{\hat{\theta}}$ is the consumer's acceptance probability in state θ when recommended $\hat{\theta}$. Denote by $\mu^{\hat{\theta}}$ the consumer's belief about $\theta = H$ following $\hat{\theta}$ recommendation. Consider each of the following subcases.

Subcase i: Exactly one of the treatment is not recommended at $\theta = L$. Suppose that L -treatment is not recommended at $\theta = L$. Then the consumer believes it is less likely that the true state is $\theta = H$ relative to her prior μ_0 following H -treatment recommendation

($\mu^H < \mu_0$) and she believes for certain the true state is $\theta = H$ following L -treatment recommendation ($\mu^L = 1$). Because the expert randomizes between two treatments at $\theta = H$, his profit must be the same for both treatments at $\theta = H$, i.e.,

$$\phi_H^H(p_H - c_H) = \phi_H^L(p_L - c_H)$$

The expert's expected profit, written in H -treatment recommendation, is then

$$\mu_0 \phi_H^H(p_H - c_H) + (1 - \mu_0) \phi_L^H(p_H - c_L).$$

Now choose $\mathbf{p}' = (v_L, p_H)$, which is in the price range for Case iv. Since ϕ_H is the consumer's best response given her belief, we may write $\phi^H = \phi(\mu^H, \mathbf{p})$. Similar to our discussion in Case iv, $\phi^H = \phi(\mu^H, \mathbf{p})$ must be unique from the lemon condition. Also, at price \mathbf{p}' , pooling H -treatment is a weakly dominant strategy for the expert and his profit is uniquely

$$\mu_0 \phi_H(\mu_0, \mathbf{p}') (p_H - c_H) + (1 - \mu_0) \phi_L(\mu_0, \mathbf{p}') (p_H - c_L).$$

Because both $\phi(\mu^H, \mathbf{p})$ and $\phi(\mu_0, \mathbf{p}')$ are unique given the same price for the recommended treatment (p_H), from Lemma* $\mu^H < \mu_0$ implies $\phi(\mu^H, \mathbf{p}) < \phi(\mu_0, \mathbf{p}')$. This implies

$$\begin{aligned} \pi(\mathbf{p}) &= \mu_0 \phi_H^H(p_H - c_H) + (1 - \mu_0) \phi_L^H(p_H - c_L) \\ &< \mu_0 \phi_H(\mu_0, \mathbf{p}') (p_H - c_H) + (1 - \mu_0) \phi_L(\mu_0, \mathbf{p}') (p_H - c_L) \\ &= \pi(\mathbf{p}'), \end{aligned}$$

so that the expert would never choose \mathbf{p} in any equilibrium. The case for H -treatment not recommended at $\theta = L$ can be shown symmetrically by choosing $\mathbf{p}'' = (v_L, p_L)$ instead (with $p_L \in [c_H, v_H]$ is the price for L -treatment in the original price profile \mathbf{p}).

Subcase ii: Both of the treatment is recommended at $\theta = L$. If distinct treatments are recommended with distinct probabilities in each state, then the consumer holds distinct beliefs, say $\mu^H < \mu_0 < \mu^L$, which puts us back to *Subcase i*. If both of the treatments are recommended with equal probabilities in each state, then consumer do not update beliefs $\mu^H = \mu^L = \mu_0$. This puts us back to Case iv.

All the remaining cases are either mirror images of the previously discussed cases or gives the expert 0 profit. This concludes the proof. \square