

Information Acquisition in the Market for Lemons

Kyungmin Kim^{*} Guram Lobzhanidze[†] Lucie Xu[‡]

October 19, 2026

Abstract

We consider a dynamic decentralized market in which sellers possess private information about the quality of their unit and buyers can acquire information about the quality of their randomly matched partner. We assume that buyers' information acquisition is fully flexible subject to uniformly posterior-separable information costs. We characterize stationary market equilibria and show that complementarity between buyers' screening and sellers' trading incentives can cause equilibrium multiplicity, and lower information costs can be harmful to buyers and even Pareto worsening.

JEL Classification Codes: C72, C78, D82,

Keywords: Adverse selection; flexible information acquisition; UPS information costs

1 Introduction

Advances in computing, communication, and data processing have substantially reduced the costs of acquiring knowledge and transmitting information. For an individual decision maker facing a fixed environment, the benefits are straightforward. Consider a buyer who is uncertain about the value of an asset. *Ceteris paribus*, a reduction in information costs necessarily (albeit not always strictly) improves his expected payoff: he can obtain better information, or retain his previous information and decision rule at a lower cost. The same reasoning applies to all other contexts, such as a firm assessing a prospective employee or an investor evaluating a project.

Technological progress, however, rarely lowers information costs for only one individual. Other buyers gain access to the same improvements and adjust their information acquisition and trading behavior as well. Their collective response changes which assets remain available, the

^{*}Department of Economics, Emory University. Email: kyungmin.kim@emory.edu

[†]Department of Economics, Emory University. Email: guram.lobzhanidze@emory.edu

[‡]Department of Economics, Emory University. Email: lucie.xu@emory.edu

prices sellers are willing to accept, and the value of continuing to search. Thus, a buyer does not simply face a cheaper version of his original decision problem; he faces a different decision problem. This paper studies these market-wide adjustments and their economic consequences. Do the private benefits of cheaper information survive in market equilibrium? Can buyers themselves be worse off when all of them can acquire information at lower costs? In other words, what are the general equilibrium effects of declining information costs?

We consider a canonical dynamic decentralized version of Akerlof’s market for lemons.¹ New buyers and sellers enter in each period, and market participants are randomly matched. Each seller privately knows the quality of her object, which is either high or low, and there are positive gains from trade for both qualities. Upon meeting a seller, a buyer can acquire information about quality before making a take-it-or-leave-it price offer. If the seller accepts, both parties exit; otherwise, they return to the market and meet new partners. Buyers do not observe sellers’ trading histories, so their prior beliefs depend on the market composition of available sellers. We focus on stationary market equilibria and examine the effects of a common reduction in buyers’ information costs.

Unlike most existing studies of dynamic adverse selection (e.g., [Kaya and Kim, 2018](#); [Zou, 2025](#)), we assume that buyers’ information acquisition is fully flexible. Following the modern information design approach, each buyer can choose any experiment about the matched seller’s quality, subject to a uniformly posterior-separable (UPS) cost function ([Caplin et al., 2022](#)): the cost of an experiment is proportional to its expected reduction in uncertainty, with a common parameter, λ , measuring the unit cost of information. Buyers thus choose not only whether to acquire information, but also how much and what to learn.

This flexibility allows us to isolate strategic information-acquisition incentives from ad hoc restrictions on the available signals. Even at a given level of expenditure, buyers can choose among experiments that reveal different information, tailoring their learning to the tradeoffs between making a high offer, making a low offer, and continuing to search. The UPS specification also makes comparative statics particularly natural: a reduction in λ proportionally lowers experiment costs at any given prior, while buyers optimally adjust both the amount and the content of their information. By contrast, lowering the fixed cost of observing a given signal changes the incentive to acquire it but not its content conditional on acquisition. Modeling technological progress as an improvement in an exogenous signal instead requires specifying how that content changes. We change the cost of information rather than prescribe the information buyers receive, allowing their experiments to respond endogenously to both lower costs and the market adjustments they induce.

¹Our model builds upon the seminal stationary framework of [Moreno and Wooders \(2010\)](#) in which (uninformed) buyers do not receive any information about (informed) sellers. See, e.g., [Janssen and Roy \(2002\)](#) for a centralized market model, [Blouin and Serrano \(2001\)](#), [Camargo and Lester \(2014\)](#), and [Moreno and Wooders \(2016\)](#) for non-stationary market models, and [Hörner and Vieille \(2009\)](#) and [Kim \(2017\)](#) for different observability assumptions about sellers’ histories.

We characterize stationary market equilibria. When adverse selection is severe, buyers necessarily acquire information: a no-learning equilibrium does not exist for any finite unit cost. Under mild adverse selection, by contrast, a no-learning equilibrium exists if and only if information is sufficiently costly. In an informative equilibrium, optimal learning takes a binary form. The favorable outcome induces the buyer to offer the high price, equal to high-type sellers' reservation price. Following the unfavorable outcome, He offers low-type sellers' reservation price, makes a losing offer that both types reject, or randomizes between these alternatives. This structure allows us to characterize screening equilibria according to whether buyers make either, both, or neither of two errors: offering the high price to a low-type seller and failing to offer it to a high-type seller.

The analysis yields three main economic results. First, unlike the no-information and perfect-information benchmarks, imperfect screening can support multiple stationary equilibria. The source is complementarity between buyers' screening and low-type sellers' trading incentives. In a pool rich in high-type sellers, buyers can optimally demand strong evidence of high quality before making a high offer. Low-type sellers then have little reason to wait for such an offer, which lowers their reservation price and encourages them to trade quickly, sustaining a favorable pool. Conversely, less stringent screening gives low-type sellers a greater chance of receiving a high offer. Their stronger incentive to wait raises their reservation price and slows their exit, sustaining a less favorable pool. Different screening strategies and market compositions can therefore be mutually consistent under the same fundamentals.

Second, a reduction in information costs can be harmful to (uninformed) buyers, overturning the direct private benefit that motivates our analysis.² In fact, it can even be Pareto worsening, lowering both buyers' and low-type sellers' payoffs. More effective screening reduces low-type sellers' chances of receiving a high price and therefore lowers their expected payoffs. At the same time, it can make high-type sellers trade relatively quickly, leaving a less favorable pool for subsequent matches. The resulting deterioration in buyers' trading opportunities can outweigh their direct gains from lower information costs. Thus, the same improvement in information technology can benefit an individual buyer when others' behavior is held fixed yet harm buyers when adopted throughout the market.

Finally, cheaper information need not increase the amount of information acquired, measured by the expected reduction in uncertainty per match. When inexpensive screening makes the pool heavily concentrated in one type, buyers can learn quality almost perfectly while resolving little uncertainty. An increase in information costs may then make the pool sufficiently more uncertain that the expected amount of information acquired rises, despite the higher unit cost. This result

²We also show that a decrease in information costs can lower social surplus in our model. The result, however, arises in the isolated bilateral bargaining setting as well: if trade is always socially desirable and the two parties would trade without information then more (or, cheaper) information necessarily lowers social surplus.

distinguishes the precision of learning from the uncertainty available to be resolved, as well as from actual information expenditure. Once again, the market-wide effect cannot be understood by holding the individual buyer’s decision problem fixed.

1.1 Related Literature

Our paper contributes to the literature on dynamic adverse selection. We build on the stationary market of [Moreno and Wooders \(2010\)](#), in which decentralized trade mitigates the lemons problem without additional buyer information. [Kim \(2017\)](#) shows that market composition and endogenous outside options can alter the welfare effects of information about sellers’ time on the market. [Kaya and Kim \(2018\)](#) and [Daley and Green \(2012\)](#) study private buyer signals and public news, respectively. We endogenize buyers’ private experiments and examine how a common change in their costs affects the market.

The closest methodological antecedent is [Kim et al. \(2024\)](#), who study firms’ flexible information acquisition about a worker and show that cheaper learning has ambiguous effects on unemployment duration. Their informational link across firms operates through the worker’s observed duration. Here, histories are unobserved, and buyers are linked through the stationary pool of sellers and both sides’ continuation opportunities. This distinction allows us to examine whether buyers themselves benefit from cheaper information once market conditions adjust.

The market-composition channel is closely related to [Zou \(2025\)](#). In his model, expected resale illiquidity encourages buyers to acquire a given noisy signal; screening then removes high-quality assets relatively quickly and worsens the remaining pool. We share this feedback but abstract from resale: successful buyers exit, whereas unsuccessful buyers continue searching. Moreover, buyers choose the entire experiment rather than whether to acquire a fixed signal. Our focus is on stationary screening and the effects of information costs, rather than transitions into and out of information traps.

More broadly, the distinction between private and social returns to information is well established ([Hirshleifer, 1971](#)). [Glode et al. \(2019\)](#) study information acquisition as an arms race, and [Fishman and Parker \(2015\)](#) show how valuation and market prices can generate multiple equilibria. Bargaining models likewise demonstrate that information choices interact with the terms of trade ([Dang, 2008](#); [Chatterjee et al., 2025](#); [Ravid, 2020](#)). Our contribution is to study the additional feedback through the market that unsuccessful traders return to, with flexible screening, reservation prices, continuation values, and seller composition jointly determined. The central question is not merely whether information is socially valuable, but whether its private beneficiaries remain better off after a market-wide reduction in its cost.

The remainder of the paper is organized as follows. Section 2 presents the model. Sections 3

and 4 study the no-information and perfect-information benchmarks and their implementation. Section 5 characterizes imperfect-information equilibria, and Section 6 examines multiplicity. Section 7 studies the effects of declining information costs. Section 8 concludes.

2 Model

2.1 Physical Environment

Time is discrete and extends from $-\infty$ to ∞ . At the beginning of each period, a unit measure of buyers (he) and a unit measure of sellers (she) enter the market. Each seller owns an indivisible object of either high quality (H) or low quality (L). We will refer to a seller who owns a low/high **quality** object as low/high **type** seller. Let q_0 denote the fraction (measure) of high-type sellers in the entry population. A type- a object, $a \in \{H, L\}$, yields utility v_a to the buyer and entails production cost c_a (incurred upon trade) for the seller. We assume that higher quality is more valuable ($v_H > v_L$), more costly ($c_H > c_L$), and that there are gains from trade for both qualities ($v_a > c_a$). Seller type is privately known to sellers, while buyers are not informed, but can acquire information about it.

At the end of each period, each buyer is randomly matched with a seller. The buyer has no information about the seller's history, but can learn about the seller's type at costs defined shortly. Then, the buyer makes a price offer, which the seller can accept or reject. If the offer is accepted, trade occurs and both exit the market. If the offer is rejected, the buyer and seller remain in the market. All individuals' discount factor is given by $\delta \in (0, 1)$.

To focus on the most economically salient cases, we maintain the following three assumptions.³

Assumption 1 $\delta(c_H - c_L) > (v_L - c_L)$.

To understand this assumption, notice that v_L is the highest price a buyer (knowing the seller's type) could offer to a low-type seller, while c_H is the least price a buyer should offer to trade with a high-type seller. The assumption states that low-type sellers are sufficiently patient and, therefore, willing to reject v_L if they expect to receive c_H in the next period. Note that this assumption holds only when (so implies) $c_H > v_L$.

Assumption 2 $v_H - c_H \geq v_L - c_L$.

In other words, there are more gains from trade when the seller is of type h than when she is of type ℓ . This assumption allows us to reduce the number of cases to consider, and the subsequent analysis

³The subsequent analysis can be modified or extended to relax each of these assumptions in the straightforward manner. However, it would be at the significant cost of increasing the number of cases to consider.

and results can be adapted to the opposite case where $v_L - c_L \geq v_H - c_H$ in a straightforward manner.

Assumption 3 $\delta [(1 - q_0)(v_L - c_L) + q_0(v_H - c_H)] > (1 - q_0)(v_L - c_L)$.

This assumption ensures that sellers are sufficiently patient relative to the incoming measure of high-type sellers and, therefore, have a strong incentive to wait for c_H . Its economic meaning becomes more transparent in the following condition, which is sufficient for [Assumption 3](#) given [Assumption 2](#):

$$\delta [(1 - q_0)(v_L - c_L) + q_0(v_L - c_L)] > (1 - q_0)(v_L - c_L) \iff \delta + q_0 > 1.$$

2.2 Information Acquisition

2.2.1 Flexible Information Acquisition

Each buyer can learn about his matched seller with full flexibility. To be specific, let q^* denote the (equilibrium) proportion of high-type sellers in the market. Then, the buyer can choose *any* distribution F over posterior beliefs $q \in [0, 1]$ as long as it satisfies the following Bayes-plausibility constraint:

$$q^* = \int q dF(q). \tag{1}$$

Clearly, getting no information is equivalent to choosing a degenerate distribution at q^* , while perfect information corresponds to choosing a binary distribution over $\{0, 1\}$.

We assume that buyers face a uniformly posterior separable (UPS) cost function when acquiring information ([Caplin et al., 2022](#)). Given the equilibrium belief q^* , the cost of choosing a distribution F is:

$$C(F, q^*) = \lambda(\Phi(q^*) - \int \Phi(\mu) dF(\mu)) \tag{2}$$

where Φ is a strictly concave function over $[0, 1]$ and the coefficient λ measures the unit cost of information. We normalize $\Phi(0) = \Phi(1) = 0$ without loss and assume that Φ is twice continuously differentiable (with derivative ϕ) on $(0, 1)$. Two leading examples we will use are quadratic costs (when $\Phi(\mu) = \mu(1 - \mu)$) and Shannon entropy costs (when $\Phi(\mu) = -\mu \ln(\mu) - (1 - \mu) \ln(1 - \mu)$).

Each distribution F describes buyers' *unconditional* beliefs after information acquisition. From each seller's perspective, more important is buyers' beliefs *conditional on* her true type, that is, what beliefs buyers facing her will arrive at. For each $a \in \{H, L\}$, we let F_a denote the distribution

of their conditional beliefs that correspond to the distribution F . Given any F , F_L and F_H should satisfy the following two consistency conditions:

$$F(\mu) = q^* F_H(\mu) + (1 - q^*) F_L(\mu) \text{ and } \frac{\mu}{1 - \mu} = \frac{q^*}{1 - q^*} \frac{dF_H(x)}{dF_L(x)}.$$

This implies that F_L and F_H can be obtained from F as follows:

$$dF_H(\mu) = \frac{\mu}{q^*} dF(\mu) \text{ and } dF_L(\mu) = \frac{1 - \mu}{1 - q^*} dF(\mu).$$

2.3 Market Equilibrium

We study steady-state equilibria of this market. To simplify the definition of equilibrium, we begin with two general observations.

First, buyers never offer a price strictly above high-type sellers' reservation value, c_H . This follows from the standard Diamond paradox argument: Let $\bar{p}$ denote the highest price buyers would offer. Due to discounting, high-type sellers would be willing to accept any price above $(1 - \delta)c_H + \delta\bar{p}$. This suggests that $\bar{p} \leq (1 - \delta)c_H + \delta\bar{p}$, which is equivalent to $\bar{p} \leq c_H$. Note that this implies that the reservation price of high-type sellers is always equal to c_H , and they never earn positive surplus.

Second, letting p denote low-type sellers' reservation price,⁴ it incurs no loss of generality to assume that each seller accepts any price weakly above her reservation price, and buyers offer only c_H , p , or a losing price (below p).⁵ In what follows, we use $\sigma_H(\mu)$ to denote the probability that buyers offer c_H , and $\sigma_L(\mu)$ to denote the probability that they offer p , when his belief after information acquisition is equal to μ . We also let V_B denote buyers' market utility—their expected payoff when they enter, or return to, the market—and $u(\mu)$ represent their expected payoff when they face a seller who is the high type with probability μ (after their information acquisition). Since high-type sellers accept only c_H , while low-type sellers accept both p and c_H ,

$$u(\mu) = \max \{V_B, (1 - \mu)(v_L - p) + \mu V_B, \mu v_H + (1 - \mu)v_L - c_H\}. \quad (3)$$

Definition 1 A tuple $(q^*, p, V_B, F^*(\cdot), \sigma_H(\cdot), \sigma_L(\cdot))$ comprises a steady-state market equilibrium if the following conditions hold:

⁴Due to discounting and the fact that buyers never offer strictly more than c_H , $p < c_H$ necessarily holds, because even if a low-type seller will receive c_H for sure in the next period, her reservation price is $c_L + \delta(c_H - c_L) = (1 - \delta)c_L + \delta c_H < c_H$.

⁵Alternatively, one can assume that buyers would offer only c_H or p , but low-type sellers may not accept p with probability 1. It is straightforward to show that this alternative assumption leads to the same market outcome as in our main analysis once low-type sellers' rejecting p is replaced with (or interpreted as) buyers' offering a losing price.

(i) *Optimal offers: For each $\mu \in [0, 1]$, $\sigma_H(\mu) > 0$ only when*

$$\mu v_H + (1 - \mu)v_L - c_H \geq \max\{V_B, (1 - \mu)(v_L - p) + \mu V_B\},$$

and $\sigma_L(\mu) > 0$ only when

$$(1 - \mu)(v_L - p) + \mu V_B \geq \max\{V_B, \mu v_H + (1 - \mu)v_L - c_H\}.$$

(ii) *Optimal information acquisition: F^* is a solution to*

$$\max_F \int u(\mu) dF(\mu) - C(F) \text{ s.t. } q^* = \int \mu dF(\mu).$$

(iii) *Buyers' market utility:*

$$V_B = \delta \left(\int u(\mu) dF^*(\mu) - C(F^*) \right).$$

(iv) *Low-type sellers' reservation price:*

$$p = c_L + \delta \left(\int \sigma_H(\mu) dF_L^*(\mu)(c_H - c_L) + \int (1 - \sigma_H(\mu)) dF_L^*(\mu)(p - c_L) \right).$$

(v) *Steady-state condition:*

$$\frac{q_0}{1 - q_0} = \frac{q^*}{1 - q^*} \frac{\int \sigma_H(\mu) dF_H^*(\mu)}{\int (\sigma_H(\mu) + \sigma_L(\mu)) dF_L^*(\mu)}.$$

In **Definition 1**, (i) and (ii) capture the optimality of buyers' offer and information acquisition strategies, respectively. (iii) and (iv) mean that all players' market utilities should be computed recursively, by taking into account their continuation actions and resulting payoffs. Note that in (iii), $u(\mu)$ includes V_B (see (3)). In addition, in (iv), we use the fact that low-type sellers are indifferent between accepting and rejecting p (or receiving p and a losing price offer), so for the purpose of computing their expected payoff, we can assume that they only face c_H or p , not a losing price; observe that in (iv), it is as if low-type sellers receive offer p with $\int (1 - \sigma_H(\mu)) dF_L^*(\mu)$, instead of probability $\int \sigma_L(\mu) dF_L^*(\mu)$. Finally, (v) ensures that the distribution of sellers' types stays constant over time: the left-hand side represents the ratio of high-type sellers to low-type sellers in the entry population, while the right-hand side captures the corresponding ratio in the exiting population.

3 Equilibrium with No Information Acquisition

We first consider the benchmark case where buyers do not acquire information and then study the condition under which information acquisition is indeed irrelevant in our model.

3.1 Benchmark: Market with No Information Acquisition

In the absence of information acquisition, our model reduces to that of [Moreno and Wooders \(2010\)](#).⁶ By definition, buyers' strategies do not depend on their posterior beliefs (as there is no updating at all), so we use σ_H and σ_L , instead of $\sigma_H(\mu)$ and $\sigma_L(\mu)$, to represent buyers' offer strategies.

The following proposition fully describes the market outcome.

Proposition 1 *Suppose buyers do not acquire any information about sellers, and let $\hat{q} := (c_H - v_L)/(v_H - v_L)$. If $q_0 \leq \hat{q}$ then $q^* = \hat{q}$, $p = v_L$, $V_B = 0$, and σ_H and σ_L are such that*

$$v_L = c_L + \delta [\sigma_H(c_H - c_L) + (1 - \sigma_H)(v_L - c_L)] \text{ and } \frac{q_0}{1 - q_0} = \frac{\hat{q}}{1 - \hat{q}} \frac{\sigma_H}{\sigma_H + \sigma_L}.$$

If $q_0 > \hat{q}$ then $q^* = q_0$, $\sigma_H = 1$, $\sigma_L = 0$, $V_B > 0$, and $p = c_L + \delta(c_H - c_L) > v_L$.

If adverse selection is mild, in the sense that the proportion of high-type sellers in the entry population is relatively large, then all buyers offer c_H , so all sellers trade in their first match and the proportion of high-type sellers, q^* , coincides with that of the entry population, q_0 . In this case, buyers also obtain positive surplus, which amounts to $V_B = \delta [q_0 v_H + (1 - q_0) v_L - c_H]$.

If adverse selection is severe (i.e., $q_0 \leq \hat{q}$) then in equilibrium buyers must randomize between c_H and p : If they would offer only c_H then $q^* = q_0$, in which case they obtain strictly negative utility of $V_B = \delta [q_0 v_H + (1 - q_0) v_L - c_H] < 0$. If they would offer only p then, since high-type sellers never trade, q^* would be close to 1, in which case buyers would strictly prefer offering c_H to p , which is a contradiction. In addition, low-type sellers' reservation price, p , should coincide with v_L : If $p > v_L$ then buyers would never offer p , so trade takes place only at c_H . Then, $q^* = q_0$, in which case buyers have no incentive to offer c_H , as well. This means that low-type sellers would never trade, in which case $p = c_L$, which is a contradiction. If $p < v_L$, then buyers would never offer a losing price, as they strictly prefer trading and earning positive surplus earlier. But

⁶One minor difference is that matching and trade take place at the beginning of each period in [Moreno and Wooders \(2010\)](#), but at the end in our model. The only economic consequence of this difference is that buyers' market utility does not depend on whether they are new entrants or returning buyers who fail to trade with their matched sellers in our model; in [Moreno and Wooders \(2010\)](#), new entrants' market utility is higher than existing buyers', because of one free match-and-trade opportunity at the entry time. Our model can be interpreted as the discrete-time analogue to the continuous-time version of [Kim \(2017\)](#), in which new entrants should also wait to meet sellers.

then, low-type sellers always trade in their first match and exit fast, in which case the steady-state condition can no longer be sustained. Combining the above two, it also follows that $V_B = 0$ and $q^* = \hat{q}$.

3.2 Condition for No Information Acquisition

Next, we consider our model of information acquisition and identify the condition under which buyers choose not to acquire information (so the market outcome remains as in [Proposition 1](#)). Formally, we study an individual buyer's incentive to acquire information in the market equilibrium given in [Proposition 1](#). As explained in [Section 2](#), an individual buyer's optimal information acquisition strategy is given by

$$\max_F \int u(\mu) dF(\mu) - C(F) = \int (u(\mu) + \lambda\Phi(\mu)) dF(\mu) - \lambda\Phi(q^*) \text{ s.t. } q^* = \int \mu dF(\mu).$$

As is well known by now, this problem can be solved through the method of concavification ([Aumann et al., 1995](#); [Kamenica and Gentzkow, 2011](#)): Let $U(\mu) := u(\mu) + \lambda\Phi(\mu)$. The optimal value of the above problem at each prior q^* coincides with the value of the concave closure of U (ignoring the constant $-\lambda\Phi(q^*)$), and any optimal signal induces only those posteriors at which U meets the concave closure. A corollary of this is that no information (an uninformative signal) is optimal if and only if a tangent line of U at prior q^* stays uniformly above U .

3.2.1 Severe Adverse Selection

In [Proposition 1](#), if $q_0 \leq \hat{q}$ then $q^* = \hat{q}$ and $V_B = 0$. In this case, $u(\cdot)$ in (3) simplifies to

$$u(\mu) = \max \{0, \mu v_H + (1 - \mu)v_L - c_H\} = \begin{cases} 0 & \text{if } \mu \leq q^* \\ \mu v_H + (1 - \mu)v_L - c_H & \text{if } \mu > q^*. \end{cases}$$

In other words, $u(\cdot)$ is a piecewise linear function with an upward kink at q^* , as depicted in [Figure 1a](#). Combined with strict concavity of Φ , this implies that the function $U = u + \lambda\Phi$ has a shape depicted in [Figure 1b](#), with two humps, one below q^* and the other above q^* . The logic of concavification then implies that there exists a unique optimal signal that induces two posteriors, $\underline{\mu}$ and $\bar{\mu}$, such that the tangent lines at the two points coincide and stay uniformly above U . Of course, this implies that no information (equivalently, choosing the degenerate distribution at q^*) is never optimal, leading to the following result.

Proposition 2 *If $q_0 \leq \hat{q}$ then, irrespective of the value of λ , there does not exist an equilibrium in which buyers do not acquire information.*




Figure 1: Individual net payoff with and without information acquisition technology

The intuition behind this result is as follows: If $q_0 \leq \hat{q}$ then, as shown in [Proposition 1](#), buyers are indifferent between offering c_H and $p = v_L$. In this case, an arbitrarily small amount of information is beneficial to a buyer, as it enables him to break the tie between c_H and v_L . With flexible information acquisition, the buyer does not have to incur substantial information costs and so necessarily benefits from the information acquisition opportunity, regardless of the unit cost of information, λ .

3.2.2 Mild Adverse Selection

Next, consider the case where $q_0 > \hat{q}$. Absent information acquisition, $q^* = q_0$ and buyers always offer c_H (i.e., $\sigma_H = 1$), strictly preferring it to p and a losing price. Since $p = (1 - \delta)c_L + \delta c_H > v_L$ ([Assumption 1](#)), buyers never offer p , even if their belief falls to a very low value after information acquisition. Meanwhile, buyers' market utility is positive: $V_B = \delta [q^*v_H + (1 - q^*)v_L - c_H] > 0$. This suggests that $u(\cdot)$ becomes

$$u(\mu) := \max\{V_B, \mu v_H + (1 - \mu)v_L - c_H\} = \begin{cases} V_B & \text{if } \mu \leq \hat{\mu} \\ \mu v_H + (1 - \mu)v_L - c_H & \text{if } \mu > \hat{\mu} \end{cases}$$

where $\hat{\mu}$ is the value such that

$$V_B = \hat{\mu}v_H + (1 - \hat{\mu})v_L - c_H \iff \hat{\mu} := \frac{V_B + c_H - v_L}{v_H - v_L}.$$

The function $u(\cdot)$ is, again, piecewise linear with one upward kink. An important difference from the previous case is that the kink is not at q^* , but at $\hat{\mu}$ that is strictly smaller than q^* (because




Figure 2: Optimality of information acquisition.

V_B includes δ with q^*). This makes it possible that the tangent line of U at q^* stays uniformly above U , and thus acquiring no information is optimal for buyers. The following result shows that for any given $q_0 (> \hat{q})$, this is indeed the case if and only if λ exceeds a certain threshold.

Proposition 3 *If $q_0 > \hat{q}$ then there exists $\bar{\lambda}(> 0)$ such that there exists an equilibrium in which buyers do not acquire information if and only if $\lambda \geq \bar{\lambda}$.*

Intuitively, since $\hat{\mu} < q^*$, the buyer benefits from information acquisition only when he acquires a substantial amount of information. In particular, his information can change his offer (from c_H to a losing price) only when his belief becomes smaller than $\hat{\mu}$, whose associated information cost is bounded away from 0. Therefore, a buyer can benefit from information acquisition if and only if the unit cost of information is relatively small.

4 Equilibrium with Perfect Information Acquisition

This section studies another benchmark case where buyers obtain perfect information about sellers, which can be interpreted as the opposite case to the previous no-information case. As in [Section 3](#), we first characterize the market outcome given perfect information about sellers and then identify the condition under which the outcome can be sustained in our model (i.e., each individual buyer indeed has an incentive to obtain perfect information).

4.1 Market with (Costly) Perfect Information

Under perfect information, the Diamond paradox applies to each type, so buyers never offer more than c_a to type- a sellers. It then follows that the low-type seller's reservation price coincides with her reservation value (i.e., $p = c_L$) and each buyer can extract full trading surplus ($v_a - c_a$) from each type- a seller. **Assumption 2** then implies that a buyer, who can never obtain more than $v_H - c_H$, will always trade with a high-type seller. However, he may not always trade with a low-type seller, because he has an incentive to wait to meet a high-type seller.

To make this benchmark directly comparable to our model, we assume that a buyer should incur the cost of obtaining perfect information whenever he meets a seller. By our cost function specification, the cost amounts to

$$\lambda(\Phi(q^*) - (1 - q^*)\Phi(0) - q^*\Phi(1)) = \lambda\Phi(q^*),$$

where the equality is due to our normalization, namely, that $\Phi(0) = \Phi(1) = 0$.

Let $\bar{\delta}(< 1)$ denote the value such that

$$\begin{aligned} v_L - c_L &= \bar{\delta} [(1 - q_0)(v_L - c_L) + q_0(v_H - c_H) - \lambda\Phi(q^*)] \\ \iff \bar{\delta} &:= \frac{v_L - c_L}{(1 - q_0)(v_L - c_L) + q_0(v_H - c_H) - \lambda\Phi(q^*)}. \end{aligned}$$

In words, if $\bar{\delta}$ is a buyer's discount factor then he is just indifferent between trading with a low-type seller in the current period (so earning $v_L - c_L$) and waiting until the next period and meeting a high-type seller with probability q_0 . If $\delta \leq \bar{\delta}$ then a buyer would not be willing to forgo today's trading opportunity with a low-type seller to be able to meet a high-type seller in the next period, so all agents trade immediately. If $\delta > \bar{\delta}$ and $q^* = q_0$ then buyers would not be willing to trade with low-type sellers, which cannot be the case in steady state. In this case, q^* would be adjusted so as to provide a proper trading incentive for buyers.

The following proposition provides a comprehensive characterization for the market outcome under perfect information.

Proposition 4 *Suppose buyers perfectly learn about their matched sellers. If $\delta \leq \bar{\delta}$ then $q^* = q_0$ and buyers offer c_L to low-type sellers with probability 1. If $\delta > \bar{\delta}$ then q^* is the value such that*

$$v_L - c_L = \delta [(1 - q^*)(v_L - c_L) + q^*(v_H - c_H) - \lambda\Phi(q^*)]$$

and buyers offer c_L to low-type sellers with probability σ_L^* (and a losing price with probability

$1 - \sigma_L^*)$ where σ_L^* is the value such that

$$\frac{q_0}{1 - q_0} = \frac{q^*}{1 - q^*} \frac{1}{\sigma_L^*} \iff \sigma_L^* := \frac{1 - q_0}{q_0} \frac{q^*}{1 - q^*}.$$

In the former case where $\delta \leq \bar{\delta}$, each buyer trades with his first matched seller, regardless of her type. Therefore, buyers' market utility is given by

$$V_B = \delta [(1 - q_0)(v_L - c_L) + q_0(v_H - c_H) - \lambda\Phi(q^*)].$$

By contrast, in the latter case where $\delta > \bar{\delta}$, a buyer may not trade if he meets a low-type seller. In this case, buyers' market utility should equal $v_L - c_L$. To see why this is the case, suppose $V_B < v_L - c_L$, so that buyers prefer trading with a low-type seller to waiting until the next period. Then, all sellers trade immediately, so $q^* = q_0$. This yields a contradiction because

$$\delta > \bar{\delta} \iff V_B = \delta [(1 - q_0)(v_L - c_L) + q_0(v_H - c_H) - \lambda\Phi(q^*)] > v_L - c_L.$$

Next, suppose $V_B > v_L - c_L$. Then, low-type sellers never trade (exit), so the market would be swamped by them (i.e., $q^* = 0$ in steady state). But then, a buyer would have to wait arbitrarily long to meet a high-type seller, in which case his expected payoff would be close to 0 (strictly below $v_L - c_L > 0$), which is a contradiction.

4.2 Condition for Perfect Information Acquisition

We now consider an individual buyer's incentive and determine the condition under which he indeed wishes to learn perfectly about his matched seller.

Suppose a buyer's (posterior) belief about his matched seller is $\mu \in [0, 1]$. Under perfect information, type- a sellers' reservation price is equal to c_a . Therefore, if a buyer's (posterior) belief about his matched seller is $\mu \in [0, 1]$ then his indirect utility is given by

$$u(\mu) = \max\{V_B, \mu V_B + (1 - \mu)(v_L - c_L), \mu v_H + (1 - \mu)v_L - c_H\}.$$

Whether $V_B < v_L - c_L$ or $V_B = v_L - c_L$ (equivalently, whether $\delta < \bar{\delta}$ or not), u is a piecewise linear function with one upward kink. This implies that the same concavification logic as in [Section 3.2](#) directly applies to the current problem. Applying the fact that perfect information (a binary distribution with support $\{0, 1\}$) is optimal if and only if the chord connecting between $(0, u(0) + \lambda\Phi(0))$ and $(1, u(1) + \lambda\Phi(1))$ lies uniformly above $u + \lambda\Phi$ (i.e., the line should be the concave closure of $u + \lambda\Phi$ as seen in [Figure 3](#)), the following result ensues.




Figure 3: Perfect information acquisition is optimal.

Proposition 5 *There exists an equilibrium in which buyers obtain perfect information about sellers if and only if*

$$\begin{aligned}
 &-(v_L - c_L - V_B) + \lambda\phi(0) \leq v_H - c_H + \lambda\Phi(1) - (v_L - c_L + \lambda\Phi(0)) \leq v_H - v_L + \lambda\phi(1) \quad (4) \\
 &\iff \lambda\phi(0) \leq v_H - c_H - V_B \text{ and } \lambda(-\phi(1)) \leq c_H - c_L.
 \end{aligned}$$

Proof. Notice that the middle expression in (4) represents the slope of the chord that connects between $(0, u(0) + \lambda\Phi(0))$ and $(1, u(1) + \lambda\Phi(1))$. Therefore, the first inequality means that the slope weakly exceeds the slope of $u + \lambda\Phi$ at 0, while the second means that the slope is not larger than the slope of $u + \lambda\Phi$ at 1. Due to a two-hump shape of $u + \lambda\Phi$, this is a necessary and sufficient condition for the chord to uniformly stay above $u + \lambda\Phi$. ■

Notice that (4) always holds if $\lambda = 0$, because $V_B < v_H - c_H < v_H - v_L$. Of course, this just means that a perfect-information equilibrium exists if information is costless. **Proposition 5** shows that the equilibrium exists even when $\lambda > 0$. As is intuitive, (4) holds only when λ is sufficiently small: Since Φ is strictly concave, $\phi(0) > \phi(1)$. Therefore, (4) holds only when $\lambda(\phi(0) - \phi(1)) \leq v_H - c_L - V_B$.

It is also worth noting that (4) requires ϕ to be bounded: If not, $\phi(0) = \infty$ or $\phi(1) = -\infty$, so at least one inequality in (4) fails. In other words, if the information cost function satisfies a type of Inada condition (e.g., when Φ is the entropy function such that $\Phi(\mu) = -\mu \log(\mu) - (1 - \mu) \log(1 - \mu)$) then there never exists a perfect-information equilibrium.

5 Equilibrium with Imperfect Information Acquisition

This section considers a non-extreme, and more representative, case where buyers obtain some imperfect (but non-negligible) information about sellers. We focus on the case where buyers' beliefs stay in the interior $(0, 1)$, which we refer to as “interior screening equilibrium” hereafter, relegating the (one-sided boundary) cases where buyers' beliefs may become equal to 0 or 1 to [Appendix B](#).

Buyer optimality. By the usual optimal pricing logic, it suffices to consider the following three prices for buyers' price offers: a losing price (strictly below p), p (low-type sellers' reservation price), and c_H (high-type sellers' reservation price). Therefore, the buyer's indirect payoff as a function of his posterior belief μ is given by

$$u(\mu) := \max\{V_B, \mu V_B + (1 - \mu)(v_L - p), \mu v_H + (1 - \mu)v_L - c_H\}.$$

As before, an individual buyer's information acquisition problem is given by

$$\max_F \int (u(\mu) + \lambda\Phi(\mu)) dF(\mu) - \lambda\Phi(q^*) \quad \text{s.t.} \quad \int \mu dF(\mu) = q^*.$$

Again, $u + \lambda\Phi$ has a two-hump “piecewise concave” shape, which leads to the following result via the usual concavification argument.

Lemma 1 *In interior screening equilibrium, each buyer chooses a binary signal such that the resulting posteriors, $\underline{\mu}(< q^*)$ and $\bar{\mu}(> q^*)$, satisfy⁷*

$$u'(\underline{\mu}) + \lambda\phi(\underline{\mu}) = \frac{u(\bar{\mu}) + \lambda\Phi(\bar{\mu}) - (u(\underline{\mu}) + \lambda\Phi(\underline{\mu}))}{\bar{\mu} - \underline{\mu}} = u'(\bar{\mu}) + \lambda\phi(\bar{\mu}).$$

If the realized posterior is $\bar{\mu}$ then the buyer offers c_H (i.e., $\sigma_H(\bar{\mu}) = 1$). If the realized posterior is $\underline{\mu}$ then the buyer offers a losing price (i.e., $\sigma_L(\underline{\mu}) = \sigma_H(\underline{\mu}) = 0$) if $V_B > v_L - p$ and p (i.e., $\sigma_L(\underline{\mu}) = 1$) if $V_B < v_L - p$.

⁷Allowing for the possibility that $\underline{\mu} = 0$ or $\bar{\mu} = 1$, this condition can be written as

$$u'(\underline{\mu}) + \lambda\phi(\underline{\mu}) \leq \frac{u(\bar{\mu}) + \lambda\Phi(\bar{\mu}) - (u(\underline{\mu}) + \lambda\Phi(\underline{\mu}))}{\bar{\mu} - \underline{\mu}} \leq u'(\bar{\mu}) + \lambda\phi(\bar{\mu}),$$

with the first inequality holding with equality if $\underline{\mu} = 0$ and the second one holding with equality if $\bar{\mu} = 1$.

Buyers' market utility. Note that buyers' optimal strategy in **Lemma 1** (i.e., $\underline{\mu}$ and $\bar{\mu}$) depends on V_B (and q^*). **Lemma 1**, in turn, implies that V_B is determined as follows:

$$V_B = \delta \left[(1 - \tau) \max\{V_B, (1 - \underline{\mu})(v_L - p) + \underline{\mu}V_B\} + \tau(\bar{\mu}v_H + (1 - \bar{\mu})v_L - c_H) - C(\underline{\mu}, \bar{\mu}; q^*) \right], \quad (5)$$

where τ is the value that makes the two posteriors, $\underline{\mu}$ and $\bar{\mu}$, consistent with q^* , that is,

$$(1 - \tau)\underline{\mu} + \tau\bar{\mu} = q^* \iff \tau = \frac{q^* - \underline{\mu}}{\bar{\mu} - \underline{\mu}} \quad (6)$$

and $C(\underline{\mu}, \bar{\mu}; q^*)$ represents the cost of inducing a binary distribution over $\{\underline{\mu}, \bar{\mu}\}$ from q^* , that is,

$$C(\underline{\mu}, \bar{\mu}; q^*) = \lambda \left[\Phi(q^*) - (1 - \tau)\Phi(\underline{\mu}) - \tau\Phi(\bar{\mu}) \right].$$

Low-type sellers' reservation price. By the definition of her reservation price, a low-type seller is indifferent between accepting and rejecting p , and thus she obtains the same payoff whether her matched buyer offers p or a losing price. This means that for the purpose of deriving p , we can assume that a seller's (next) matched partner would offer either c_H or p (not a losing price), so she would definitely trade in the next period. It then suffices to compute the probability that she will be offered c_H .

By **Lemma 1**, a seller is offered c_H if and only if her matched buyer realizes posterior $\bar{\mu}$. Importantly, the probability of this event depends on the seller's underlying type, because a buyer is more likely to observe a good realization from a high-type seller than from a low-type seller. To be formal, let β_a denote the probability for type $a \in \{H, L\}$. Then, by Bayes' rule,

$$\underline{\mu} = \frac{q^*(1 - \beta_H)}{q^*(1 - \beta_H) + (1 - q^*)(1 - \beta_L)} \quad \text{and} \quad \bar{\mu} = \frac{q^*\beta_H}{q^*\beta_H + (1 - q^*)\beta_L},$$

which can be explicitly solved as

$$\beta_L = \frac{q^* - \underline{\mu}}{\bar{\mu} - \underline{\mu}} \frac{1 - \bar{\mu}}{1 - q^*} \quad \text{and} \quad \beta_H = \frac{q^* - \underline{\mu}}{\bar{\mu} - \underline{\mu}} \frac{\bar{\mu}}{q^*}. \quad (7)$$

Combining this with the above reservation price logic pins down low-type sellers' reservation price as follows:

$$p - c_L = \delta [\beta_L(c_H - c_L) + (1 - \beta_L)(p - c_L)] \iff p = c_L + \frac{\delta\beta_L}{1 - \delta(1 - \beta_L)}(c_H - c_L). \quad (8)$$

Steady-state condition. In each period, a high-type seller trades if, and only if, her matched buyer realizes posterior $\bar{\mu}$ and, therefore, offers c_H . Conditional on her type, this event occurs with probability β_H , where β_H is given as in (7). Meanwhile, a low-type seller trades not only when her matched partner offers c_H (following posterior $\bar{\mu}$), but also when he realizes posterior $\underline{\mu}$ and then offers p . For the market proportion of high-type sellers, q^* , to stay constant, the ratio of these two probabilities should equate the corresponding ratio of the incoming population. This leads to the following condition for steady-state equilibrium:

$$\frac{q^*}{1 - q^*} \frac{\beta_H}{\beta_L + (1 - \beta_L)\sigma_L(\underline{\mu})} = \frac{q_0}{1 - q_0}. \quad (9)$$

The steady-state equilibrium condition, (9), highlights the role of information acquisition in our model. Absent information acquisition, as shown in **Proposition 1**, in equilibrium buyers either offer only c_H (if $q_0 \geq \hat{q}$), or randomize among a losing price, $p = v_L$, and c_H (if $q_0 < \hat{q}$). This is no longer the case in our model. First, it is possible that $\sigma_L(\underline{\mu}) = 1$, so that buyers never offer a losing price. This can be sustained in our model because high-type sellers receive offer c_H more frequently than low-type sellers (i.e., $\beta_H > \beta_L$), and thus low-type sellers' immediate trade (always trading in their first match) does not necessarily drive up q^* to an unsustainably high level. Second, to the contrary, it is possible that $\sigma_L(\underline{\mu}) = 0$, in which case buyers offer either c_H (after realizing $\bar{\mu}$) or a losing price (after $\underline{\mu}$); note that in the no-information case, if $q_0 \geq \hat{q}$ then buyers *always* offer c_H , not randomizing between c_H and a losing price. Although trade occurs only at c_H , it does not imply $q^* = q_0$ in our model because, once again due to information acquisition, $\beta_H > \beta_L$, so $q^* < q_0$ even when $\sigma_L(\underline{\mu}) = 0$.

The following result puts together all results so far.

Proposition 6 *An interior screening equilibrium consists of a tuple $(V_B, p, q^*, \underline{\mu}, \bar{\mu}, \beta_L, \beta_H, \sigma_L(\underline{\mu}))$ that satisfies **Lemma 1**, (5), (7), (8), and (9).*

We conclude this section by illustrating how one can establish the existence of (informative) steady-state market equilibrium. We begin by taking $V_B \in [0, v_L - c_L]$. Then, all equilibrium conditions other than (5) can be used to determine the other equilibrium variables; in fact, this part can be generalized so as to accommodate the possibility that $\underline{\mu} = \bar{\mu} = q^*$ (no information), $\underline{\mu} = 0$, or $\bar{\mu} = 1$. Plugging those values into (5) produces an equilibrium value of V_B , denoted by $\Phi(V_B)$. Clearly, a market equilibrium corresponds to a fixed point of Φ , whose existence follows from the fact that Φ is continuous, $\Phi(0) \geq 0$, and $\Phi(v_L - c_L) \leq v_L - c_L$.




Figure 4: Equilibrium multiplicity under quadratic cost $\Phi(q) = q(1 - q)$. Parameters: $v_H = 4$, $v_L = 1$, $c_H = 2.5$, $c_L = 0$, $\delta = 0.95$, $q_0 = 0.4$. For $\lambda \in [\lambda_1, \lambda_2]$ (shaded) three equilibria coexist, shown in blue, orange, and green.

6 Equilibrium Multiplicity

Figure 4 shows how the two equilibrium values, q^* and p , depend on λ (the unit cost of information) for the case with quadratic information costs (i.e., $\Phi(\mu) = \mu(1 - \mu)$). Its most intriguing feature is that there may exist multiple equilibria; in the gray area, there are two or three equilibria.⁸

To understand this equilibrium multiplicity, suppose q^* is relatively large, that is, the market is mainly comprised of high-type sellers; this corresponds to the equilibrium that exist when λ is close to 0 in Figure 4. In this case, buyers would set $\bar{\mu}(> q^*)$ high, meaning that they would offer c_H only when they are sufficiently convinced that their matched partner is the high type. This suggests that p (low-type sellers' reservation price) would be low and close to c_L ; if $\bar{\mu} = 1$ (i.e., buyers offer c_H only when they are certain that the seller is the high type) then low-type sellers would never receive c_H , in which case $p = c_L$. This makes trading with a low-type seller attractive for buyers, so they surely offer p (instead of a losing price) after observing $\underline{\mu}$. Then, low-type sellers trade immediately, regardless of buyers' signals about them, which is necessary to sustain a high value of q^* .

Now, consider the opposite case where q^* is relatively small, which corresponds to the equilibria that exist when λ is sufficiently large in Figure 4. In this case, $\bar{\mu}$ also remains relatively small, which implies that low-type sellers face a significant chance of receiving offer c_H . Then, p would be significantly larger than c_L . In particular, it could be that $v_L - p \leq V_B$, so buyers would not necessarily offer p (after observing posterior $\underline{\mu}$). This slows down low-type sellers' trade, so q^* indeed stays low.

To put it another way, equilibrium multiplicity in our market model stems from the novel

⁸With quadratic costs, equilibrium multiplicity always arises for a non-negligible λ interval. One can obtain this result by first identifying the lowest value of λ with which an interior screening equilibrium exists and then proving that a second interior equilibrium exists for a slightly larger λ .

complementarity between buyers' screening and (low-type) sellers' trading incentives. The more precisely buyers screen sellers (offering c_H with a higher posterior), the less are low-type sellers willing to wait for a high price. Conversely, precise screening (offering c_H only with a sufficiently high posterior) is less costly and, therefore, more likely to arise when low-type sellers trade faster and so q^* is already higher.

The above complementarity also explains why equilibrium multiplicity is present only when λ takes an intermediate value. If λ is sufficiently small then buyers necessarily acquire precise information, and thus a low value of q^* cannot be sustained in equilibrium. Conversely, if λ is sufficiently large then rigorous screening becomes too expensive for buyers, so a high value of q^* cannot be sustained in equilibrium. Notice that these patterns are consistent with [Figure 4](#), namely, that only the equilibrium with $\bar{\mu} = 1$ exists when λ is below the gray area, while only the equilibrium with a low value of q^* (and a high value of p) exists when λ is above the gray area.

When there are multiple equilibria, buyers and low-type sellers have opposite preferences over them: buyers are better off, while low-type sellers are worse off, when q^* is higher. As illustrated above, when q^* is higher, buyers offer c_H to low-type sellers less frequently, which lowers low-type sellers' reservation price (equivalently, their market utility). Buyers benefit from higher q^* in two ways: (i) they are more confident (with higher $\bar{\mu}$) about the seller's type when they offer c_H ; (ii) they can offer a lower price when they observe $\underline{\mu}$ and offer p . These opposing preferences make overall welfare rankings ambiguous in general: when there are three equilibria (in the gray area of [Figure 4](#)), total surplus can be maximized by any of them (with top, middle, or bottom q^*).⁹

7 Effects of Declining Information Costs

This section examines the economic effects of declining information costs. In particular, we address how reducing the unit cost of information, λ , influences the market share of high-type sellers (equivalently, the severity of the lemons problem in the market), market surplus, and the total amount of information acquired.¹⁰

Market share of high-type sellers. [Figure 5](#) shows how the equilibrium market share of high-type sellers, q^* , varies according to λ .¹¹ Its most distinguishing feature is that q^* does not change

⁹Total surplus for each cohort can be defined to be $(1 - q_0)(p - c_L) + V_B$, because high-type sellers always obtain zero surplus and $p - c_L$ is an individual low-type seller's market utility by the definition of p .

¹⁰For the comparative statics figures that follow, we adjust parameters to narrow the multiplicity region, so that the behavior on welfare, market duration, and information acquisition can be seen clearly away from the region of coexisting equilibria.

¹¹[Figure 5](#) considers the case where adverse selection is severe, so $q^* > q_0$ in the no information benchmark (see [Proposition 1](#)). However, the same pattern emerges even if adverse selection is mild.




Figure 5: Equilibrium share of high-type sellers on the market over λ .¹²

in a simple monotone or quasi-concave fashion; and its complexity is not due to equilibrium multiplicity, because a similar pattern emerges even if there is always a unique equilibrium, as shown in the right panel for the case of entropy-based costs.

To understand the increasing-decreasing-increasing pattern in Figure 5, first consider the case where λ is sufficiently close to 0. Then, buyers obtain precise (almost perfect) information about sellers. Due to Assumption 2, buyers are more eager to trade with high-type sellers, and thus they may forgo trading opportunities with low-type sellers (i.e., $\sigma_L(\mu) < 1$). This renders high-type sellers trading faster than low-type sellers, in which case q^* stays low. In this case, an increase of λ raises q^* because it would reduce buyers' screening power and, therefore, makes it more costly to pass the trading opportunity with low-type sellers, as well as directly lowers high-type sellers' frequency to receive c_H .

If λ continues to rise then buyers would acquire less information and, in particular, choose a lower value of $\bar{\mu}$. Then, low-type sellers would receive c_H more frequently, which raises their reservation price p . This reduces buyers' incentives to trade with low-type sellers, slowing down their exit. Consequently, q^* may fall. Eventually, however, q^* can rise again. When λ is arbitrarily large (i.e., $\lambda = \infty$), buyers simply cannot screen sellers and, therefore, offer c_H at a common rate for both types of sellers. A decrease in λ relaxes this problem; high-type sellers would receive c_H relatively more frequently than low-type sellers, which has the effect of lowering q^* .

Market surplus. Figure 6 illustrates how market participants' market utilities (left) and total surplus (right) depend on λ . An increase in λ is always beneficial to low-type sellers, because it weakens buyers' screening power, thereby raising the probability that they (mistakenly) offer c_H to low-type sellers. Interestingly, it is not necessarily harmful to buyers; as shown in the left panel, if λ is sufficiently large then V_B strictly increases in λ . This is due to the market composition

¹²Figure 5a uses $v_L = 0.5$ and $q_0 = 0.7$; all other parameters are as in Figure 4. Figure 5b uses the benchmark parameters of Figure 4 under the Shannon entropy cost $\Phi(q) = -q \ln(q) - (1 - q) \ln(1 - q)$.




Figure 6: Buyer–seller surplus division (left) and total surplus (right) across information cost λ .¹³

effect illustrated above, namely, that the share of high-type sellers in the market rises when λ is sufficiently large. Clearly, this is another unique feature of our market model: in the standard bilateral trading setting, an increase in λ can never benefit the buyer.

The right panel of **Figure 6** highlights the subtle economic effects of information acquisition by showing that total surplus is minimized when λ is neither too small nor too large. If λ is sufficiently small then (low-type) sellers cannot enjoy large information rents, which contributes to market efficiency. If λ takes an intermediate value then buyers’ screening incentives collide with sellers’ information rents, which results in significant trade delay and lower market efficiency.

Duration on the market. **Figure 7** depicts how long sellers (left) and buyers (right) stay on the market. Mathematically, this is the reciprocal of the per-period probability of trade. Expected duration can be interpreted as a measure of market liquidity. As noted above, when information is cheap, buyers have an incentive to delay trade if a low posterior is realized. Therefore, both buyers and low-type sellers are expected to remain on the market. As information becomes more costly, this incentive fades and the market becomes more liquid. For sufficiently high costs of information, all agents remain on the market for longer, but for different reasons. Buyers screen less precisely, causing low-type sellers to want to delay trade in hopes of receiving a mistaken high price offer from buyers in the future. This also means that the high posterior is realized less frequently. This leads to fewer high price offers, slowing down trade for high-type sellers. Thus, both types of sellers and buyers trade more slowly as λ increases.

Total information acquired. Finally, we consider how buyers’ information acquisition depends on λ , which is depicted in **Figure 8**. It is not surprising that, as shown in the left panel, buyers’ information costs initially increase but eventually decrease: If λ is close to 0 then so are buyers’

¹³Figure 6 uses $v_L = 0.5$, $\delta = 0.6$, and $q_0 = 0.85$; all other parameters are as in **Figure 4**.

¹⁴Figure 7 uses $v_L = 0.5$ and $\delta = 0.8$; all other parameters are as in **Figure 4**.




Figure 7: Sellers' (left) and buyers' expected durations on the market.¹⁴




Figure 8: Effect on spending on information.¹⁵

information costs, regardless of the precision of buyers' information. If λ is sufficiently large then buyers simply choose to acquire little information. The right panel, however, shows that buyers choose to acquire more information as λ initially rises from 0. To see this, suppose λ is sufficiently small. In this case, buyers get almost perfect information, in which case the reduction in uncertainty is mainly determined by $\Phi(q^*)$. In addition, under [Assumption 2](#), q^* is relatively small and increasing in λ . Combined with the structure of Φ , this implies that the expected reduction in uncertainty increases in λ .

8 Conclusion

This paper analyzes a dynamic decentralized market for lemons in which buyers engage in flexible information acquisition subject to UPS information costs. We characterize steady-state market equilibria and show that complementarity between buyers' screening and sellers' trading incentives

¹⁵Figure 8 uses $v_L = 0.5$; all other parameters are as in [Figure 4](#).

can cause equilibrium multiplicity. We also demonstrate that a reduction in information costs can be harmful to buyers and does not necessarily lead to more information being produced.

A natural next step is to allow the cost of information to change over time. We have treated λ as fixed and studied how the market responds as it varies. In practice, information technology improves over time, so it would be important to study how the market changes as λ falls, focusing on a transition path rather than comparing stationary economies. If cheaper information can leave buyers worse off in steady-state comparisons, it is natural to ask whether, and how, this plays out as information technology improves in real time. This is particularly relevant in the multiple equilibria region. As λ changes, equilibria can appear or disappear, so the market may be forced from one equilibria to another. Characterizing the transition dynamics across equilibria is a direction for future work.

References

- Aumann, Robert J, Michael Maschler, and Richard E Stearns**, *Repeated games with incomplete information*, MIT press, 1995.
- Blouin, Max R and Roberto Serrano**, “A decentralized market with common values uncertainty: Non-steady states,” *The Review of Economic Studies*, 2001, 68 (2), 323–346.
- Camargo, Braz and Benjamin Lester**, “Trading dynamics in decentralized markets with adverse selection,” *Journal of Economic Theory*, 2014, 153, 534–568.
- Caplin, Andrew, Mark Dean, and John Leahy**, “Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy,” *Journal of Political Economy*, 2022, 130 (6), 1676–1715.
- Chatterjee, Kalyan, Miaomiao Dong, and Tetsuya Hoshino**, “Bargaining and Information Acquisition,” *American Economic Journal: Microeconomics*, August 2025, 17 (3), 75–102.
- Daley, Brendan and Brett Green**, “Waiting for News in the Market for lemons,” *Econometrica*, 2012, 80 (4), 1433–1504.
- Dang, Tri Vi**, “Bargaining with endogenous information,” *Journal of Economic Theory*, May 2008, 140 (1), 339–354.
- Fishman, Michael J and Jonathan A Parker**, “Valuation, adverse selection, and market collapses,” *The Review of Financial Studies*, 2015, 28 (9), 2575–2607.
- Glode, Vincent, Christian C Opp, and Xingtang Zhang**, “On the efficiency of long intermediation chains,” *Journal of Financial Intermediation*, 2019, 38, 11–18.

- Hirshleifer, Jack**, “Suppression of inventions,” *Journal of Political Economy*, 1971, 79 (2), 382–383.
- Hörner, Johannes and Nicolas Vieille**, “Public vs. private offers in the market for lemons,” *Econometrica*, 2009, 77 (1), 29–69.
- Janssen, Maarten CW and Santanu Roy**, “Dynamic trading in a durable good market with asymmetric information,” *International Economic Review*, 2002, 43 (1), 257–282.
- Kamenica, Emir and Matthew Gentzkow**, “Bayesian Persuasion,” *American Economic Review*, October 2011, 101 (6), 2590–2615.
- Kaya, Ayça and Kyungmin Kim**, “Trading Dynamics with Private Buyer Signals in the Market for Lemons,” *The Review of Economic Studies*, October 2018, 85 (4), 2318–2352.
- Kim, Jeong Ho (John), Kyungmin Kim, and Marilyn Pease**, “Unemployment Duration Under Flexible Information Acquisition,” *International Economic Review*, February 2024, 65 (1), 471–503.
- Kim, Kyungmin**, “Information about sellers’ past behavior in the market for lemons,” *Journal of Economic Theory*, May 2017, 169, 365–399.
- Moreno, Diego and John Wooders**, “Decentralized Trade Mitigates the Lemons Problem,” *International Economic Review*, May 2010, 51 (2), 383–399.
- and —, “Dynamic markets for lemons: Performance, liquidity, and policy intervention,” *Theoretical Economics*, 2016, 11 (2), 601–639.
- Ravid, Doron**, “Ultimatum Bargaining with Rational Inattention,” *The American Economic Review*, 2020, 110 (9), pp. 2948–2963.
- Zou, Junyuan**, “Information Traps in Over-the-Counter Markets,” 2025. Working paper, INSEAD. Revise and resubmit, *Journal of Finance*.

A Omitted Proofs

A.1 Proof of Proposition 3

Proof. For each $\lambda > 0$, let $\underline{q}$ be the largest value in $[0, \hat{\mu}]$ such that

$$\begin{aligned} u'(q_0) + \lambda\phi(q_0) &= \frac{u(q_0) + \lambda\Phi(q_0) - (u(\underline{q}) + \lambda\Phi(\underline{q}))}{q_0 - \underline{q}} \\ \iff c_H - v_L + V_B - \underline{q}(v_H - v_L) &= \lambda [(\Phi(q_0) - \Phi(\underline{q})) - \phi(q_0)(q_0 - \underline{q})]. \end{aligned}$$

Notice that the left-hand side is affine decreasing and reaches 0 when $\underline{q} = \hat{\mu}$. The right-hand side is always positive and strictly convex in $\underline{q}$, both because Φ is strictly concave. Finally, clearly, the equality necessarily holds for some $\underline{q} (< \hat{\mu})$ if λ is sufficiently small, while it does not (i.e., the left-hand side is smaller than the right-hand side for all $\underline{q} \leq \hat{\mu}$) if λ is sufficiently large. This implies that there exists $\bar{\lambda} (> 0)$ such that $\underline{q}$ is well defined if and only if $\lambda \leq \bar{\lambda}$. Observe that, by construction, if $\lambda = \bar{\lambda}$ then

$$\left[u(\underline{q}) + \lambda\Phi(\underline{q}) \right]' = \lambda\phi(\underline{q}) = u'(q_0) + \lambda\phi(q_0) = \frac{u(q_0) + \lambda\Phi(q_0) - (u(\underline{q}) + \lambda\Phi(\underline{q}))}{q_0 - \underline{q}}.$$

The desired result then follows by the concavification argument as in [Section 3.2.1](#). If $\lambda \geq \bar{\lambda}$ then the tangent line at q_0 stays uniformly above $u + \lambda\Phi$, which implies that a buyer has no incentive to acquire information (see [Figure 2a](#)). To the contrary, if $\lambda < \bar{\lambda}$ then $u(q_0) + \lambda\Phi(q_0)$ is strictly below the value of the concave envelope of $u + \lambda\Phi$ at q_0 (see [Figure 2b](#)). In this case, it is optimal for a buyer to adopt a binary signal. ■

A.2 Proof of Proposition 4

Proof. By [Assumption 2](#), trading with a high-type seller yields the largest available surplus, so buyers always trade with high types. It remains to determine σ_L^* and q^* .

Case 1: $\delta \leq \bar{\delta}$. At $q^* = q_0$ a buyer weakly prefers trading with a low-type seller now to waiting, so $\sigma_L^* = 1$. Both types then trade in their first match, and the stock composition equals the entry composition, $q^* = q_0$.

Case 2: $\delta > \bar{\delta}$. At $q^* = q_0$ a buyer strictly prefers to wait, so low types cannot all trade in equilibrium (they would otherwise accumulate without bound). Thus $\sigma_L^* \in (0, 1)$, which requires

the buyer to be indifferent between trading and waiting at the equilibrium composition:

$$v_L - c_L = \delta[(1 - q^*)(v_L - c_L) + q^*(v_H - c_H) - \lambda\Phi(q^*)],$$

which determines q^* . Under perfect information $\beta_H = 1$ and $\beta_L = 0$, so the steady-state condition reduces to

$$\frac{q^*}{1 - q^*} \frac{1}{\sigma_L^*} = \frac{q_0}{1 - q_0} \iff \sigma_L^* = \frac{1 - q_0}{q_0} \frac{q^*}{1 - q^*}.$$

■

B Equilibrium Conditions for One-Sided Corner Cases

We now set up two intermediate cases where one posterior is at the corner and the other is interior. This means that the concave closure touches $U(\mu)$ at one boundary point and is tangent at another point. This yields a tangency equality at the interior posterior, and a slope inequality at the corner.

B.1 Upper corner: $\underline{\mu} \in (0, 1), \bar{\mu} = 1$

In this case, the buyer offers c_H only when certain that the seller is high type (the buyer is overly cautious in offering the high price). Since the low-type sellers will never receive the offer c_H , their reservation value equals c_L . At the interior posterior $\underline{\mu}$, the buyer offers $p = c_L$, so low-type sellers trade immediately. This means that high-type sellers accumulate on the market, driving $q^* > q_0$. In principle, at $\underline{\mu}$, the buyer can offer c_L (so the low-type seller trades immediately) or a losing price (so no trade occurs and the buyer waits). This pricing choice depends on $v_L - c_L$ and V_B . In any steady-state equilibrium, $V_B \leq v_L - c_L$. If $V_B > v_L - c_L$, the buyer would always refuse to trade with low types, who, never receiving c_H , never trade and accumulate without bound, driving the market composition to all low types, which causes V_B to fall below $v_L - c_L$, causing a contradiction. So, we analyze the two cases that can arise: $V_B < v_L - c_L$, where the buyer strictly prefers to offer c_L , and $V_B = v_L - c_L$, where the buyer is indifferent and may mix between c_L and a losing price.

Case 1: $V_B < v_L - c_L$. At posterior $\underline{\mu}$, $u(\mu) = (1 - \mu)(v_L - c_L) + \mu V_B$, meaning $u' = V_B - (v_L - c_L)$. At the upper corner $\bar{\mu} = 1$, the slope is given by $v_H - v_L$. The slope of the chord

is $\frac{U(1)-U(\underline{\mu})}{1-\underline{\mu}}$. This means that the two optimality conditions are:

$$V_B - (v_L - c_L) + \lambda\phi(\underline{\mu}) = \frac{U(1) - U(\underline{\mu})}{1 - \underline{\mu}} \quad (\text{tangency at } \underline{\mu}), \quad (10)$$

$$\frac{U(1) - U(\underline{\mu})}{1 - \underline{\mu}} \leq v_H - v_L + \lambda\phi(1) \quad (\text{corner condition at } 1). \quad (11)$$

We know that $U(1) = v_H - c_H + \lambda\Phi(1)$ and $U(\underline{\mu}) = (1 - \underline{\mu})(v_L - c_L) + \underline{\mu}V_B + \lambda\Phi(\underline{\mu})$. This means that (10) can be written as

$$\lambda[\Phi(\underline{\mu}) + (1 - \underline{\mu})\phi(\underline{\mu}) - \Phi(1)] = v_H - c_H - V_B. \quad (12)$$

The LHS is the gap at $\mu = 1$ between Φ and its tangent line at $\underline{\mu}$, which is positive by strict concavity. The RHS is the payoff wedge between the high-type trade surplus $v_H - c_H$ and the market utility V_B .

From Bayes plausibility, we know that $\tau = \frac{q^* - \underline{\mu}}{1 - \underline{\mu}}$. We can use this in the steady state equation as well as the buyer's market utility:

$$\frac{q_0}{1 - q_0} = \frac{q^* - \underline{\mu}}{(1 - q^*)(1 - \underline{\mu})} \quad (13)$$

$$V_B = \delta \left(\left(\frac{1 - q^*}{1 - \underline{\mu}} \right) [(1 - \underline{\mu})(v_L - c_L) + \underline{\mu}V_B] + \left(\frac{q^* - \underline{\mu}}{1 - \underline{\mu}} \right) (v_H - c_H) - C^* \right), \quad (14)$$

where $C^* = \lambda[\Phi(q^*) - (1 - \tau)\Phi(\underline{\mu})]$.

Finally, substituting the tangency value of the chord slope into (11), we have the following feasibility condition

$$\lambda[\phi(\underline{\mu}) - \phi(1)] \leq v_H - c_L - V_B. \quad (15)$$

An upper-corner equilibrium is a triple $(\underline{\mu}, q^*, V_B)$ satisfying (12)–(14), with $\underline{\mu} \in (0, 1)$ and $q^* \in (q_0, 1)$, provided that (15) holds and $V_B < v_L - c_L$. Note that the seller condition V_L disappears because $\beta_L = 0$, which means that $p - c_L = \delta[0 \cdot (c_H - c_L) + 1(p - c_L)] \Rightarrow p = c_L$ - it just gives us the pricing condition for the low-type seller.

Case 2: $V_B = v_L - c_L$. In this case, the buyer mixes at $\underline{\mu}$, offering price c_L with probability σ_L and a losing price with probability $1 - \sigma_L$. If the buyer offers c_L , their expected payoff is $(1 - \underline{\mu})(v_L - c_L) + \underline{\mu}V_B$. Otherwise, the losing price offer results in no trade, and the buyer's

expected payoff is just V_B . Setting the two equal simplifies to

$$V_B = v_L - c_L. \quad (16)$$

The tangency condition follows from Case 1, but simplifies to

$$\lambda\phi(\underline{\mu}) = \frac{U(1) - U(\underline{\mu})}{1 - \underline{\mu}}. \quad (17)$$

Substituting $U(1)$ and $U(\underline{\mu})$ gives

$$\lambda[\Phi(\underline{\mu}) + (1 - \underline{\mu})\phi(\underline{\mu}) - \Phi(1)] = v_H - c_H - (v_L - c_L). \quad (18)$$

Notice that the RHS is the difference between the high-type and low-type trade surpluses. This means that the buyer acquires information up to the point where the marginal information cost matches the surplus gain from distinguishing the two types of sellers.

Next, the buyer's market utility is given by

$$V_B = v_L - c_L = \delta \left[\tau(v_H - c_H) + (1 - \tau)V_B - C^* \right], \quad (19)$$

where $\tau = \frac{q^* - \underline{\mu}}{1 - \underline{\mu}}$ and $C^* = \lambda[\Phi(q^*) - (1 - \tau)\Phi(\underline{\mu})]$. Since the buyer is indifferent, both offers at $\underline{\mu}$ yield V_B , so the mixing probability σ_L does not appear here. Instead, it is pinned in the steady state condition below:

$$\frac{q_0}{1 - q_0} = \frac{\beta_H}{\sigma_L} \frac{q^*}{1 - q^*} \Rightarrow \sigma_L = \beta_H \frac{q^*}{1 - q^*} \frac{1 - q_0}{q_0}, \quad (20)$$

where $\beta_H = \frac{q^* - \underline{\mu}}{q^*(1 - \underline{\mu})}$. For $\sigma_L \leq 1$, we need $\frac{q^* - \underline{\mu}}{(1 - q^*)(1 - \underline{\mu})} \leq \frac{q_0}{1 - q_0}$, meaning that the high-type exit flow odds ratio must not exceed the entrant ratio of high types to low types.

Substituting the tangency chord slope into the corner condition (11) and using $V_B = v_L - c_L$, the feasibility condition becomes

$$\lambda[\phi(\underline{\mu}) - \phi(1)] \leq v_H - v_L. \quad (21)$$

An upper-corner equilibrium is a tuple $(\underline{\mu}, q^*, V_B, \sigma_L)$ satisfying (18)–(20), with $\underline{\mu} \in (0, 1)$ and $q^* \in (0, q_0)$, provided that (21) holds.

B.2 Lower corner: $\underline{\mu} = 0, \bar{\mu} \in (0, 1)$

In this case, the low-type seller is perfectly identified. Probability of receiving low signal upon a match with high type equals 0. On the high signal $\bar{\mu}$, the buyer offers price c_H . Therefore, high types only face the offers c_H and all of them exit the market. On the low signal, the buyer offers price $p > c_L$. The price exceeds c_L because by waiting, the low-type seller has a positive probability of being offered $c_H > p$ in a future period. We again look at different cases that depend on $v_L - p$ and V_B .

Case 1: $V_B < v_L - p$. Although the buyer captures less on low types, they still offer p with probability 1 when they receive signal $\underline{\mu}$ (since $V_B < v_L - p$), which means that $\sigma_L = 1$. This leads to $q_0 = q^*$, meaning that the equilibrium composition of high-type sellers in the market is equal to the entrant proportion.

At the lower corner, $u(0) = v_L - p$. At the upper posterior, the slope is given by $u'(\mu) = v_H - v_L$. The chord slope is given by $\frac{U(\bar{\mu}) - U(0)}{\bar{\mu}}$, where $U(0) = v_L - p + \lambda\Phi(0)$. Thus, the two optimality conditions are

$$V_B - (v_L - p) + \lambda\phi(0) \leq \frac{U(\bar{\mu}) - (v_L - p) - \lambda\Phi(0)}{\bar{\mu}} \quad (\text{corner condition at } 0), \quad (22)$$

$$v_H - v_L + \lambda\phi(\bar{\mu}) = \frac{U(\bar{\mu}) - (v_L - p) - \lambda\Phi(0)}{\bar{\mu}} \quad (\text{tangency at } \bar{\mu}). \quad (23)$$

Rearranging (23) and using $u(\bar{\mu}) = \bar{\mu}v_H + (1 - \bar{\mu})v_L - c_H$, the interior posterior solves

$$\lambda[\Phi(\bar{\mu}) - \bar{\mu}\phi(\bar{\mu}) - \Phi(0)] = c_H - p. \quad (24)$$

Notice that the LHS is λ times the gap at $\mu = 0$ between Φ and its tangent line at $\bar{\mu}$. This gap is positive by strict concavity, and the tangency condition sets it equal to $c_H - p$, the gap between the high type's price and the low type's reservation price.

From Bayes plausibility, we know that $\tau = \frac{q^*}{\mu}$. The steady state condition gives $q^* = q_0$ and buyer's market utility is given by

$$V_B = \delta[\tau(\bar{\mu}v_H + (1 - \bar{\mu})v_L - c_H) + (1 - \tau)(v_L - p) - C^*], \quad (25)$$

where $C^* = \lambda[\Phi(q_0) - \tau\Phi(\bar{\mu})]$.

In this case, the seller's market utility is given by

$$V_L = p - c_L = \delta[\beta_L(c_H - c_L) + (1 - \beta_L)(p - c_L)]$$

Solving for p explicitly gives

$$p = c_L + \frac{\delta\beta_L}{1 - \delta(1 - \beta_L)}(c_H - c_L). \quad (26)$$

Finally, substituting the tangency value of the chord slope into (23), we have the following feasibility condition

$$\lambda[\phi(0) - \phi(\bar{\mu})] \leq v_H - p - V_B. \quad (27)$$

The lower-corner equilibrium is characterized by $(\bar{\mu}, V_B, p)$ satisfying (24)–(26), given $\bar{\mu} \in (q_0, 1)$ and $q^* = q_0$, provided that (27) and $V_B < v_L - p$ holds. Note that $\bar{\mu} > q_0$ is required for the high signal to be informative. So when q_0 is close to 1, this leaves little room for information, and the corner may not exist. In addition, when q_0 is high, this means that there are few low types in the pool of sellers. This means that perfectly identifying them (driving $\underline{\mu} \rightarrow 0$) is costly and yields little benefit, which tightens (27) and the corner can again cease to be optimal.

Case 2: $V_B = v_L - p$. In this case, the buyer is indifferent between trading at $p > c_L$ and waiting for a new match (potentially with a high type) next period. Thus, the buyer mixes at $\underline{\mu}$, offering price $p > c_L$ with probability σ_L and losing price with probability $1 - \sigma_L$.

The indifference condition is given by $V_B = v_L - p$, which comes from the buyer indifference between offering p and the losing price. The tangency condition is unchanged from Case 1 (24) as is the reservation price (26) and the buyer's market utility (25). However, since the low-type sellers do not trade with probability $1 - \sigma_L$, they exit more slowly, so $q^* < q_0$. The steady state condition then pins down σ_L

$$\sigma_L = \frac{1}{1 - \beta_L} \left[\frac{q^*}{1 - q^*} \frac{1 - q_0}{q_0} - \beta_L \right]. \quad (28)$$

Substituting the tangency chord slope into the corner condition and using $V_B = v_L - p$, the feasibility condition simplifies to

$$\lambda[\phi(0) - \phi(\bar{\mu})] \leq v_H - v_L. \quad (29)$$

The lower-corner equilibrium is a tuple $(\bar{\mu}, q^*, V_B, p, \sigma_L)$ satisfying (24)–(26) and (28), with $q^* \leq q_0 \leq \bar{\mu}$ and $\sigma_L \in (0, 1)$, provided that (29) and indifference $V_B = v_L - p$ holds.

Case 3: $V_B > v_L - p$. When $V_B > v_L - p$, the buyer never offers reservation price p ($\sigma_L = 0$) when receiving low signal. High types always generate high signal, but the low-type sellers can

also generate high signal with probability $\beta_L = \frac{q^*}{\bar{\mu}} \frac{1-\bar{\mu}}{1-q^*}$, so there is a chance for them to trade, just much slower than the high types. Thus in equilibrium we must have $q^* < q_0$. Also, as trade only happens at c_H , in steady state we have $\bar{\mu} = q_0$. The steady state condition under this scenario, therefore, is

$$\frac{q_0}{1-q_0} = \frac{q^*}{1-q^*} \frac{1}{\beta_L} \quad (30)$$

Using $\bar{\mu} = q_0$ and $u(0) = V_B$, the tangency condition for the interior posterior is

$$v_H - v_L + \lambda\phi(q_0) = \frac{q_0v_H + (1-q_0)v_L - c_H + \lambda\Phi(q_0) - V_B - \lambda\Phi(0)}{q_0} \quad (31)$$

while the inequality condition at the corner posterior is

$$\lambda(\phi(0) - \phi(q_0)) \leq v_H - v_L \quad (32)$$

Using $\bar{\mu} = q_0$, buyer's market utility equals

$$V_B = \delta \left[\tau(q_0v_H + (1-q_0)v_L - c_H) + (1-\tau)V_B - C^* \right] \quad (33)$$

where $\tau = \frac{q^*}{q_0}$ and $C^* = \lambda \left[\Phi(q^*) - \tau\Phi(q_0) - (1-\tau)\Phi(0) \right]$.

Equilibrium in this case is the pair (q^*, V_B) satisfying (31) and (33). Moreover, the primitives of the model must satisfy (32). Finally, the reservation price of low-type sellers must satisfy the inequality condition characteristic to this case. That is $p = c_L + \frac{\delta\beta_L}{1-\delta(1-\beta_L)}(c_H - c_L)$ must be such that $V_B > v_L - p$ holds.

B.3 Closed-form conditions under quadratic costs

When the information cost is quadratic, $\Phi(\mu) = \mu(1-\mu)$ (so $\phi(\mu) = 1-2\mu$), the tangency conditions solve in closed form, and the resulting posteriors, together with the feasibility conditions, yield explicit two-sided windows on the information cost λ . These conditions are necessary for the respective corner equilibria; global optimality of the induced signal imposes further restrictions.

Upper corner. When $V_B < v_L - c_L$, the tangency (12) reduces to $\lambda(1-\underline{\mu})^2 = v_H - c_H - V_B$, so

$$\underline{\mu} = 1 - \sqrt{\frac{v_H - c_H - V_B}{\lambda}}. \quad (34)$$

Requiring $\underline{\mu} \in (0, 1)$ gives the lower bound $\lambda > v_H - c_H - V_B$, while the feasibility condition (15), which under quadratic costs reads $2\sqrt{\lambda(v_H - c_H - V_B)} \leq v_H - c_L - V_B$, gives an upper bound. Together,

$$v_H - c_H - V_B < \lambda \leq \frac{(v_H - c_L - V_B)^2}{4(v_H - c_H - V_B)}, \quad (35)$$

which is nonempty if and only if $v_H - c_H - V_B < c_H - c_L$.

In the indifference case ($V_B = v_L - c_L$), (34) becomes

$$\underline{\mu} = 1 - \sqrt{\frac{(v_H - c_H) - (v_L - c_L)}{\lambda}}, \quad (36)$$

so information must be costly relative to the surplus gain from screening: $\lambda > (v_H - c_H) - (v_L - c_L)$.

Lower corner. When $V_B < v_L - p$, the tangency (24) reduces to $\lambda \bar{\mu}^2 = c_H - p$, so

$$\bar{\mu} = \sqrt{\frac{c_H - p}{\lambda}}. \quad (37)$$

Combining this with the reservation price (26) and requiring $\bar{\mu} \in (q_0, 1)$, a lower-corner equilibrium requires

$$c_H - c_L < \lambda < \frac{(1 - \delta)(c_H - c_L)}{q_0^2}, \quad (38)$$

which is nonempty if and only if

$$q_0 < \sqrt{1 - \delta}. \quad (39)$$

The lower corner can thus arise only when the pool is not too high-type-heavy; as buyers grow more patient ($\delta \rightarrow 1$), the threshold $\sqrt{1 - \delta}$ shrinks, and the corner survives for a narrower range of entrant compositions.

When $V_B = v_L - p$, the tangency is unchanged from Case 1, giving $\bar{\mu} = \sqrt{(c_H - p)/\lambda}$, and the feasibility condition (29) reduces, using $\phi(0) - \phi(\bar{\mu}) = 2\bar{\mu}$, to

$$2\lambda\bar{\mu} \leq v_H - v_L, \quad \text{i.e.} \quad \bar{\mu} \leq \frac{v_H - v_L}{2\lambda}. \quad (40)$$

Here, $\bar{\mu}$ remains a free root coupled to the reservation price p and the composition $q^* < q_0$; the tangency and feasibility take closed form, but $(\bar{\mu}, p, q^*, \sigma_L)$ are jointly determined by the remaining conditions.

In the case where $V_B > v_L - p$, $\bar{\mu} = q_0$ is fixed and the tangency (31) solves linearly for the buyer's market utility:

$$V_B = v_L - c_H + \lambda q_0^2. \quad (41)$$

The corner condition (32) reduces, using $\phi(0) - \phi(q_0) = 2q_0$, to a single upper bound on the information cost:

$$\lambda \leq \frac{v_H - v_L}{2q_0}. \quad (42)$$

Given V_B , the recursion (33) determines q^* , and the defining inequality of this case, $V_B > v_L - p$ with p the reservation price (26), must hold. Since $\bar{\mu}$ is pinned at the prior, Case 3 admits the most explicit characterization of the three: the buyer's value and the feasibility bound are given in closed form directly, without solving for an interior posterior.

In the strict cases of both corners, the quadratic specification delivers a two-sided window on λ : information must be costly enough that the buyer does not over-acquire (pushing the interior posterior to the boundary), yet not so costly that the corner condition fails. The mixing and boundary case instead yield one-sided bounds, with the remaining equilibrium objects jointly determined. In all cases, these conditions are necessary, but not sufficient, as the induced signal must additionally be a global optimum of the buyer's problem.