Bargaining and Information Acquisition[*]

Kalyan Chatterjee[†] Miaomiao Dong[‡] Tetsuya Hoshino[§]

September 4, 2024

Abstract

We consider an ultimatum game where the value of the object being sold to the buyer is either high or low. The seller knows the value but the buyer does not. The value to the seller is zero. We introduce the option for the buyer to acquire costly information after an offer is made. This information either confirms the high value or provides no information. As the cost of information vanishes, the buyer gets all the surplus in a refinement of perfect Bayesian equilibrium although the option to acquire information is never used. Moreover, this signal structure is optimal for the buyer.

Keywords: Bargaining, asymmetric information, information acquisition.

JEL codes: C72, C78, D82, D83, L15.

---

[*] We thank Jack Fanning, Maxim Ivanov, and the participants at the Birmingham conference on economic theory, the 33rd Stony Brook International Conference on Game Theory, and the 22nd SAET conference for their helpful comments. We also thank the editor, Alexander Wolitzky, and anonymous referees for their valuable suggestions. All remaining errors are our own.

[†] Pennsylvania State University. Email: `kchatterjee@psu.edu`

[‡] Pennsylvania State University. Email: `miaomiao.dong@psu.edu`

[§] ITAM. Email: `tetsuya.hoshino@itam.mx`

1 Introduction

One of the main questions that a bargaining model seeks to address is how the bargaining surplus is divided (if utilities are transferable). Confining ourselves to bilateral bargaining between a seller and a buyer, there is a substantial body of literature on the factors that determine a player's share of surplus. One factor is the power to make proposals, as demonstrated by the ultimatum bargaining game. In this game, the seller makes a take-it-or-leave-it offer of a surplus division, and the buyer either accepts or rejects it. If no agreement is reached, both parties receive zero payoffs. In the unique subgame perfect equilibrium of this game, the seller gains the full surplus. If the buyer makes an offer instead, then the result is reversed, and the buyer receives the full surplus.[1] Another factor in determining the division of surplus is the presence of private information. In a dynamic bargaining game, if the buyer is privately informed about her valuation of the object but the seller is uninformed about it, the Coase Conjecture comes into play. Even if the seller makes all offers, as the frequency of making offers increases, the seller's commitment power vanishes and the buyer captures all the surplus in the limit.

We study how a different factor—the buyer's access to information—affects ultimatum bargaining over a single indivisible object. The object's value to the buyer, or its quality, is either high (H) or low (L), with $H > L > 0$.[2] The value to the seller is 0. The seller (he) has *both* the power to make a take-it-or-leave-it offer and private information about the quality, while the buyer (she) is *initially uninformed* about the quality but has *cheap access to information*. Specifically, she has the option to acquire information about the quality after observing the seller's offer, with a cost depending on the accuracy of the information. The information takes two possible signal realizations: One conclusively reveals the high quality, while the other provides no information. She is either informed of the value being high or is uninformed. We call this information structure H-focused.[3,4] Neither the choice of accuracy nor the signal realization is observed by the seller. Later, we will justify the H-focused signal structure.

---

[1] In a dynamic bargaining game with time discounting, if the seller makes all the offers, the buyer gains no share in any subgame perfect equilibrium. In alternating-offer bargaining games, two players alternate in making offers. If their discount factors are equal, then in every period, the player proposing the offer gains a greater share than the player responding to the offer; as the discount factor tends to one, their division will be equal (Rubinstein, 1982). If a player is randomly selected as a proposer (rather than alternately), then the division of surplus depends on the players' (subjective) probabilities of being recognized as a proposer (Yildiz, 2003).

[2] Our results remain true for any finite number of qualities.

[3] Feinberg and Skrzypacz (2005) assume H-focused information. Their information structure is exogenous, but ours is endogenous.

[4] The accuracy of the H-focused information is the probability that the buyer learns a high signal conditional on the object's value being high.

An example we have in mind is selling a used car. The seller knows his car's quality, which is the value to the buyer. If he does not sell the car, he discards it and gains payoff zero regardless of the quality. After observing the offer by the seller, the buyer may check the car by hiring a mechanic or test-drive it to evaluate the quality.[5]

We are interested in the limit equilibrium as the buyer's cost of information vanishes—that is, she has cheap access to information. To put our result in perspective, we consider the following two benchmark cases. (i) In the full-information benchmark where the object's quality is common knowledge between the seller and the buyer, the buyer must get zero surplus in the unique equilibrium. (ii) In the null-information benchmark where the buyer has no option to acquire information, there is a continuum of equilibria, where the buyer's surplus ranges from zero to the maximum possible.

Our main result—formally stated as Theorem 1—is as follows:

> In *every* equilibrium satisfying a reasonable requirement, discussed below, in the limit as the buyer's cost of information vanishes, the seller offers the lowest price regardless of the object's quality and the buyer gains full surplus; moreover, on the equilibrium path, the buyer never acquires any information.

Our result says that the buyer's having cheap access to information gives the buyer the full surplus and thus reverses the existing results of surplus division. Neither the seller's sole power to make a (take-it-or-leave-it) offer nor the presence of his private information about the object's value gives him any of the surplus.

The main takeaway (from our stylized model and restriction to pure learning strategies) is that having cheap access to information that reveals only high quality may be significantly better than knowing everything from the start or knowing nothing to the end. Recent information technology has given consumers cheap access to information. The information accessibility will keep prices low enough that a buyer, in the end, will not have to acquire any information.

Lastly, we show that as the information cost vanishes, the buyer's equilibrium payoff under the H-focused signal structure is weakly higher than in all equilibria and strictly higher than in some equilibrium under an L-focused signal structure, which is analogously defined, or under a combination of the two signal structures. Thus, if the buyer can commit to an information structure before our bargaining game starts, she would want to choose the H-focused signal structure.

---

[5] We note that in practice sellers do not usually subsidize the cost of verification—only buyers bear this cost. For instance, Federal Trade Commission (2022) recommends that buyers obtain an independent inspection before purchasing a used car, as dealers may not disclose all mechanical problems. In the used car market, the burden of proof lies with the buyer, who must take additional steps to verify the quality of the product independently.

The equilibria we focus on are Pareto-undominated perfect Bayesian equilibria (PBEs) in which the buyer uses a pure information acquisition strategy. This assumption is crucial for our results. There is a Pareto-dominated PBE with pure information acquisition strategies and a PBE with mixed information acquisition strategies, in both of which the buyer gains almost zero surpluses.

Related Literature Our study contributes to the literature on bargaining with information acquisition. The baseline setting in the papers mentioned below is the same as ours: The seller makes a take-it-or-leave-it offer that the buyer either accepts or rejects; the buyer does not know the object's value (its quality) but can acquire information about it. Roesler and Szentes (2017) study a buyer-optimal signal structure, which maximizes the buyer's expected payoff. They assume that the seller does not know the quality but observes the buyer's (costless) signal structure before making an offer. In the buyer-optimal signal structure, the buyer is imperfectly informed about the quality. Ravid, Roesler, and Szentes (2021) study a similar model but assume that the seller does not know the buyer's (costly) signal structure. They show that when the cost of information is zero, there exist multiple Pareto-ranked equilibria, while when the cost is positive but tends to zero, equilibria converge to the Pareto-worst equilibrium of the zero-cost benchmark. Our study differs from these papers: In our model, the seller knows the object's quality and the buyer may acquire information *after* the seller makes a price offer. As a result, as the cost vanishes, the buyer gains the full surplus in our model.

Ravid (2020) assumes, like us, that the seller knows the object's quality but assumes, unlike us, that the buyer has to pay a cost to learn an offered price, not only the quality.[6] In his model, the buyer's learning takes place after the seller's offer. He shows that in the presence of the learning cost, the buyer benefits from being imperfectly informed. As the cost vanishes, a difference becomes clear: In his model, the buyer is fully informed and the buyer's payoff vanishes, but in our model, she stays uninformed and the buyer's payoff remains large.

We compare our study with these studies in Table 1. While the existing studies, in the framework of information design or rational inattention, allow the buyer to flexibly choose any kind of signal, we assume that she learns H-focused information. The cost of flexible information may depend on the buyer's prior belief, but the cost of H-focused information does not. For our result, it suffices that the buyer can acquire information after the offer, but allowing pre-offer acquisition as well does not alter the results.

Wolitzky (2023) discusses a version of an ultimatum bargaining game with unobservable

---

[6] In our model, a buyer perfectly observes an offered price without cost.

	RS	RRS	R	Ours
Does the seller know the object's quality?	no	no	yes	yes
Does the seller know the buyer's information structure?	yes	no	no	no
When can the buyer learn, before or after the seller's offer?	before	before	after	before/after
What kind of information does the buyer learn?	flexible	flexible	flexible	focused
Is it costly for the buyer to learn about the object's quality?	no	yes	yes	yes
Is it costly for the buyer to observe the seller's offer?	no	no	yes	no
How much information does the buyer learn as cost vanishes?	imperfect[7]	imperfect	perfect	never

Note: RS stands for Roesler and Szentes (2017), RRS for Ravid, Roesler, and Szentes (2021), and R for Ravid (2020).

Table 1: bargaining with information acquisition

offers and no private information. The responder exogenously gets an imperfect signal of the proposer's offer on the basis of which she makes accept/reject decisions. In particular, since the offers are unobservable, they might be greater than the surplus.

Cheng and Kim (2024) provide a characterization of Blackwell-monotone information costs and study, as an application, our bargaining model. Their analysis is complementary to ours. They dispense with the H-focused signal structure but focus on specific cost functions (a quadratic cost and an absolute-linear cost). By contrast, we restrict the buyer to the H-focused signal structure but do not make parametric assumptions on cost functions.[8] Given that agents may not be able to access all experiments in reality, these results demonstrate the significance of the type of signal they are able to access and the associated costs.

Our model is not a lemons model because the seller's payoff in the event of no trade does not depend on the object's quality. The H-focused information structure used in our paper has some similarities to the one used in Feinberg and Skrzypacz (2005). In their paper, the seller is initially uninformed about the buyer's valuation, which is either high or low. Then, with positive probabilities, he is exogenously informed about whether the buyer's valuation is high or stays uninformed. They show that the buyer's (second-order) uncertainty about the seller's belief changes the outcomes. In our paper, the buyer, who is initially uninformed of her valuation, has the option to acquire information, but *in equilibrium, she does not acquire any* and thus remains uninformed throughout the game.

Our setting has some similarity with the informed principal problem (Myerson, 1983; Maskin and Tirole, 1990), but unlike the informed principal problem, in our model the seller, who knows his type, does not propose a mechanism, and the buyer, who is not initially informed, commits to a particular information structure.[9]

---

[7] The cost of information is zero in their case.

[8] We also relax the H-focused signal assumption and the two-type assumptions.

[9] Rubinstein (1985) in a different bargaining model with incomplete information, about two types, confined himself to pure-strategy equilibria and used a belief where any off-equilibrium offer signals the low type, as in some of our equilibria. Bikhchandani (1992) used the same model but imposed a monotonicity condition

Beyond bargaining problems, information acquisition is shown to have implications on equilibrium outcomes. In contract theory, an agent may acquire costly information about a payoff-relevant state before signing a contract offered by a principal. Kessler (1998) considers the agent who can gather costly information before the principal offers a contract. She argues that the agent may benefit from having less than perfect information. Crémer and Khalil (1992) assume that after the principal offers a contract, the agent can gather costly information. They show that the optimal contract never induces the agent to acquire information, which echoes our result that on path the buyer never acquires information.[10] Information acquisition by agents has been analyzed in other applications, such as auction theory.[11] A common feature of these studies is that a buyer acquires information at exogenously given times; we show that our results continue to hold if we allow the buyer to endogenously choose when to acquire information (Proposition 3).

Our study is also related to early studies on the Coase conjecture (e.g., Fudenberg, Levine, and Tirole, 1985; Gul, Sonnenschein, and Wilson, 1986). They consider dynamic bargaining models where the seller is uninformed about the buyer's valuation, and show that the seller gains no surplus even though he makes all the offers. Although we study static (ultimatum) bargaining, while theirs are dynamic bargaining—the seller gets no surplus even though only the seller can make offers in both models.

Layout The remainder of this paper is organized as follows. Section 2 sets up the model, and Section 3 presents the main results. Section 4 examines key assumptions. In particular, we show the optimality of the H-focused signal structure. Section 5 concludes.

2 Model

A seller (S, he) and a buyer (B, she) bargain over a single indivisible object. S observes the object's quality and then makes a take-it-or-leave-it price offer that B either accepts or rejects. Before making the decision, she can acquire costly signals about the quality.

Timeline There are three stages:

Time 0: Nature draws the object's quality $\mathbf{v}$ from the set $V \equiv \{H, L\}$, with $H > L > 0$,

---

that ruled out these beliefs and derived randomised-strategy equilibria.

[10] Crémer and Khalil (1994) and Crémer, Khalil, and Rochet (1998) assume that the agent may gather information before the principal offers a contract.

[11] There are studies that discuss bidders acquiring costly information about their valuations (e.g, Persico, 2000; Compte and Jehiel, 2007; Shi, 2012). These studies assume that the seller commits to a selling mechanism after which the bidders decide how much information to acquire.

according to a common prior, $\mathbb{P}(\mathbf{v} = H) = \pi \in (0,1)$ and $\mathbb{P}(\mathbf{v} = L) = 1-\pi$.[12] Abusing notation, we denote the prior by the same notation π. S observes the realized quality $v \in V$ (where S is called type v) and then offers price p to B. B observes the offered price p but not the quality v. The object's value to S is zero (regardless of v), while the value to B is v.

Time 1: B acquires a signal $\mathbf{x}$, which is H-focused. That is, a signal realization is either H (which conclusively reveals that $\mathbf{v} = H$) or N (which we interpret as B's observing nothing). She chooses accuracy $q \in [0,1]$, which gives the following conditional distributions:

$$\begin{aligned} \mathbb{P}(\mathbf{x} = H \mid \mathbf{v} = H) &= q, & \mathbb{P}(\mathbf{x} = H \mid \mathbf{v} = L) &= 0, \\ \mathbb{P}(\mathbf{x} = N \mid \mathbf{v} = H) &= 1-q, & \mathbb{P}(\mathbf{x} = N \mid \mathbf{v} = L) &= 1. \end{aligned}$$

Then, B observes a signal realization x.

Time 2: B decides whether or not to buy the object at price p. Payoffs are then realized, and the game ends.

Information Cost The H-focused signal is costly to acquire, and we assume that the cost depends on the accuracy q chosen by B. Let $c : [0,1] \to \mathbb{R}_+$ be the cost function that is continuously differentiable and strictly increasing: $c'(q) > 0$ for each $q \in (0,1)$.

We analyze the case where this cost vanishes. We parameterize the cost by a unit cost parameter $\lambda > 0$, assuming that the cost of the signal with accuracy q is equal to $\lambda c(q)$. Then, we will analyze the case $\lambda \to 0$.

Payoffs Let v be the realized quality. Let p be the price that S chooses and q be the accuracy that B chooses. If B buys then S receives payoff p and B receives payoff $v-p-\lambda c(q)$; otherwise, S receives payoff 0 and B receives payoff $-\lambda c(q)$.

Strategies S chooses a strategy $\sigma : V \to \Delta(\mathbb{R}_+)$, where $\sigma(\cdot \mid v)$ denotes a price distribution when S is of type v.

B chooses accuracy q at time 1 and makes a purchase decision and time 2. At time 1, B observes an offered price p, thereby having a history $h_1 = p$. At time 2, B knows his choice of accuracy q and also the signal realization x, thereby having a history $h_2 = (p, q, x)$. Let $H_1 = \mathbb{R}_+$ be the set of time-1 histories and $H_2 = H_1 \times [0,1] \times \{H, N\}$ be the set of time-2

---

[12] As discussed in Section 4, the main result remains true for any finite set V if the lowest value is strictly positive.

histories. We define B's strategy as a function $\beta : H_1 \cup H_2 \to [0,1]$ such that (i) $\beta(h_1)$ is the choice of accuracy at history $h_1 = p$, and (ii) $\beta(h_2)$ is the probability of buying at history h_2. In the main model, we assume that B does not randomize her information acquisition. We call such a strategy a pure-learning strategy. We revisit this assumption in Section 4.

Equilibrium Our equilibrium concept is perfect Bayesian equilibrium (Fudenberg and Tirole, 1991a,b). An assessment is a tuple (β, σ, μ), where β is B's strategy, σ is S's strategy, and μ is a belief system (about S's type $\mathbf{v}$). Since S's type is binary, we write μ for B's belief of $\mathbf{v} = H$. We write $\mu(h_t)$ for her belief at her history $h_t \in H_1 \cup H_2$. Let $\mu(h_1, q)$ be B's belief after observing history $h_1 = p$ and choosing accuracy q but before observing a signal realization x_2.

Definition 1. An assessment $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$ is a perfect Bayesian equilibrium (PBE) if it satisfies the following conditions:

1. β^* is optimal for B given (σ^*, μ^*), and σ^* is optimal for S given (β^*, μ^*).
2. μ^* is obtained from π on path given (β^*, σ^*), using Bayes's rule.
3. $\mu^*(h_1) = \mu^*(h_1, q)$ for all h_1 and all q.

The third condition, which is called the "no-signaling-what-you-don't-know" condition, requires that B's own deviation not change her belief.

We focus on **Pareto-undominated** PBEs, where the equilibrium payoff vector (consisting of type-H S's, type-L S's, and B's payoffs) is not Pareto-dominated by the payoff vector of any other PBE. In the sequel, unless mentioned otherwise, an **equilibrium** refers to a PBE that is Pareto-undominated.

Benchmark We discuss two benchmarks to put our main result in perspective. The first is the full-information benchmark where the seller and the buyer have common knowledge about the object's quality. This benchmark has a unique equilibrium, and on the equilibrium path, each type v of S offers price v and B accepts both offers. Hence, B gets zero surplus. The second is the null-information benchmark where B has no option to acquire information. This benchmark has a continuum of equilibria, and each type of S offers a price between L and B's expected value $L + \pi(H - L)$, and B accepts it. In these equilibria, B's surplus ranges from 0 to $\pi(H - L)$, which is the maximum possible value, and we call it the full surplus.

3 Main Result

We are interested in the case where B can cheaply acquire information. With positive costs of information $\lambda > 0$, there exists an equilibrium in which each type of S offers price L and B purchases at price L;[13] S has the smallest surplus L, and B has the greatest surplus $\pi(H - L)$. Our main result shows that as λ vanishes, this outcome arises as the limit equilibrium outcome.

Theorem 1. *For each $\epsilon > 0$, there exists some $\bar{\lambda} > 0$ such that for each $\lambda < \bar{\lambda}$,* every *Pareto-undominated PBE $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$ is a pooling equilibrium such that:*

1. *Each type v of S offers a price $p^* \in [L, L + \epsilon)$ with probability 1: $\sigma^*(\{p^*\} \mid v) = 1$.*
2. *On the equilibrium path, B acquires no information and buys at price p^* with probability 1.*

Thus, as the cost parameter λ tends to zero, B extracts full surplus, while each type of S receives the lowest possible payoff L.

Theorem 1 demonstrates the power of having cheap access to information. That is, it gives B (almost) full surplus, although S has the sole power to make a (take-it-or-leave-it) offer and privately knows the object's value. Specifically, B's option of acquiring information deters S from offering a high price, although B never exercises the option. Moreover, giving the learning option to B is better (for B) than giving full information directly to B, because the latter leads to the full-information benchmark, in which B gains zero surplus.[14] An implication of our result is that recent information technology, which has given consumers cheap access to information, can help keep prices low enough that a buyer, in the end, will not have to acquire any information.

We discuss the intuition of our result. With the H-focused signal structure, the signal realization H is a "buy" signal for B. A key feature of this signal structure is that type L of S will never be able to generate the "buy" signal if B were to collect information.[15] Therefore, any price that type L charges must result in zero information acquisition; otherwise, type L would not generate any sale and would deviate to charging price L. Since there is no need

---

[13] B assigns probability 1 to type L for any prices $p \neq L$ (and buys only at price $p \leq L$).

[14] Giving cheap access to information to B is also better than giving free access to information, in the sense that the former option selects the best outcome for B from the equilibria of the free-information benchmark. Indeed, this benchmark has the following continuum of equilibria: for any price $p \in [L, H]$, type H of S offers price p and type L offers price L; if price L is offered, B buys for sure without acquiring information, while if $p \neq L$ is offered, B perfectly learns the type of S and buys the item if and only if S is of type H. Given any off-path price, B assigns probability 1 to type L and buys only when the price is at most L.

[15] We discuss the importance of this feature for the optimality of H-focused signals in Section 4.2.

to acquire information if an on-path price is charged only by type H, this means that B will not acquire information after any on-path offer.

Next, we show that any PBE $\mathcal{E}$ in which S randomizes over multiple equilibrium prices and B never acquires information on the path is Pareto-dominated by another PBE $\mathcal{E}'$. We construct such a PBE $\mathcal{E}'$ by having every type of S offer the lowest price of the PBE $\mathcal{E}$ and B buys at that price for sure without acquiring information. By construction, B and S have weakly higher payoffs in $\mathcal{E}'$ than in $\mathcal{E}$. Moreover, trade is efficient in $\mathcal{E}'$, but not in $\mathcal{E}$, because B must reject some offer with a positive probability in $\mathcal{E}$ so that S is willing to randomize over multiple prices. Thus, B or S gains a strictly higher payoff in $\mathcal{E}'$ than in $\mathcal{E}$.

Finally, we show that in any (Pareto-undominated) PBE, in which S offers one pooling price p^λ, we have $p^\lambda \to L$ as the cost parameter λ vanishes. Since B acquires no information and buys for sure after the equilibrium price p^λ, any deviation to acquiring almost perfect information $q \approx 1$ and buying only when $\mathbf{x} = H$ must be non-profitable. The benefit of this deviation is that B can avoid buying the low-quality good, with value L, at price $p^\lambda > L$, but B forgoes the surplus $H - p^\lambda$ if she gets a "do not buy" signal $x = N$ when the object's value is H. That is, her net benefit, compared to the equilibrium strategy, is $(1-\pi)(p^\lambda - L) - \pi(1-q)(H - p^\lambda) - \lambda c(q)$. Since no deviation is profitable, the net benefit is at most zero. Since the cost of almost perfect information converges to 0 as λ vanishes, the benefit $(1-\pi)(p^\lambda - L)$ must also converge to 0 and thus the price p^* must tend to L.

Remark 1. The requirement of Pareto dominance has bite. There exists a *Pareto-dominated* separating equilibrium in which B acquires no information and gains zero surplus: Each type v of S charges price v, while B buys the object (with probability 1) when the offered price is L or lower, and buys with probability L/H when the offered price is H. This PBE is Pareto-dominated by a pooling equilibrium in which both types of S offer price L. $\square$

4 Discussion

We discuss the assumptions for Theorem 1 and extend the theorem in various directions. First, we show that Theorem 1 holds in settings in which S has any finite number of types (Section 4.1), B has different timing of information acquisition (Section 4.3), or B has a noisy H-focused signal structure (Section 4.7). Second, we show that B prefers the H-focused signal structure over others (Section 4.2). Third, we discuss the importance of pure learning strategies (Section 4.4). Our result, however, does not hold if the value of the object to S is correlated with its value to B (Section 4.5) or if there are not always gains from trade (Section 4.6).

4.1 Number of Seller's Types

Our baseline model assumes that the object's quality v, or the type of S, has only two possible values, $H > L > 0$. This setting simplifies the presentation, but is unnecessary for Theorem 1. Indeed, we show that Theorem 1 remains true for any finite type space, provided that the lowest value L is strictly positive, which ensures that there are always gains from trade. Specifically, the object's quality $\mathbf{v}$ is drawn from the set $V = \{H_1, H_2, \ldots, H_{K-1}, L\}$, with $H_1 > H_2 > \cdots > H_{K-1} > L > 0$, according to a common prior, $\mathbb{P}(\mathbf{v} = v) = \pi_v \in (0,1)$ for all v.

A non-trivial task is how to generalize the notion of the H-focused signal structure. The key idea is to bunch all types except the lowest type into one group, and the generalized H-focused signal conclusively reveals whether S's type is in that group or not. Specifically, a signal realization is either H or N, and B chooses accuracy $q \in [0,1]$, which gives the following conditional distributions: for any $v \in V \setminus \{L\}$,

$$\begin{aligned} \mathbb{P}(\mathbf{x} = H \mid \mathbf{v} = v) &= q, & \mathbb{P}(\mathbf{x} = H \mid \mathbf{v} = L) &= 0, \\ \mathbb{P}(\mathbf{x} = N \mid \mathbf{v} = v) &= 1 - q, & \mathbb{P}(\mathbf{x} = N \mid \mathbf{v} = L) &= 1. \end{aligned}$$

As in the two-type model, $\mathbf{x} = H$ conclusively reveals that $\mathbf{v} \neq L$ and $\mathbf{x} = N$ is interpreted as B's observing nothing. However, $\mathbf{x} = H$ is uninformative regarding whether type $\mathbf{v}$ is any of $H_1, H_2, \ldots, H_{K-1}$. That is, $\mathbb{P}(\mathbf{v} = H_k \mid \mathbf{x} = H)/\mathbb{P}(\mathbf{v} = H_{k'} \mid \mathbf{x} = H) = \pi_{H_k}/\pi_{H_{k'}}$ for all k, k'. Note that this setting includes the baseline two-type model as a special case when $K = 2$ and $H_1 = H$.

In this generalized setting, Theorem 1 holds true without any modification. Our proof depends on the property that B cannot distinguish the type of S, and this property is specific to the generalized H-focused signal structure. Hence, B would choose such a structure given that B can choose (and commit to) her signal structure. Nevertheless, we note that it is an open question whether a more complex signal structure that is close to the generalized H-focused signal structure still gives our result.[16]

4.2 Optimality of H-Focused Signal

So far we have assumed an H-focused signal structure. We now consider why B may want to commit to it in advance. With the H-focused signal structure, as the cost parameter goes to zero, B obtains the full surplus in the limit equilibrium. With other information structures, we show that there are Pareto-undominated PBEs in which B gets almost zero

---

[16] We appreciate an anonymous referee for raising this question.

surplus. Therefore, if B can choose among information structures, she will choose the H-focused signal structure as her payoff is weakly higher than in all equilibria and strictly higher than in some equilibria under other information structures.

H-Focused Signal versus L-Focused Signal Similar to the H-focused signal structure, we define the **L-focused signal structure**, with two possible realizations L and N, by the following conditional distributions:

$$\begin{aligned} \mathbb{P}(\mathbf{x} = L \mid \mathbf{v} = H) &= 0, & \mathbb{P}(\mathbf{x} = L \mid \mathbf{v} = L) &= q, \\ \mathbb{P}(\mathbf{x} = N \mid \mathbf{v} = H) &= 1, & \mathbb{P}(\mathbf{x} = N \mid \mathbf{v} = L) &= 1 - q. \end{aligned}$$

The L-focused signal may conclusively reveal low quality in a similar way that the H-focused signal may conclusively reveal high quality. Assume that the same cost for the L-focused signal as the cost of the H-focused signal. That is, B incurs the cost $\lambda c(q)$ when choosing accuracy q in the L-focused signal structure.

Proposition 1. *Under the L-focused signal structure, for each $\epsilon > 0$, there exists some $\bar{\lambda} > 0$ such that for each $\lambda < \bar{\lambda}$, there exists* some *equilibrium $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$ such that the payoff of B is in the interval $(0, \epsilon)$ and the payoff of each type v of S is in the interval $(v - \epsilon, v + \epsilon)$.*

Proposition 1 implies that if B can commit to either H- or L-focused signal structure before time 0, she will choose the H-focused one because her payoff under the H-focused one is weakly higher than that in all Pareto-undominated PBEs and strictly higher than in some of them under the L-focused one.

Here we discuss the intuition for Proposition 1. As we have mentioned, the key feature of the H-focused signal structure is that type L of S cannot generate a "buy" signal, which results in any equilibrium being a pooling equilibrium. However, with the L-focused signal structure, type L can generate a "buy" signal, because now B may buy the object if she observes a realization $x = N$. This intuition suggests the existence of an "almost separating" equilibrium when the cost of information is sufficiently low. Indeed, we prove Proposition 1 by constructing a semi-separating equilibrium such that (i) type H offers price $p^* \approx H$ (with probability 1), and type L offers price L with a probability close to 1 and price p^* with the complementary probability, and (ii) on the equilibrium path, B acquires information (by choosing accuracy q_L^* bounded away from 0) and conditions her purchase decision on the signal realization. This equilibrium is "almost separating" since different types offer different prices with high probabilities. In the proof in Appendix A.2, we show that B's payoff is close to zero in this equilibrium.

H-Focused Signal versus Combined Signal Now we allow B to combine the two signal structures. We define a **combined signal structure**, which includes H- and L-focused signal structures as special cases. A combined signal structure has three realizations H, L, and N, and the following conditional distributions:

$$\begin{aligned} &\mathbb{P}(\mathbf{x} = H \mid \mathbf{v} = H) = q_H, && \mathbb{P}(\mathbf{x} = L \mid \mathbf{v} = L) \ = q_L, \\ &\mathbb{P}(\mathbf{x} = N \mid \mathbf{v} = H) = 1 - q_H, && \mathbb{P}(\mathbf{x} = N \mid \mathbf{v} = L) = 1 - q_L. \end{aligned}$$

As before, the realizations H and L conclusively reveals high quality and low quality respectively. Indeed, an H-focused (resp. L-focused) signal structure is the combined signal structure with accuracy $q_L = 0$ (resp. $q_H = 0$). Assume that the cost of choosing q_H of an H signal is $\lambda c(q_H)$ and the cost of choosing q_L of an L signal is $\lambda c(q_L)$.

Proposition 2. *Assume that the cost function c is quadratic: $c(q) = \lambda q^2/2$. Under the combined signal structure, for each $\epsilon > 0$, there exists some $\bar{\lambda} > 0$ such that for each $\lambda < \bar{\lambda}$, there exists* some *equilibrium $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$ such that the payoff of B is in the interval $(0, \epsilon)$ and the payoff of each type v of S is in the interval $(v - \epsilon, v + \epsilon)$.*

Proposition 2 implies that *when the cost is quadratic, if B can commit to either H-focused or combined signal structure before time* 0*, she will choose the H-focused one.* We show that the equilibrium constructed for Proposition 1 remains an equilibrium under a combined signal structure. B has no incentive to deviate to acquire the H-focused signal. The reason is that in these equilibria, after any on-path offer, B's expected value of the good exceeds the price. If B were to deviate to the H-focused signal, she would need to choose a relatively high accuracy q_H for the signal to be useful (i.e., to buy if $x = H$ and not to buy if $x = N$), which is not as profitable as choosing the L-focused signal.

4.3 Acquiring Pre-Offer Information and Post-Offer Information

We have so far assumed that B acquires information only after observing S's price offer. In this subsection, we ask whether Theorem 1 remains true (i.e., whether B extracts full surplus in the limit $\lambda \to 0$) when B is allowed to acquire pre-offer information as well as post-offer information. We will show that Theorem 1 still holds in this variant.

We modify the baseline model, which is introduced in Section 2, so that at time 0 (before observing an offered price p), B chooses accuracy $q_0 \in [0, 1]$ to acquire an H-focused signal $\mathbf{x}_0$, which is drawn according to the following conditional distributions:

$$\mathbb{P}(\mathbf{x}_0 = H \mid \mathbf{v} = H) = q_0, \qquad \mathbb{P}(\mathbf{x}_0 = H \mid \mathbf{v} = L) = 0,$$

$$\mathbb{P}(\mathbf{x}_0 = N \mid \mathbf{v} = H) = 1 - q_0, \qquad \mathbb{P}(\mathbf{x}_0 = N \mid \mathbf{v} = L) = 1.$$

At the end of time 0, B observes both a signal realization x_0 and an offered price p. We do not modify the stages at time 1 and time 2 of our model in Section 2.

We modify histories according to the change. We introduce time-0 history $h_0 = \varnothing$. We also include pre-offer accuracy q_0 and a signal realization x_0 in histories h_1 and h_2. That is, we have time-1 history $h_1 = (p, q_0, x_0)$ and time-2 history $h_1 = (p, q_0, x_0, q, x)$. Accordingly, B's strategy β is a function such that (i) $\beta(\varnothing)$ is the choice of accuracy at history $h_0 = \varnothing$, (ii) $\beta(h_1)$ is the choice of accuracy at history h_1, and (iii) $\beta(h_2)$ is the probability of buying at history h_2.

Payoff structures are modified in a straightforward way. Both the pre-offer and post-offer signals are costly, and we assume that B incurs a total cost $\lambda c(q_0) + \lambda c(q)$ when she chooses pre-offer accuracy q_0 and post-offer accuracy q.

B's option of acquiring pre-offer information does not alter the main result (Theorem 1) of the baseline model, as shown in the following proposition.

Proposition 3. *For each $\epsilon > 0$, there exists some $\bar{\lambda} > 0$ such that for each $\lambda < \bar{\lambda}$,* every *Pareto-undominated PBE $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$ is a pooling equilibrium such that:*

1. *Each type v of S offers a price $p^* \in [L, L+\epsilon)$ with probability 1 (i.e., $\sigma^*(\{p^*\} \mid v) = 1$).*
2. *On the equilibrium path, B acquires neither pre- nor post-offer information and buys at price p^* with probability 1.*

Thus, as the cost parameter λ tends to zero, B extracts full surplus, while each type of S receives the lowest possible payoff L.

The equilibrium outcome identified in Proposition 3 is essentially the same as that in Theorem 1. These results, therefore, demonstrate that B chooses to stay uninformed throughout the bargaining process and also extracts (almost) full surplus, even if she has an option of acquiring information anytime—i.e., before and after S offers a price.

4.4 Randomizing Information Acquisition

We have assumed that B does not randomize her choice of accuracy (but may randomize her purchase decision). In this restriction, cheap access to information yields (almost) full surplus to B in every equilibrium (Theorem 1). This raises the question as to whether B's advantage from having cheap access to information might survive in equilibria in which B randomizes information acquisition. We find that as the cost parameter λ tends to zero, B's equilibrium payoff set in the limit is $[0, \pi(H-L)]$, as is shown in Proposition 4. That is, any payoff between zero surplus and full surplus can be obtained in some limit equilibrium. For simplicity, we refer to a PBE such that B mixes between acquiring information and not acquiring information on the equilibrium path as a *mixed-learning equilibrium*.

Proposition 4. *Suppose that the cost function c of information satisfies $c'(0) = 0$. For each $\epsilon > 0$, take any $u_B \in (\epsilon, \pi(H-L) - \epsilon)$. There exists some $\bar{\lambda} > 0$ such that for each $\lambda < \bar{\lambda}$, there exists a mixed-learning equilibrium in which B's payoff equals u_B; moreover, it is Pareto-undominated among all mixed-learning equilibria.*

The proof is in the Online Appendix C.

A mixed-learning equilibrium of this proposition is a semi-separating equilibrium that is "almost separating." Fix a sufficiently small $\lambda > 0$, and take any $p \in (L, H)$. Then, here is a mixed-learning equilibrium in which B's payoff is close to $\pi(p - L)$:

- For S, type H charges the price p with probability 1, while type L charges the price p with a high probability y and the price L with the complementary probability $1 - y$.
- For B, if the equilibrium price p is offered, then with probability L/p, B buys the object without acquiring any information and with probability $1 - L/p$, B acquires almost full information (at accuracy q close to 1) and buys if and only if she learns the signal realization $\mathbf{x} = H$. For all the other price offers, B buys if and only if the price is at most L.

This PBE is almost separating on the side of S. Indeed, the probability of type L offering price p (i.e., mimicking type H) converges to 0 as the cost parameter λ vanishes. In addition, B may or may not acquire information. That is, she randomizes between accuracies 0 and $q(p, \lambda)$. This accuracy $q(p, \lambda)$, if B opts for acquiring information, also converges to 1.

This contrast between the pure and mixed-learning equilibria may be reminiscent of Bagwell (1995). He argues that in a Stackelberg-like setting where the follower's observation of the leader's action is slightly noisy, the first-mover advantage is eliminated in pure-strategy subgame perfect equilibrium, but it is preserved in some mixed-strategy equilibrium. However, our mechanisms are different. In Bagwell (1995), imperfect observation of the leader's action increases the leader's incentive to deviate to lower production. In our paper, the low-type seller has an incentive to mimic the high-type seller, forcing the high-type seller to offer low prices if the buyer uses pure-learning strategies.

4.5 Correlated Value with H-Focused Signals

Theorem 1 may not hold in the "lemons" variant, where the value of the object to S is correlated with its value to B. To illustrate this point, assume that the seller's reservation value for the item is αv, for each $v \in \{L, H\}$, with $\alpha \in [0, 1)$. Consider the separating equilibrium in which type v offers price v; if L is offered, B buys with probability one, and if H is offered, B buys with probability $\frac{(1-\alpha)L}{H-\alpha L}$. With independent values (i.e., $\alpha = 0$), this equilibrium is Pareto-dominated by the pooling equilibrium in which both types of S charge price L and B buys for sure at this price without acquiring information. With such correlated values, the separating equilibrium is not Pareto-dominated by any pooling equilibrium if λ is sufficiently low. To see this, suppose by contradiction that a pooling equilibrium with some equilibrium price p Pareto dominates the separating equilibrium. Then, the payoff of type H of S in the pooling equilibrium, which is $p - \alpha H$, must be weakly higher than his payoff in the separating equilibrium, which is $(1-\alpha)H \cdot \frac{(1-\alpha)L}{H-\alpha L}$. That is, $p \geq L + \frac{\alpha(H-L)}{H-\alpha L}(H - L)$ must hold. If λ is sufficiently low, then B would strictly prefer to acquire information, which means that type L of S will not make any sale by charging this price, which contradicts the assumption that p is a pooling price.

4.6 Gains from Trade

Our main model has assumed that there always exist gains from trade. That is, the lowest quality is assumed to be $L > 0$. This assumption is crucial: If it is violated (i.e., if $L \leq 0$) then Theorem 1 does not hold. Indeed, the following equilibrium exists: for S, type L mixes between offering prices 0 and $p_H \in (L, H)$, where p_H is sufficiently close to H, and type H offers price p_H with probability 1; for B, if any (positive) price $p \neq p_H$ is offered, then B assigns probability 1 to type L and does not buy, and if price p_H is offered then B acquires information and buys if and only if the signal realization is $\mathbf{x} = H$. In this equilibrium, B's surplus is close to zero.

4.7 Noisy H-Focused Signal Structure

Our main model has assumed that the H-focused signal $\mathbf{x}$ is conclusive in the sense that the realization $\mathbf{x} = H$ reveals type $\mathbf{v} = H$. Under this assumption, type L of S cannot "fool" B because he cannot generate the "buy" signal $\mathbf{x} = H$. This assumption significantly simplifies the analysis, but we show that it is unnecessary. In particular, we show that our main result holds if the signal structure is almost H-focused in the sense that the probability ratio $\mathbb{P}(\mathbf{x} = H \mid \mathbf{v} = L)/\mathbb{P}(\mathbf{x} = H \mid \mathbf{v} = H)$ is less than L/H. This result extends Theorem 1 (in which the said ratio equals zero), and the proof is similar to that of the original Theorem 1 but is more involved. See Appendix B.

5 Conclusion

We have considered an ultimatum bargaining problem with a seller (the proposer) who is informed about the quality of his good and a buyer (the responder) who is initially uninformed about the quality. We allow the buyer to acquire information about the quality after the seller's proposal at a cost. Compared to the benchmark cases, we obtain striking conclusions. In the full information case, where the object's quality is common knowledge, the seller uses his proposal power to obtain full surplus in the unique equilibrium. In the null information case, where the buyer has no option to acquire information about the quality, there is a continuum of Pareto-undominated equilibria, including one where the buyer's surplus is zero. However, in our case, where the buyer has cheap access to information about the quality, in all Pareto-undominated equilibria, the buyer obtains the full surplus although she never acquires any information.

A Proofs

A.1 Theorem 1

We prove Theorem 1 for the general model in which S has K possible types for any $K \geq 2$. The extended model is in Section 4.1. Readers who are only interested in the two-type model with types $H > L > 0$ may assume, in the proof below, that $K = 2$ and $H_1 = H$ and replace π_{H_1} with π.

Step 1 We show that every Pareto-undominated PBE is such that all types of S offer one price (with probability 1). That is, we prove that any PBE $\mathcal{E}$ in which the set of equilibrium prices is not a singleton is Pareto-dominated by some PBE $\mathcal{E}'$, which we will construct from the PBE $\mathcal{E}$.

Lemma 1. *Given any PBE* $\hat{\mathcal{E}} = (\hat{\beta}, \hat{\sigma}, \hat{\mu})$*, B chooses accuracy* $\hat{\beta}(p) = 0$ *after any on-path price offer* $p \in \bigcup_{v \in V} \operatorname{supp}(\hat{\sigma}(\cdot \mid v))$*.*

In words, this lemma claims that for any PBE that has multiple equilibrium prices, B does not acquire any information on the equilibrium path.

Proof. Let a PBE $\hat{\mathcal{E}} = (\hat{\beta}, \hat{\sigma}, \hat{\mu})$ be given. To prove the lemma, it suffices to show that B chooses accuracy $\hat{\beta}(p) = 0$ after any price offer $p \in \operatorname{supp}(\hat{\sigma}(\cdot \mid L))$ that type L of S may offer. This is because, by the definition of the H-focused signal structure, the signal $\mathbf{x}$ is uninformative about types H_k for all k.

If B is offered price $p = L, H_1$, then she does not acquire information for any $\lambda > 0$. Then, we consider the case where B is offered a price $p \in (L, H_1) \cap \operatorname{supp}(\hat{\sigma}(\cdot \mid L))$, which type L may offer. Suppose, by contradiction, that B chooses accuracy $q > 0$ after that p. Then, B's purchase decision depends on the realization of her signal $\mathbf{x}$. That is, B must buy if $\mathbf{x} = H$ but never buys if $\mathbf{x} = N$. This is because if her purchase decision were independent of the realization of $\mathbf{x}$, she would profitably deviate to choosing zero accuracy, which saves the information cost without changing her purchase decision. Since the conditional probability of $\mathbf{x} = N$ given $\mathbf{v} = L$ is 1, type L's profit must be zero. Hence, type L has a profitable deviation of offering price L or a price slightly less than that.[17] This is a contradiction. ■

Lemma 2. *Given any PBE $\mathcal{E} = (\beta, \sigma, \mu)$ such that the set $\bigcup_{v \in V} \operatorname{supp}(\sigma(\cdot \mid v))$ is not a singleton, let $\underline{p} = \inf\{\bigcup_{v \in V} \operatorname{supp}(\sigma(\cdot \mid v))\}$ be the infimum equilibrium price. Let $\mathcal{E}'$ be the assessment defined as follows:*

1. *Each type v of S offers price $\underline{p}$.*
2. *B acquires no information on the equilibrium path.*
   - (a) *At any (on- or off-path) history h_1, B chooses accuracy $q = 0$.*
   - (b) *At history h_2, if $p = \underline{p}$ or $p \leq L$ then B buys with probability 1; if $p \in (L, H_1] \setminus \{\underline{p}\}$ then she assigns probability 1 to type L, chooses accuracy $q = 0$, and buys with probability 0.*

This assessment $\mathcal{E}'$ is a pooling PBE.

Proof. In the assessment $\mathcal{E}'$, any type of S has no profitable deviation. This is because given B's belief and strategy, any type v of S makes a profit of $\underline{p}$ if he offers price $\underline{p}$ but makes a less profit if he charges any other price.

In order to prove the lemma, it suffices to show that after the price offer $\underline{p}$, B optimally chooses accuracy $q = 0$ and buys with probability 1. As shown in the proof of Lemma 1, whenever B chooses any $q > 0$, her purchase decision must be such that she buys if $\mathbf{x} = H$ but never buys if $\mathbf{x} = N$. Thus, it suffices to show that B's payoff when choosing $q = 0$ and buying for sure is at least her payoff when choosing $q > 0$ and buying if and only if $\mathbf{x} = H$. The former payoff is $\mathbb{E}[\mathbf{v}] - \underline{p}$, where $\mathbb{E}[\mathbf{v}]$ is the ex-ante expected value of the object. The latter payoff is $\sum_{k=1}^{K-1} \pi_{H_k} q(H_k - \underline{p}) - \lambda c(q)$.[18] By Lemma 1, B chooses accuracy $q = 0$ after

---

[17] In any equilibrium, if price L is offered, B must buy with probability one because otherwise, type L can give B a strict incentive to buy by cutting the price infinitesimally and has a profitable deviation.

[18] This payoff is $\pi q(H - \underline{p}) - \lambda c(q)$ for the two-type model.

any price offered in the PBE $\mathcal{E}$. Therefore, if she buys with a non-zero probability at the offered price, then her payoff is also equal to the payoff from buying for sure. Thus, her payoff in $\mathcal{E}$ is equal to the payoff from buying for sure (without acquiring information) at every equilibrium price. For B to be willing to acquire no information after any equilibrium price offer in $\mathcal{E}$, B's payoff when choosing $q = 0$ and buying for sure must be weakly greater than her payoff when choosing $q \in (0, 1]$ and buying if and only if $\mathbf{x} = H$. This implies that for any $p \in \bigcup_{v \in V} \operatorname{supp}(\sigma(\cdot \mid v))$,

$$\pi_L(p)(L - p) + \sum_{k=1}^{K-1} \pi_{H_k}(p)(H_k - p) \geq \sum_{k=1}^{K-1} \pi_{H_k}(p) q (H_k - p) - \lambda c(q),$$

where $\pi_v(p)$ is B's posterior probability that S is of type $v \in V$ after the price offer p in the PBE $\mathcal{E}$. Since $p \geq \underline{p}$,

$$\pi_L(p)(L - \underline{p}) + \sum_{k=1}^{K-1} \pi_{H_k}(p)(H_k - \underline{p}) \geq \sum_{k=1}^{K-1} \pi_{H_k}(p) q (H_k - \underline{p}) - \lambda c(q),$$

Multiplying both sides by the probability that price p is offered in the PBE $\mathcal{E}$ and summing up over all prices p, we obtain that

$$\mathbb{E}[\mathbf{v}] - \underline{p} \geq \sum_{k=1}^{K-1} \pi_{H_k} q (H_k - \underline{p}) - \lambda c(q)$$

for all $q \in (0, 1]$. It is, therefore, concluded that the assessment $\mathcal{E}'$ is a (pooling) PBE. ■

Lemma 3. *Let $\mathcal{E}$ and $\mathcal{E}'$ be the PBEs in Lemma 2. Then, $\mathcal{E}$ is Pareto-dominated by $\mathcal{E}'$.*

Proof. First, we show that B's payoff in $\mathcal{E}'$ is weakly higher than in $\mathcal{E}$. The reason is that in both $\mathcal{E}$ and $\mathcal{E}'$, B's payoff must be equal to the payoff from buying the object with probability 1 (without acquiring information) whenever an equilibrium price is offered. Since the equilibrium price $\underline{p}$ in $\mathcal{E}'$ is, by assumption, lower than or equal to any equilibrium price in $\mathcal{E}$, B's payoff in $\mathcal{E}'$ is weakly higher than in $\mathcal{E}$.

Second, we show that S's profit in $\mathcal{E}'$ is weakly higher than that in $\mathcal{E}$. Note that since B does not acquire information after any equilibrium price in equilibrium $\mathcal{E}$ (by Lemma 1), every type of S must obtain the same equilibrium payoff, or else a type with a lower payoff would find it profitable to deviate to an equilibrium price offered by another type with a higher payoff. Thus, it suffices to show that S's profit in $\mathcal{E}$ is at most $\underline{p}$ regardless of his type, because his profit in $\mathcal{E}'$ is $\underline{p}$. To show this, suppose, by negation, that his profit is $\Pi > \underline{p}$. Since $\underline{p}$ is, by definition, the infimum of the equilibrium prices in $\mathcal{E}$, $\mathcal{E}$ must have some price $p' < (\Pi + \underline{p})/2$. Then, his profit from offering price p' is at most p', which is strictly lower than Π, but it contradicts the fact that p' is an equilibrium price and Π is the equilibrium profit for S.

Lastly, we show that either B or one of the types of S has a strictly higher payoff in $\mathcal{E}'$ than in $\mathcal{E}$. In order to show this, it suffices to show that trade is efficient in $\mathcal{E}'$ but it is not in $\mathcal{E}$. The efficiency in $\mathcal{E}'$ is immediate because each type of S offers the same price $\underline{p}$ and then B buys for sure without acquiring information. It remains to show that trade is not efficient in $\mathcal{E}$. Since B does not acquire information and buys for sure after any equilibrium price offers (Lemma 1), each type of S must be indifferent between any two equilibrium prices. Each type v of S is indifferent between PBE prices $p_h, p_l \in \bigcup_{v \in V} \operatorname{supp}(\sigma(\cdot \mid v))$ with $p_h > p_l$. Both prices must bring the same expected profit, $p_h y_h = p_l y_l$, where y_h, y_l denote the probabilities that B buys when prices p_h, p_l are offered, respectively. It implies that $y_h = p_l y_l / p_h < y_l \leq 1$, which means that trade is not efficient in $\mathcal{E}$. ■

By Lemmas 2 and 3, we conclude that any PBE $\mathcal{E}$ such that there are at least two equilibrium prices is Pareto-dominated by another PBE $\mathcal{E}'$. Hence, every Pareto-undominated PBE is a PBE such that all types of S offer one price (with probability 1).

Step 2 To complete the proof, we show that for any PBE in which S offers one price (regardless of types), we show that the price approaches L as the cost parameter λ vanishes.

Consider our bargaining model with a cost parameter $\lambda > 0$. By Lemmas 2 and 3, any Pareto-undominated equilibrium $\mathcal{E}^\lambda = (\beta^\lambda, \sigma^\lambda, \mu^\lambda)$ has a unique equilibrium price, denoted p^λ. That is, $\bigcup_{v \in V} \operatorname{supp}(\sigma^\lambda(\cdot \mid v)) = \{p^\lambda\}$. Our goal is to show that $\lim_{\lambda \to 0} p^\lambda = L$. For this purpose, we find an upper bound $\bar{p}^\lambda$ of the equilibrium price p^λ and then show that for each $\epsilon > 0$, there exists $\bar{\lambda} > 0$ such that for each $\lambda < \bar{\lambda}$, $\bar{p}^\lambda < L + \epsilon$.

In a PBE such that all types of S offer the same price $p^\lambda > L$, B acquires no information and buys with probability 1 at any on-path histories (Lemma 1). Hence, her equilibrium payoff is $\mathbb{E}[\mathbf{v}] - p^\lambda$, where $\mathbb{E}[\mathbf{v}]$ is the ex-ante expected value of the object. Consider the case where B deviates to a strategy such that B chooses accuracy $\beta(p^\lambda) = q > 0$ after the price offer p^λ and buys (with probability 1) if she observes a signal realization $\mathbf{x} = H$ but never buys if $\mathbf{x} = N$. This deviation yields B's payoff $(1 - \pi_L) q (\mathbb{E}[\mathbf{v} \mid \mathbf{x} = H] - p^\lambda) - \lambda c(q)$, where $\mathbb{E}[\mathbf{v} \mid \mathbf{x} = H]$ is the conditional expected value of the object given the signal $\mathbf{x} = H$. Since the equilibrium payoff is at least the payoff from the deviation, it must be that for any $q > 0$,

$$\mathbb{E}[\mathbf{v}] - p^\lambda \geq (1 - \pi_L) q (\mathbb{E}[\mathbf{v} \mid \mathbf{x} = H] - p^\lambda) - \lambda c(q).$$

Since $(1-\pi_L)\mathbb{E}[\mathbf{v} \mid \mathbf{x}=H] = \sum_{k=1}^{K-1} \pi_{H_k} H_k$ by Bayes' rule, it follows that

$$p^\lambda \leq \bar{p}^\lambda \equiv L + \frac{1-q}{1-(1-\pi_L)q} \sum_{k=1}^{K-1} \pi_{H_k}(H_k - L) + \frac{\lambda c(q)}{1-(1-\pi_L)q}. \tag{1}$$

Fix any $\epsilon > 0$. There exists $\underline{q} < 1$ such that for any $q \in (\underline{q}, 1)$,

$$\frac{1-q}{1-(1-\pi_L)q} \sum_{k=1}^{K-1} \pi_{H_k}(H_k - L) < \frac{\epsilon}{2}.$$

Fo each $q \in (\underline{q}, 1)$, there exists $\bar{\lambda} > 0$ such that for any $\lambda < \bar{\lambda}$,

$$\frac{\lambda c(q)}{1-(1-\pi_L)q} < \frac{\epsilon}{2}$$

Evaluating (1) with these inequalities, we have $\bar{p}^\lambda < L + \epsilon$, which completes the proof.

A.2 Proposition 1

The proof is by construction. Consider the following assessment:

- Type H of S offers price $p^* \in (L, H)$, where p^* is close to H. Type L offers prices p^* and L with probabilities $y^* \in (0,1)$ and $1-y^*$ respectively, where y^* is close to 0.
- If B observes price p^*, then she chooses accuracy $\beta(p^*) = q^* > 0$ and buys if $\mathbf{x} = N$ and does not buy if $\mathbf{x} = L$. If B observes price $p \neq p^*$, then she assigns probability 1 to S being of type L and thus chooses accuracy $\beta(p) = 0$ and buys if and only if the offered price is at most L.

First, we show that there is a tuple (p^*, q^*, y^*), with price p^* close to H, such that the above assessment is a PBE. Then, we prove that B's equilibrium payoff is close to zero.

There are two conditions for the above assessment to be a PBE. First, type L of S must be indifferent between offering prices L and p^*. Type L makes a sale only if either (i) he offers price L or (ii) he offers price p^* and B observes a "buy" signal realization $\mathbf{x} = N$. Price L gives profit L, while price p^* gives profit $(1-q^*)p^*$, where $1-q^*$ is the probability that B observes $\mathbf{x} = N$. Hence,

$$L = p^*(1-q^*), \tag{2}$$

which implies that type H strictly prefers to offer price p^*. Second, B's choice of accuracy must be optimal given her belief. It suffices to consider the case where she observes price p^*

(because otherwise, she assigns probability 1 to type L). The posterior probability that B assigns to type H when observing price p^* is, by Bayes' rule,

$$\pi_1 = \frac{\pi_0}{\pi_0 + (1-\pi_0)y^*}.$$

B's equilibrium choice of accuracy q^* solves

$$\max_q \quad \pi_1(H - p^*) + (1-\pi_1)((1-q)(L-p^*) + q \times 0) - \lambda c(q),$$

where (i) if S is of type H then with probability 1, B observes $\mathbf{x} = N$ and buys to get payoff $H - p^*$, and (ii) if S is of type L then with probability $1-q$, B observes $\mathbf{x} = N$ and buys to get payoff $L - p^*$, while with probability q, B observes $\mathbf{x} = L$ and does not buy. The first-order condition for the above problem is that $(1-\pi_1)(p^* - L) = \lambda c'(q^*)$. Substituting the above π_1 into this first-order condition, we obtain that

$$\frac{(1-\pi_0)y^*}{\pi_0 + (1-\pi_0)y^*}(p^* - L) = \lambda c'(q^*).$$

Since $p^* - L = p^* q^*$ by equation (2), it follows that

$$\frac{(1-\pi_0)y^*}{\pi_0 + (1-\pi_0)y^*} p^* q^* = \lambda c'(q^*). \tag{3}$$

Next, we construct some (p^*, q^*, y^*) that satisfies conditions (2) and (3) and gives B a payoff close to zero. By condition (2), p^* is uniquely determined by q^*. We fix any small enough $\epsilon > 0$, and let $q^* = 1 - \frac{L}{H} - \epsilon$, which, as we shall see, brings desirable properties. To show the existence of the remaining parameter y^*, we define the following function:

$$\phi(y) = \frac{(1-\pi_0)y}{\pi_0 + (1-\pi_0)y} p^* q^* - \lambda c'(q^*).$$

By definition, any solution to equation $\phi(y) = 0$ in the interval $[0,1]$ is the desired y^*. Since the function ϕ is continuous, to show the existence of a solution, it suffices to show that $\phi(0)\phi(1) \leq 0$. It is immediate that $\phi(0) = -\lambda c'(q^*) < 0$. Take any λ small enough that for each $q \in (1 - \frac{L}{H} - 2\epsilon, 1 - \frac{L}{H})$,

$$(1-\pi_0)L > \frac{\lambda c'(q)}{q}.$$

Then,

$$\phi(1) = q^*\left((1-\pi_0)p^* - \frac{\lambda c'(q^*)}{q^*}\right) \geq q^*\left((1-\pi_0)L - \frac{\lambda c'(q^*)}{q^*}\right) > 0.$$

Hence, there is some $y^* \in (0,1)$ that satisfies equation (3). By this argument, we know that the above (p^*, q^*, y^*) with accuracy $q^* = 1 - \frac{L}{H} - \epsilon$ makes the assessment, described at the beginning of the proof, a (semi-separating) equilibrium. It then remains to show that B's surplus is close to 0 in this equilibrium. Since $p^* = H/(1 + \frac{H}{L}\epsilon) < H$, which we obtain by substituting $q^* = 1 - \frac{L}{H} - \epsilon$ into equation (2), we find that price p^* is close to H (from below). For any small λ and ϵ, we must have probability y^* close to 0.

This semi-separating equilibrium is "almost separating" in the sense that type H of S offers price p^* (close to H) with probability 1 and type L of S offers price L with probability close to 1. Therefore, B's equilibrium payoff is close to zero.

A.3 Proposition 2

Step 1 We show that after any price offer, it is not optimal for B to choose positive accuracy for both the H- and L-focused signals. The intuition is that if a signal does not change B's purchase decision irrespective of the realization, then it is not worthwhile to acquire it. Suppose by contradiction that there is some PBE in which after some equilibrium price p, B chooses $q_L > 0$ and $q_H > 0$. This is optimal for B only if $p \in (L, H)$. There are three signal realizations L, N, H. Since the signal realizations L and H conclusively reveal qualities L and H, respectively, and since price p is in (L, H), it must be that B chooses not to buy after realization L (since the benefit from buying $L - p$ is strictly negative) and to buy after realization H (since the benefit from buying $H - p$ is strictly positive). There are two cases regarding the purchase decision after realization N: (a) B buys with positive probability, and (b) B buys with zero probability. In Case (a), B's payoff is

$$\pi_p(H-p) + (1-\pi_p)(1-q_L)(L-p) - c(q_L) - c(q_H) \tag{4}$$

If B chooses not to acquire the H-focused signal (i.e., chooses accuracy $\tilde{q}_H = 0$ for the H-focused signal) and keeps accuracy q_L for the L-focused signal, and buys if N occurs and does not buy if L occurs, then B's payoff is

$$\pi_p(H-p) + (1-\pi_p)(1-q_L)(L-p) - c(q_L),$$

which is strictly higher than the payoff in (4). Moreover, buying after signal realization N is optimal because the expected value of the object after realization N, $\frac{\pi_p H+(1-\pi_p)(1-q_L)L}{\pi_p+(1-\pi_p)(1-q_L)}$, exceeds the price p, by the fact that the expression in (4) is positive. Therefore, not acquiring the H-focused signal is a profitable deviation. Similarly, in Case (b), deviating to not acquiring the L-focused signal is profitable.

Step 2 We show that the PBE constructed for the L-focused signals (Proposition 1) remains a PBE even when B acquires a combined signal. It suffices to show that in the assessment, B has no incentive to deviate to choose accuracy $q_H > 0$ (and $q_L = 0$ for the L-focused signal). The reason is that by Step 1, after any price offer, it is never optimal for B to choose $q_H > 0$ and $q_L > 0$; if deviating to $q_H > 0$ and $q_L = 0$ is not profitable for any $q_H > 0$, then acquiring only the L-focused signal is optimal. Since the PBE constructed for the L-focused signals (Proposition 1) is a PBE, the accuracy for the L-focused signal is optimal for B. In the rest of the proof, we compare B's payoff in this PBE constructed for the L-focused signals (Proposition 1) to B's payoff if B deviates to any $q_H > 0$ and.

Given an offered price p, let π_p denote B's posterior belief and $\mathbb{E}_{\pi_p}[\mathbf{v}]$ be her expected value. If B chooses accuracy q_L for the L-focused signal (which can be optimal only if she buys when she does not observe signal L), B's payoff, denoted by $U_L(q_L)$, is given by

$$U_L(q_L) = \pi_p(H-p) + (1-\pi_p)(1-q_L)(L-p) - c(q_L). \tag{5}$$

If B chooses accuracy q_H for the H-focused signal, (which can be optimal only if she does not buy when she does not observe signal H), B's payoff, denoted by $U_H(q_H)$, is given by

$$U_H(q_H) = \pi_p q_H(H-p) - c(q_H). \tag{6}$$

Then,

$$\begin{aligned} U_L(q_L) - U_H(q_H) &= \pi_p(1-q_H)(H-p) + (1-\pi_p)(1-q_L)(L-p) + c(q_H) - c(q_L) \\ &= (1-q_L)(\mathbb{E}_{\pi_p}[\mathbf{v}] - p) - \pi_p(q_H - q_L)(H-p) + c(q_H) - c(q_L), \end{aligned}$$

where we note that $\mathbb{E}_{\pi_p}[\mathbf{v}] = \pi_p(H-p) + (1-\pi_p)(L-p)$.

Now we show that B strictly prefers not to acquire the H-focused signal.

Lemma 4. *Assume that $c(q) = \lambda q^2/2$. For any small enough $\lambda > 0$, it holds that $\max_{q_L} U_L(q_L) > \max_{q_H} U_H(q_H)$, where an optimal q_L is strictly positive.*

Proof. Let $q_L \in \operatorname{argmax}_{q'_L} U_L(q'_L)$ and $q_H \in \operatorname{argmax}_{q'_H} U_H(q'_H)$. Since they satisfy the first-order conditions,

$$(1-\pi_p)(p-L) = \lambda q_L. \tag{7}$$

$$\pi_p(H-p) = \lambda q_H, \tag{8}$$

where we note that $c'(q) = \lambda q$. By subtracting equation (7) from equation (8),

$$\mathbb{E}_{\pi_p}[\mathbf{v}] - p = \lambda(q_H - q_L). \tag{9}$$

Hence,

$$\begin{aligned} U_L(q_L) - U_H(q_H) &= (1-q_L)\lambda(q_H-q_L) - \pi_p(q_H-q_L)(H-p) + \frac{\lambda}{2}(q_H-q_L)(q_H+q_L) \\ &= (q_H-q_L)\left[(1-q_L)\lambda - \pi_p(H-p) + \frac{\lambda}{2}(q_H+q_L)\right]. \end{aligned}$$

Since $\pi_p(H-p) = \lambda q_H$ by equation (8),

$$U_L(q_L) - U_H(q_H) = \lambda(q_H-q_L)\left[1 - \frac{q_H+q_L}{2}\right]. \tag{10}$$

Hence, $U_L(q_L) > U_H(q_H)$ if and only if $q_H > q_L$ because the square bracket is strictly positive for any $q_H, q_L \in [0,1]$.

We show that for any small $\epsilon > 0$, there exists a small λ such that $q_H > q_L$. Note that the price p and the L-focused accuracy q_L satisfies (2), where we replace p^* and q^* with p and q_L. This is because (2) holds at the PBE under the L-focused signal. By (2) and (7), $(1-\pi_p)p = \lambda$. By (2), $1-\pi_p = \lambda x_L/L$, with $x_L \equiv 1-q_L$. Thus,

$$\mathbb{E}_{\pi_p}[\mathbf{v}] - p = H - (1-\pi_p)(H-L) - p = H - \lambda(H-L)\frac{x_L}{L} - \frac{L}{x_L}. \tag{11}$$

Recall, from the proof of Proposition 1, that $q_L \in (1-\frac{L}{H}-\epsilon, 1-\frac{L}{H})$ in the PBEs under the L-focused signal. Equivalently, $x_L \in (\frac{L}{H}, \frac{L}{H}+\epsilon)$. We define a function $\tilde{\lambda} : [\frac{L}{H}, \frac{L}{H}+\epsilon] \to \mathbb{R}_+$ by $\tilde{\lambda}(x) = \tilde{a}x + \tilde{b}$ with slope $\tilde{a} > 0$ such that $\tilde{\lambda}(\frac{L}{H}) = 0$ and

$$-\frac{H-L}{L}\left(\frac{L}{H}+\epsilon\right)\tilde{a} - \frac{H-L}{L}\tilde{\lambda}\left(\frac{L}{H}+\epsilon\right) + L\left(\frac{L}{H}+\epsilon\right)^{-2} > 0. \tag{12}$$

Let

$$\Delta(x) = H - \tilde{\lambda}(x)(H-L)\frac{x}{L} - \frac{L}{x}$$

By (11), $\Delta(x_L)$ is B's net gain from purchasing at price p when the unit cost λ equals $\tilde{\lambda}(x_L)$. For all $x \in [\frac{L}{H}, \frac{L}{H} + \epsilon]$,

$$\begin{aligned}\frac{d\Delta}{dx} &= -\frac{H-L}{L}\tilde{a}x - \frac{H-L}{L}\tilde{\lambda}(x) + \frac{L}{x^2} \\ &> -\frac{H-L}{L}\tilde{a}\left(\frac{L}{H}+\epsilon\right) - \frac{H-L}{L}\tilde{\lambda}\left(\frac{L}{H}+\epsilon\right) + L\left(\frac{L}{H}+\epsilon\right)^{-2} > 0,\end{aligned}$$

where the first and second inequalities are because $\tilde{\lambda}$ is strictly increasing and satisfies (12) respectively. Since $\Delta(\frac{L}{H}) = 0$, it follows that $\Delta(x) > 0$ for all $x \in (\frac{L}{H}, \frac{L}{H} + \epsilon]$. That is, if we take the unit cost $\tilde{\lambda}(x_L)$, B's net gain from purchasing at price p is strictly positive. By (9), $q_H > q_L$. By (10), $U_L(q_L) - U_H(q_H) > 0$ since $q_L \leq 1$ and $q_H \leq 1$. Hence, B does not want to choose any $q_H > 0$.

Since $\tilde{\lambda}$ is a strictly increasing affine function, we have, for all $\lambda \in (0, \tilde{\lambda}(\frac{L}{H} + \epsilon))$, a semi-separating equilibrium constructed above, in which B chooses accuracy $q_L = 1 - (\tilde{\lambda})^{-1}(\lambda)$ for the L-focused signal (and zero accuracy for the H-focused signal), such that B does not want to deviate to choose a positive q_H. ■

A.4 Proposition 3

Note that B does not acquire any post-offer information after observing price L or H, which is summarized below:

Lemma 5. *Every equilibrium* $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$ *is such that B chooses post-offer accuracy* $\beta^*(h_1) = 0$ *at any history* $h_1 = (p, q_0, x_0)$ *with price* $p = L, H$.[19]

We show that B does not acquire any post-offer information.

Lemma 6. *Every equilibrium* $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$ *is such that on any equilibrium path in which price* $p^* \in \bigcup_v \operatorname{supp}(\sigma^*(\cdot \mid v))$ *is offered, B chooses post-offer accuracy* $\beta^*(h_1) = 0$ *at history* $h_1 = (p^*, q_0^*, N)$ *with pre-offer accuracy* $q_0^* = \beta^*(h_0)$*, which gives a signal realization* $\mathbf{x}_0 = N$.

Proof. Suppose, by negation, that there exists some equilibrium $\mathcal{E}$ such that on some equilibrium path where price $p^* \in \bigcup_v \operatorname{supp}(\sigma^*(\cdot \mid v))$ is offered, B chooses post-offer accuracy $\beta^*(h_1) > 0$ at history $h_1 = (p^*, q_0^*, N)$. Then, $p^* \in (L, H)$ by Lemma 5.

We claim that if B acquires post-offer information on the equilibrium path, then on the continuation path, she buys if $\mathbf{x}_1 = H$ and does not if $\mathbf{x}_1 = N$. This is because if B's purchase decision were independent of a realization of $\mathbf{x}_1$, she would profitably deviate to choose post-offer accuracy $q = 0$, by which she saves the cost without changing her purchase

---

[19] We do not claim that this history h_1 is on the equilibrium path.

decision. This means that if type L of S charges price p^* then he makes zero profit because given type L, B observes a signal realization $\mathbf{x}_2 = N$ with probability 1 and thus never buys. Type L strictly prefer to charging price L to price p^*.[20] Hence, price p^* must come from type H of S and therefore reveals type H. But then, B will not acquire any post-offer information since information is useless but costly. This contradicts $\beta^*(h_1) > 0$ at history $h_1 = (p^*, q_0^*, N)$. ■

We show that on any equilibrium path, along which she acquires no post-offer information (Lemma 6), she buys at the offered price.

Lemma 7. *Take any equilibrium $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$. At any on-path history h_2 such that the offered price is strictly less than H, B buys at the offered price with probability $\beta^*(h_2) = 1$.*

Proof. By Lemma 6, B chooses post-offer accuracy $\beta^*(h_1) = 0$ at history $h_1 = (p^*, q_0^*, N)$ on any equilibrium path. Hence, the continuation time-2 history on the equilibrium path is such that $h_2 = (h_1, 0, N)$.

To show that B buys at price p^* with probability $\beta^*(h_2) = 1$, we suppose, by negation, that B buys with probability $\beta^*(h_2) < 1$. Since type L of S never charges any price at which B buys with probability 0, it must be that either (i) p^* comes only from type H of S or (ii) it comes from both types of S and $\beta^*(h_1, 0, N) \in (0, 1)$. In Case (i), p^* reveals type H; and since this lemma assumes $p^* < H$, we have $\beta^*(h_1, 0, N) = 1$. We now consider Case (ii). Since $\beta^*(h_1, 0, N) \in (0, 1)$, B is indifferent between buying and not buying; moreover, since not buying yields payoff 0 (excluding the cost of pre-offer information), B's equilibrium payoff is zero. Consider B's payoff from the deviation that B chooses accuracy q and buys the good if and only if she observes $\mathbf{x} = H$. This deviation gives B the expected payoff

$$\mu^*(p^*)[q \cdot (H - p^*) + (1 - q) \cdot 0] + (1 - \mu^*(p^*)) \cdot 0 - \lambda c(q) = \mu^*(p^*)q(H - p^*) - \lambda c(q),$$

which excludes the cost $\lambda c(\beta(\varnothing))$ of pre-offer information acquisition. Note that $\mu^*(p^*)$ denotes B's equilibrium belief of $\mathbf{v} = H$ at history h_1. Since $\mu^*(p^*) \in (0, 1)$ and $c'(0) = 0$, the optimal q that maximizes the above payoff is strictly positive (and the maximum payoff is strictly positive). This contradicts B's equilibrium payoff being equal to zero. ■

Lemma 8. *For any equilibrium $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$, B does not acquire pre-offer information and there exists some $p^* \in [L, H)$ such that*

$$\bigcup_{v \in V} \operatorname{supp}(\sigma^*(\cdot \mid v)) \setminus \{H\} = \{p^*\}$$

---

[20] When observing price L, B would be indifferent between buying and not (because she infers type $\mathbf{v} = L$). Even if she does not buy at price L, type L of S can give B a strict incentive to buy by cutting the price infinitesimally.

Proof. By Lemmas 6 and 7, B chooses post-offer accuracy $\beta^*(h_1) = 0$ on any equilibrium path h_1 of every equilibrium $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$ and, if $p < H$ is offered, B buys with probability 1, while if $p = H$ is offered, B buys with a positive probability (otherwise, type H does not offer H in equilibrium). This means that any pre-offer information is useless for B's purchase decision. Since information is costly, B acquires no pre-offer information. To show that $\bigcup_{v \in V} \operatorname{supp}(\sigma^*(\cdot \mid v)) \setminus \{H\}$ is singleton. Suppose, by negation, that $\bigcup_{v \in V} \operatorname{supp}(\sigma^*(\cdot \mid v)) \setminus \{H\}$ is not a singleton. Then, we can take any two prices $p \neq p'$ from this set. By Lemmas 6, 7 and the first part of this lemma, at prices p and p', B will acquire no information and buy for sure: $\beta^*(p) = \beta^*(p') = 0$ and $\beta^*(p, 0, N) = \beta^*(p', 0, N) = 1$. However, this means that S would profitably choose the higher price only. A contradiction. ■

Finally, we prove the following lemma.

Lemma 9. *Any PBE such that* $H \in \bigcup_{v \in V} \operatorname{supp}(\sigma^*(\cdot \mid v))$ *is Pareto-dominated.*

Proof. Consider a PBE in which $H \in \bigcup_{v \in V} \operatorname{supp}(\sigma^*(\cdot \mid v))$. Price H must come only from type H of S, because otherwise, B would not buy at price H and thus type H will not charge H. By Lemma 8, there is exactly one equilibrium price $p^* < H$.

We argue that at price H, B buys with probability $y = p^*/H$. Indeed, if $y > p^*/H$ then type L would strictly prefer to charge price H and get profit $yH > p^*$; if $y < p^*/H$, then type H would strictly prefer to charge price p^* than H. The equilibrium must satisfy the following conditions:

1. Type H of S charges price $p \in \operatorname{supp}(\sigma^*(\cdot \mid H))$, and type L charges price p^*.
2. B acquires no information on the equilibrium path:
   (a) At any on-path history h_1 or h_2, B chooses accuracy $q_1 = q_2 = 0$.
   (b) At any on-path history h_3, if $p = p^*$, B buys with probability 1; if $p = H$, she buys with probability p^*/H.

We now show that the above equilibrium is Pareto-dominated by the following pooling equilibrium:

1. Each type v of S offers price p^*.
2. B acquires no information on the equilibrium path.
   (a) At any history h_1 or h_2 (on or off path), B chooses accuracy $q = 0$.
   (b) At history h_3, if $p = p^*$ or $p = L$, B buys with probability 1; for any other price in (L, H), she assigns probability 1 to type L, does not acquire information, and buys with probability 0. ■

Note that if the former assessment is an equilibrium, so is the latter. This is because the conditions for the latter to be an equilibrium—that is, type L of S is willing to offer price p^* rather than any other price, and B is willing to buy without acquiring any information at price p^*—have to be satisfied in the former equilibrium as well. We can now see that the former equilibrium is Pareto-dominated by the latter, because each type v of S has payoff p^* and thus their payoffs are the same in both equilibria, but B's payoff is strictly higher in the latter equilibrium.

Now we prove Proposition 3. Each type v of S offers a pooling price p^* such that $\bigcup_{v \in V} \operatorname{supp}(\sigma^*(\cdot \mid v)) = \{p^*\}$ (Lemma 8). B chooses post-offer accuracy 0 and buys at price p^* with probability 1 (Lemmas 6 and 7). Then, B's equilibrium payoff is equal to

$$\mathbb{E}_{\mathbf{h}_3}[\mathbb{E}_{\mathbf{v} \sim \mu(\cdot \mid \mathbf{h}_3)}[\mathbf{v}]] - p^* - \lambda c(\beta^*(\varnothing)),$$

where $\mathbf{h}_2 = (p^*, \beta^*(\varnothing), \mathbf{x}_0, 0, N)$ is the time-2 history on the equilibrium path. This history may be possibly random because if B chooses pre-offer accuracy $\beta^*(\varnothing) > 0$ then signal $\mathbf{x}_0$ is random. Since B's belief process is a martingale, it must be that

$$\mathbb{E}_{\mathbf{h}_3}[\mathbb{E}_{\mathbf{v} \sim \mu(\cdot \mid \mathbf{h}_3)}[\mathbf{v}]] = \mathbb{E}_{\mathbf{v} \sim \pi}[\mathbf{v}].$$

Hence, B's payoff $\mathbb{E}_{\mathbf{v} \sim \pi}[\mathbf{v}] - p^* - \lambda c(\beta^*(\varnothing))$ is maximized at accuracy $\beta^*(\varnothing) = 0$.

As claimed in Proposition 3, in every equilibrium, each type v of S offers a pure-strategy pooling price p^*, and B sets pre- and post-offer accuracy equal to zero and then buys at price p^* with probability 1. By exactly the same argument as in the proof of Theorem 1, price p^* is close to L, by which B extracts full surplus as the cost parameter λ vanishes.

B Noisy H-Focused Signal Structure

The main text has assumed that B accesses the *exact* H-focused signal structure, in which the signal realization H is conclusive. That is, type L of S for sure generates the signal realization L. We make use of this property, and it significantly simplifies our argument, but here we show that Theorem 1 holds robustly even if type L can generate the signal realization H with a small probability. Formally, if B chooses accuracy q then her signal structure is as follows:

$$\begin{aligned} &\mathbb{P}(\mathbf{x} = H \mid \mathbf{v} = H) = q, && \mathbb{P}(\mathbf{x} = H \mid \mathbf{v} = L) = \eta q, \\ &\mathbb{P}(\mathbf{x} = N \mid \mathbf{v} = H) = 1 - q, && \mathbb{P}(\mathbf{x} = N \mid \mathbf{v} = L) = 1 - \eta q, \end{aligned}$$

given a small $\eta \geq 0$. This variant coincides with the original model with the exact H-focused signal structure if and only if $\eta = 0$.

Extension of Theorem 1 We show that Theorem 1 remains true for any $\eta \in [0, L/H)$. The proof is analogous to the original proof in Appendix A.1.

Lemmas 2 and 3 are independent of the assumption that $\eta = 0$ and thus remain true for any $\eta \in [0, L/H)$. Moreover, Step 2 in Appendix A.1 also remains true. Hence, it suffices to extend Lemma 1.

Lemma 10. *Consider the noisy H-focused signal structure with $\eta \in [0, L/H)$. Given any PBE $\hat{\mathcal{E}} = (\hat{\beta}, \hat{\sigma}, \hat{\mu})$, B chooses accuracy $\hat{\beta}(p) = 0$ after any on-path price offer $p \in \bigcup_{v \in V} \operatorname{supp}(\hat{\sigma}(\cdot \mid v))$.*

Proof. Let a PBE $\hat{\mathcal{E}} = (\hat{\beta}, \hat{\sigma}, \hat{\mu})$ be given. To prove the lemma, it suffices to show that B chooses accuracy $\hat{\beta}(p) = 0$ after any price offer $p \in \operatorname{supp}(\hat{\sigma}(\cdot \mid L))$ that type L of S may offer.

If B is offered price $p = L, H$ then she does not acquire information for any $\lambda > 0$. Then, we consider the case where B is offered a price $p \in (L, H) \cap \operatorname{supp}(\hat{\sigma}(\cdot \mid L))$, which type L may offer. Suppose, by contradiction, that B chooses accuracy $q > 0$ after that p. Then, B's purchase decision depends on the realization of her signal $\mathbf{x}$. That is, B must buy if $\mathbf{x} = H$ but never buys if $\mathbf{x} = N$. This is because if her purchase decision were independent of the realization of $\mathbf{x}$, she would profitably deviate to choosing zero accuracy, which saves the information cost without changing her purchase decision. However, this purchase decision implies that type L of S makes a profit of $p \cdot \eta q$ because B buys only in the event of $\mathbf{x} = H$ (that has a probability of ηq). For any $\eta \leq L/H$, we have $p \cdot \eta q < H \cdot L/H = L$. That is, type L has a profitable deviation of offering price L. A contradiction. ■

As stated above, since the remaining argument in Appendix A.1 remains valid, we have Theorem 1 under the noisy H-focused signal structure.

References

K. Bagwell. Commitment and observability in games. *Games and Economic Behavior*, 8(2): 271–280, 1995.

S. Bikhchandani. A bargaining model with incomplete information. *Review of Economic Studies*, 59(1):187–203, 1992.

X. Cheng and Y. Kim. Blackwell-monotone information costs. *Working Paper*, 2024. doi: 10.48550/arXiv.2404.15158.

O. Compte and P. Jehiel. Auctions and information acquisition: sealed bid or dynamic formats? *Rand Journal of Economics*, 38(2):355–372, 2007.

J. Crémer and F. Khalil. Gathering information before signing a contract. *American Economic Review*, 82(3):566–578, 1992.

J. Crémer and F. Khalil. Gathering information before the contract is offered: The case with two states of nature. *European Economic Review*, 38(3-4):675–682, 1994.

J. Crémer, F. Khalil, and J.-C. Rochet. Strategic information gathering before a contract is offered. *Journal of Economic Theory*, 81(1):163–200, 1998.

Federal Trade Commission. Buying a used car from a dealer, Aug. 2022. URL `https://consumer.ftc.gov/articles/buying-used-car-dealer#Inspection`. Accessed: 2024-06-27.

Y. Feinberg and A. Skrzypacz. Uncertainty about uncertainty and delay in bargaining. *Econometrica*, 73(1):69–91, 2005.

D. Fudenberg and J. Tirole. *Game Theory*. MIT Press, 1991a.

D. Fudenberg and J. Tirole. Perfect Bayesian equilibrium and sequential equilibrium. *Journal of Economic Theory*, 53(2):236–260, 1991b.

D. Fudenberg, D. K. Levine, and J. Tirole. Infinite-horizon models of bargaining with one-sided incomplete information. In A. Roth, editor, *Game Theoretic Models of Bargaining*, pages 73–98. Cambridge University Press, Cambridge, UK and New York, 1985.

F. Gul, H. Sonnenschein, and R. Wilson. Foundations of dynamic monopoly and the Coase conjecture. *Journal of Economic Theory*, 39(1):155–190, 1986.

A. S. Kessler. The value of ignorance. *RAND Journal of Economics*, pages 339–354, 1998.

E. Maskin and J. Tirole. The principal-agent relationship with an informed principal: The case of private values. *Econometrica*, 58(2):379–409, 1990.

R. B. Myerson. Mechanism design by an informed principal. *Econometrica*, 51(6):1767–1797, 1983.

N. Persico. Information acquisition in auctions. *Econometrica*, 68(1):135–148, 2000.

D. Ravid. Ultimatum bargaining with rational inattention. *American Economic Review*, 110 (9):2948–63, 2020.

D. Ravid, A.-K. Roesler, and B. Szentes. Learning before trading: on the inefficiency of ignoring free information. *Journal of Political Economy* (forthcoming), 2021.

A.-K. Roesler and B. Szentes. Buyer-optimal learning and monopoly pricing. *American Economic Review*, 107(7):2072–80, 2017.

A. Rubinstein. Perfect equilibrium in a bargaining model. *Econometrica*, 50(1):97–109, 1982.

A. Rubinstein. A bargaining model with incomplete information about time preferences. *Econometrica*, 53(5):1151–1172, 1985.

X. Shi. Optimal auctions with information acquisition. *Games and Economic Behavior*, 74 (2):666–686, 2012.

A. Wolitzky. Unobserved-offers bargaining. *American Economic Review*, 113(1):136–173, 2023.

M. Yildiz. Bargaining without a common prior—an immediate agreement theorem. *Econometrica*, 71(3):793–811, 2003.

C Proof of Proposition 4 (Not for Publication)

C.1 Case of $\lim_{q\to 1} c'(q) = \infty$

We prove Proposition 4 under the assumption that $c'(1) \equiv \lim_{q\to 1} c'(q) = \infty$, under which B will never choose accuracy 1.

Our proof consists of three steps. Consider our model given a cost parameter $\lambda > 0$. First, we derive necessary conditions that any mixed-learning PBE must satisfy. Using the properties, we construct a mixed-learning PBE, denoted $\mathcal{E}_\lambda(p^*)$ given a price p^* that S may offer, which is tractable. Second, we show that the PBE $\mathcal{E}_\lambda(p^*)$ is Pareto-undominated for a sufficiently small $\lambda > 0$. Third, we examine B's (ex-ante) expected payoff in the PBE $\mathcal{E}_\lambda(p^*)$.

Step 1 Consider our model with a cost parameter $\lambda > 0$. We derive some necessary conditions that any mixed-learning PBE must satisfy. These properties are used not only to construct a tractable mixed-learning PBE but also to prove that the PBE is Pareto-undominated.

Lemma 11. *For any mixed-learning PBE, if B randomizes information acquisition after an equilibrium price $p \in (L, H)$, the following holds after B is offered price p:*

1. *B randomizes over two accuracies 0 and $\tilde{q}(p)$; that is, her strategy β is such that $\operatorname{supp}(\beta(\cdot \mid p)) = \{0, \tilde{q}(p)\}$, where $\tilde{q} : (L, H) \to (0, 1)$ is the function defined by*

$$\lambda c'(\tilde{q}(p))\left(\frac{H-L}{H-p} - \tilde{q}(p)\right) + \lambda c(\tilde{q}(p)) = p - L. \tag{13}$$

   This implicit function $\tilde{q}$ is well-defined.

2. *B's posterior probability that S is of type H after observing the price p, denoted $\tilde{\pi}_1(p)$, satisfies equation*

$$\tilde{\pi}_1(p) = \frac{\lambda c'(\tilde{q}(p))}{H-p}. \tag{14}$$

Proof. Consider any mixed-learning PBE, at which B randomizes information acquisition after some price offer p. Suppose that she chooses to acquire information. If B chooses accuracy $q > 0$ and buys if a signal realization is $\mathbf{x} = H$ and never buys if $\mathbf{x} = N$ then her payoff is $\tilde{\pi}_1(p)q(H-p) - \lambda c(q)$. In the equilibrium, the accuracy $q = \tilde{q}(p)$ after the price p must maximize this payoff. Hence, it satisfies the first-order condition $\tilde{\pi}_1(p)(H-p) = \lambda c'(\tilde{q}(p))$, which gives the desired equation (14).

Next, suppose that she acquires no information. That is, she chooses accuracy 0, after which she buys with probability 1. Her payoff is $\tilde{\pi}_1(p)(H-p) + (1 - \tilde{\pi}_1(p))(L-p)$.

B must be indifferent between the two accuracies $\tilde{q}(p)$ and 0 since B randomizes her choice of accuracies. That is,

$$\tilde{\pi}_1(p)(1-\tilde{q}(p))(H-p)+(1-\tilde{\pi}_1(p))(L-p)+\lambda c(\tilde{q}(p))=0. \tag{15}$$

Substituting (14) into (15), we obtain the desired equation (13).

It remains to show that the implicit function $\tilde{q}$ is well-defined. That is, we show that for any $p \in (L,H)$, there exists a unique $q \in (0,1)$ that solves equation (13). Since c is strictly convex and $\frac{H-L}{H-p}-q > 1-q > 0$, the LHS of (13) is strictly increasing in $\tilde{q}(p)$. It is also continuous in $\tilde{q}(p)$. Moreover,

$$\begin{aligned}
&\lambda c'(0)((H-L)/(H-p)-0)+\lambda c(0)=0<p-L,\\
&\lambda c'(1)((H-L)/(H-p)-1)+\lambda c(1)=\infty>p-L,
\end{aligned}$$

which ensures the existence and uniqueness of q that solves (13). ■

Lemma 12. *The functions $\tilde{\pi}_1$ and $\tilde{q}$, defined by equations (13) and (14), satisfy the following properties:*

1. $\lim_{p\to L}\tilde{q}(p)=0$ *and* $\lim_{p\to H}\tilde{q}(p)=0$ *for any* $\lambda>0$.
2. $\lim_{p\to L}\tilde{\pi}_1(p)=0$ *and* $\lim_{p\to H}\tilde{\pi}_1(p)=1$ *for any* $\lambda>0$.
3. $\tilde{\pi}_1(p)$ *is continuous and strictly increasing for any* $\lambda>0$.
4. $\lim_{\lambda\to 0}\tilde{q}(p)=1$, $\lim_{\lambda\to 0}\tilde{\pi}_1(p)=1$, *and* $\lim_{\lambda\to 0}\tilde{\pi}_1'(p)=0$ *for any* $p\in(L,H)$.

Proof. The first claim is immediate from (13). We prove the second claim. Since $\lim_{p\to L}\tilde{q}(p)=0$, we have $\lim_{p\to L}\tilde{\pi}_1(p)=\frac{\lambda c'(0)}{H-L}=0$. Since $\lim_{p\to H}\tilde{q}(p)=0$ and (13) is equivalent to $\tilde{\pi}_1(p)(H-L)-\lambda c'(\tilde{q}(p))\tilde{q}(p)=p-L$, we have $\lim_{p\to H}\tilde{\pi}_1(p)=1$.

We show the third claim. Since the continuity is obvious, we prove that it is strictly increasing. By the implicit function theorem applied to the function $\tilde{q}$, as defined in (13),

$$\tilde{q}'(p)=-\frac{\lambda c'(\tilde{q}(p))\frac{H-L}{(H-p)^2}-1}{\lambda c''(\tilde{q}(p))(\frac{H-L}{H-p}-\tilde{q}(p))}. \tag{16}$$

Substituting it into (14), we have

$$\tilde{\pi}_1'(p)=\frac{1-\tilde{q}(p)\frac{\lambda c'(\tilde{q}(p))}{H-p}}{H-L-\tilde{q}(p)(H-p)}=\frac{1-\tilde{q}(p)\tilde{\pi}_1(p)}{H-L-\tilde{q}(p)(H-p)}. \tag{17}$$

Note that $\tilde{\pi}_1'(p)>0$ for any p such that $\tilde{\pi}_1(p)\le 1$. This is because both the denominator and the numerator of the RHS of (17) is strictly positive. Hence, to show that $\tilde{\pi}_1(p)<1$ for all

$p \in (L,H)$, it suffices to show that $\tilde{\pi}_1(p) < 1$ for all $p \in (L,H)$. Suppose, by negation, that there is some $\hat{p} \in (L,H)$ such that $\tilde{\pi}_1(\hat{p}) = 1$. Then, $\tilde{\pi}_1(p) > 1$ for all $p \in (\hat{p}, H)$ (because if $\tilde{\pi}_1(p) = 1$, we must have $\tilde{\pi}_1(p) > 0$). Since $\lim_{p\to H} \tilde{\pi}_1(p) = 1$ and $\tilde{\pi}_1 > 1$ on $(\hat{p}, H)$, $\tilde{\pi}_1$ must be weakly decreasing on a neighborhood of H. However, applying $\lim_{p\to H} \tilde{q}(p) = 0$ and $\lim_{p\to H} \tilde{\pi}_1(p) = 1$ to the last expression of (17), we have that $\tilde{\pi}_1'(H) > 0$, a contradiction.

We prove the fourth claim. Let $\lambda \to 0$. If $\tilde{q}(p) \not\to 1$ then the LHS of (13) would converge to zero, but the RHS is $p - L > 0$ for any $p \in (L,H)$. This is a contradiction, and thus $\tilde{q}(p) \to 1$. To show that $\tilde{\pi}_1(p) \to 1$, rewrite (13) as

$$\lambda c(\tilde{q}(p))\left[\frac{c'(\tilde{q}(p))}{c(\tilde{q}(p))}\left(\frac{H-L}{H-p} - q(p)\right) + 1\right] = p - L. \tag{18}$$

Since $\lim_{q\to 1} c'(q) = \infty$, we have $\lim_{q\to 1} \frac{c'(q)}{c(q)} = \infty$.[21] For any fixed $p \in (L,H)$, taking the limit as $\lambda \to 0$, we have $\tilde{q}(p) \to 1$, and thus the term in the square brackets of (18) goes to infinity. Since the RHS is finite, we have $\lambda c(\tilde{q}(p)) \to 0$. Using (14), we can rewrite (13) as

$$\tilde{\pi}_1(p)(H - L - \tilde{q}(p)(H-p)) + \lambda c(\tilde{q}(p)) = p - L.$$

Taking limit as $\lambda \to 0$ and applying $\tilde{q}(p) \to 1$ and $\lambda c(\tilde{q}(p)) \to 0$, we have $\tilde{\pi}_1(p) \to 1$.

Finally, taking the limit as $\lambda \to 0$ on both sides of (17) and applying $\tilde{q}(p) \to 1$ and $\tilde{\pi}_1(p) \to 1$, we have $\tilde{\pi}_1'(p) \to 1$. ■

Next, we construct a mixed-learning PBE.

Lemma 13. *Given any $\lambda > 0$, there exists some $\underline{p}_\lambda \in (L,H)$ such that for any $p^* \in (\underline{p}_\lambda, H)$, the following assessment $\mathcal{E}_\lambda(p^*)$ is a PBE:*

1. *Type H of S offers a price p^* with probability 1, and type L offers prices p^* and L with probabilities y^* and $1-y^*$, respectively, where $y^* \in (0,1)$ solves equation*

$$\tilde{\pi}_1(p^*) = \frac{\pi}{\pi + (1-\pi)y^*}. \tag{19}$$

2. *If B is offered price p^* then:*
    - *With probability $z^* = L/p^*$, B chooses accuracy 0 and buys with probability 1.*
    - *With probability $1 - z^*$, B chooses accuracy $\tilde{q}(p^*)$ and buys with probability 1 if a signal realization is $\mathbf{x} = H$ and never buys if $\mathbf{x} = N$.*

---

[21] We show that $\lim_{q\to 1} c'(q)/c(q) = \infty$. Since the claim is trivial if $\lim_{q\to 1} c(q) < \infty$, let $\lim_{q\to 1} c(q) = \infty$. Suppose, for a contradiction, that $\lim_{q\to 1} c'(q)/c(q) < \infty$. Then, for some $M > 0$, $c'(q)/c(q) \le M$ for all q sufficiently close to 1. For any $q_0 \in [0,1)$ that is sufficiently close to 1, integrating both sides, we have $\log(c(q)/c(q_0)) \le M(q - q_0)$ and thus $c(q) \le c(q_0)\mathrm{e}^{M(q-q_0)}$, but this contradicts the assumption that $\lim_{q\to 1} c(q) = \infty$.

If B is offered any price $p \neq p^*$ then she assigns probability 1 to type L and chooses accuracy 0 and buys if and only if price p is at most L.

Moreover, $\underline{p}_\lambda \to L$ as $\lambda \to 0$.

Proof. We derive the necessary and sufficient conditions for this assessment $\mathcal{E}_\lambda(p^*)$ to be a PBE. First, we note that if B is offered the price p^* then she randomizes over two accuracies 0 and $\tilde{q}(p^*)$ by Lemma 11.

Second, we derive (19). In the assessment, type L of S offers prices p^* and L with probabilities y^* and $1 - y^*$, respectively. Then, B's posterior probability (that S is of type H) at price p^* is $\frac{\pi}{\pi+(1-\pi)y^*}$. By Lemma 11, this posterior probability, which we have denoted by $\tilde{\pi}_1(p^*)$, must satisfy (14). Since these two representations must coincide,

$$\tilde{\pi}_1(p^*) = \frac{\pi}{\pi + (1-\pi)y^*},$$

which is the desired (19).

We show that there exists $\underline{p}_\lambda \in (L, H)$ such that for any $p^* \in (\underline{p}_\lambda, H)$, (19) has a solution y^*. By Lemma 12, $\tilde{\pi}_1$ is continuous and strictly increasing, and $\lim_{p\downarrow L} \tilde{\pi}_1(p) = 0$ and $\lim_{p\uparrow H} \tilde{\pi}_1(p) = 1$. Hence, there must exist a unique $\underline{p}_\lambda \in (L, H)$ such that $\tilde{\pi}_1(\underline{p}_\lambda) = \pi$, where we recall that $\pi \in (0, 1)$ is the prior probability. Then, we have $\tilde{\pi}_1(p^*) \in (\pi, 1)$ since $p^* \in (\underline{p}_\lambda, H)$ by assumption. Since the function $(0, 1) \ni y \mapsto \frac{\pi}{\pi+(1-\pi)y} \in (\pi, 1)$ is strictly decreasing and continuous, we must have some y^* that satisfies (19).

Third, we see that S has no profitable deviation. Type L is willing to randomize between prices L and p^* if and only if he gains the same profit from both prices. That is, $L = p^*z^*$ because type L makes sales only when B does not acquire information. Hence,

$$z^* = L/p^*,$$

as desired. Type H gains a profit of $p^*(z^* + (1 - z^*)q^*)$. We show that he has no profitable deviation. Indeed, any deviation would yield a profit of at most L, but $p^*(z^*+(1-z^*)q^*) > L$. This is because, for $z^* = L/p^*$, this inequality is reduced to $z^* < 1$.

Lastly, we show that $\underline{p}_\lambda \to L$ as $\lambda \to 0$. For any $p \in (L, H)$, $\tilde{q}(p) \to 1$ and $\tilde{\pi}_1(p) \to 1$ as $\lambda \to 0$ by Lemma 12. By the definition of $\underline{p}_\lambda$, it follows that $\underline{p}_\lambda \to L$. ■

Step 2 We show the Pareto-undominance of PBE $\mathcal{E}_\lambda(p^*)$, which we construct in Lemma 13.

Lemma 14. *For each $p^* \in (L, H)$, if λ is sufficiently small then the PBE $\mathcal{E}_\lambda(p^*)$ is Pareto-undominated.*

Proof. We prove this lemma in seven steps.

Step 1. In any mixed-learning PBE, B must randomize information acquisition after any equilibrium price offer $p' \in (L, H)$.

Proof: Suppose, by contradiction, that there exists a mixed-learning PBE $\mathcal{E}$ such that B does not randomize information acquisition after some equilibrium price $p' \in (L, H)$. Note that p' must be in the support of the prices offered by type H of S (because otherwise, B would never buy as she is sure that S is of type L and thus S would profitably deviate to offering price L). Next, B must choose accuracy 0 after the price offer p'. This is because otherwise, since B would acquire information for sure (as she does not randomize information acquisition), type L would have a profitable deviation of offering price L (since type L makes no sale. Hence, B chooses accuracy 0 after the price offer p'. Let $\alpha \geq 0$ be the probability that B buys the item after the price offer p'.

There is some price p after which B randomizes information acquisition, since $\mathcal{E}$ is a mixed-learning PBE. Then, p is in the support of the prices offered by both types of S, otherwise B would be sure about the type of S. Moreover, it must be that $p \in (L, H)$, otherwise B would not acquire information. By Lemma 11, if price p is offered then B randomizes over two accuracies 0 and $\tilde{q}(p)$. She chooses accuracy 0 with probability z.

Since prices p and p' are in the support of the prices offered by type H of S, his profits from offering both prices are the same; that is, $p(z+(1-z)q) = p'\alpha$, where α is the probability that B buys (when she does not acquire information). It implies $pz < p'\alpha$. Note that pz and $p'\alpha$ are type L's profits from offering prices p and p', respectively. However, since $pz < p'\alpha$, type L must strictly prefer price p', which contradicts the fact that p is in the support of the prices offered by type L of S.

Step 2. The function $\tilde{q}$, as defined in (13), is unimodal. That is, there exists a unique $p_\lambda \in (L, H)$ such that $\tilde{q}$ is strictly increasing on the interval (L, p_λ) and strictly decreasing on the interval (p_λ, H). Moreover, $p_\lambda \to L$ as $\lambda \to 0$, which implies that for any fixed $p \in (L, H)$, if λ is sufficiently small then $\tilde{q}'(p) < 0$.

Proof: Recall the derivative $\tilde{q}'(p)$ given in (16). Since the denominator of the RHS in (16) is positive, $\tilde{q}'(p)$ is positive (resp. negative) if and only if its numerator, denoted $\tilde{f}(p)$, is negative (resp. positive), where

$$\tilde{f}(p) \equiv \lambda c'(\tilde{q}(p)) \frac{H-L}{(H-p)^2} - 1.$$

There exists at most one $p_\lambda \in (L, H)$ such that $\tilde{f}(p_\lambda) = 0$, or equivalently $\tilde{q}'(p_\lambda) = 0$. This is because if $\tilde{f}(p_\lambda) = 0$ and thus $\tilde{q}'(p_\lambda) = 0$ then $\tilde{f}'(p_\lambda) = 2\lambda c'(\tilde{q}(p_\lambda))(H-L)/(H-p_\lambda)^3 > 0$.

Moreover, there exists $p_\lambda \in (L, H)$ such that $\tilde{q}'(p_\lambda) = 0$. This is because by Lemma 12, $\tilde{q}(p) \to 0$ as $p \to L$ or $p \to H$ and $\tilde{q}(p) > 0$ for any $p \in (L, H)$. Therefore, we have established the existence and uniqueness of p_λ.

Now we show that $p_\lambda \to L$ as $\lambda \to 0$. By (14), $\tilde{f}(p) = \tilde{\pi}_1(p)(H - L)/(H - p) - 1$. Since $\tilde{\pi}_1(p) \to 1$ as $\lambda \to 0$ for any $p \in (L, H)$ by Lemma 12, it follows that $\tilde{f}(p) \to \frac{p-L}{H-p} > 0$ and thus $\tilde{q}'(p) < 0$, which implies that $p_\lambda \to L$.

Step 3. For any small $\delta > 0$, there exists some $\lambda_\delta > 0$ such that if $\lambda < \lambda_\delta$ then any mixed-learning PBE has at most one equilibrium price in the interval $(L, H - \delta)$. That is, the set, $\operatorname{supp}(\bigcup_v \sigma(\cdot \mid v)) \cap (L, H - \delta)$, is a singleton or an empty set for any of S's equilibrium strategy σ.

Proof: For any $u_L \in [L, H)$, let Γ_{u_L} be the set of all PBEs such that type L's payoff is u_L. Let $p \in (L, H - \delta)$ be an equilibrium price of some PBE in Γ_{u_L}. By **Step 1** with Lemma 11, B randomizes between accuracies 0 and $\tilde{q}(p)$ after price p is offered. Moreover, the following holds. First, the probability that B acquires no information, denoted $z(p)$, satisfies $u_L = pz(p)$, since $pz(p)$ is type L's payoff from offering price p. Second, type L's payoff from offering price p, denoted $\tilde{U}_H(p)$, is

$$\tilde{U}_H(p) = pz(p) + (1 - z(p))\tilde{q}(p)p = u_L + \tilde{q}(p)(p - u_L). \tag{20}$$

Now we show that for any small $\delta > 0$, there exists some $\lambda_\delta > 0$ such that for any $\lambda < \lambda_\delta$, the function $\tilde{U}_H$ is strictly increasing on the interval $(L, H - \delta)$. Note that

$$\tilde{U}'_H(p) = \tilde{q}(p) + \tilde{q}'(p)(p - u_L) = \tilde{q}(p)\left(1 + \frac{\tilde{q}'(p)}{\tilde{q}(p)}(p - u_L)\right).$$

By (16),

$$\tilde{q}'(p) > -\frac{c'(\tilde{q}(p))(H - L)}{c''(\tilde{q}(p))(H - p)(p - L)}.$$

Since $u_L \geq L$ and $H - p > \delta$,

$$\frac{\tilde{q}'(p)}{\tilde{q}(p)}(p - u_L) > -\frac{c'(\tilde{q}(p))}{c''(\tilde{q}(p))\tilde{q}(p)} \frac{H - L}{H - p} \frac{p - u_L}{p - L} > -\frac{c'(\tilde{q}(p))}{c''(\tilde{q}(p))\tilde{q}(p)} \frac{H - L}{\delta}.$$

By Lemma 12, $\tilde{q}(p) \to 1$ as $\lambda \to 0$. Since $\frac{c'(q)}{c''(q)} \to 0$ as $q \to 1$, it follows that $\frac{c'(\tilde{q}(p))}{c''(\tilde{q}(p))\tilde{q}(p)} \to 0$ as $\lambda \to 0$. Moreover, there exists $\eta > 0$ such that $\frac{c'(q)}{c''(q)} < \frac{\epsilon}{H-L}$ for all $q \in (1 - \eta, 1)$. Recall from **Step 2** that for any $p \in (L, H)$, if λ is sufficiently small then $\tilde{q}'(p) < 0$. Therefore, there exists $\lambda_\delta > 0$ such that if $\lambda < \lambda_\delta$ then for $p = H - \delta$, $\tilde{q}(p) > 1 - \eta$ and $\tilde{q}'(p) < 0$. By the definition of p_λ, we have $\tilde{q}'(p) < 0$ for any $p \in (p_\lambda, H - \delta)$. Since $\tilde{q}(H-\delta) > 1 - \eta$, we have $\tilde{q}(p) > 1 - \eta$ for any $p \in (p_\lambda, H - \delta)$, implying that $\frac{c'(\tilde{q}(p))}{c''(\tilde{q}(p))\tilde{q}(p)} < \frac{\delta}{H-L}$ for all $p \in (p_\lambda, H - \delta)$. Hence, if $\lambda < \lambda_\delta$, then

$$\frac{\tilde{q}'(p)}{\tilde{q}(p)}(p - u_L) > -\frac{c'(\tilde{q}(p))}{c''(\tilde{q}(p))\tilde{q}(p)} \frac{H - L}{\delta} > -1,$$

which implies that $\tilde{U}_H'(p) > 0$ for all $p \in (p_\lambda, H - \delta)$.

Take any equilibrium in Γ_{u_L}. Now we show that if $\lambda < \lambda_\delta$, then there is at most one equilibrium price in $(L, H - \delta)$. Indeed, if there were two equilibrium prices p and p' in $(L, H - \delta)$, then by Step 1, B randomizes information acquisition after both prices. This implies that type H of S receives the same payoff from offering p and p' (otherwise one of the price reveals type L and thus B would not acquire information); and type H payoff from offering prices p and p' are $\tilde{U}_H(p)$ and $\tilde{U}_H(p')$, respectively. But since $\tilde{U}_H'(\cdot) > 0$ on $\in (L, H - \delta)$ (for $\lambda < \lambda_\delta$), we have $\tilde{U}_H(p) \neq \tilde{U}_H(p')$, a contradiction.

Step 4. There exists $\lambda_{p,\delta} \in (0, \lambda_\delta)$ such that if $\lambda < \lambda_{p,\delta}$ then in any mixed-learning PBE with an equilibrium price $p \in (L, H - \delta)$, type L of S offers price L with a positive probability. Moreover, $\lambda_{p,\delta}$ weakly increases in p.

Proof. Take any mixed-learning PBE with an equilibrium price $p \in (L, H - \delta)$. B's posterior probability that S is of type H after price p is offered is $\tilde{\pi}_1(p|\lambda) \equiv \tilde{\pi}_1(p)$, where in this proof we write $\tilde{\pi}_1(p|\lambda)$ in order to be explicit about its dependence on λ. By Lemma 12, $\tilde{\pi}_1(p|\lambda) \to 1$ as $\lambda \to 0$. Thus, $\tilde{\pi}_1(p|\lambda) \geq \pi$ for any sufficiently small λ. Let $\lambda_p^1 \equiv \sup\{\lambda' > 0 : \tilde{\pi}_1(p|\lambda) \geq \pi, \forall \lambda < \lambda'\}$. That is, λ_p^1 is the highest λ' such that if $\lambda < \lambda'$, then $\tilde{\pi}_1(p'|\lambda) \geq \pi$. By Lemma 12, $\tilde{\pi}_1(p|\lambda)$ is strictly increasing in p, which implies that for any $p' > p$, if $\lambda < \lambda_p^1$ then $\tilde{\pi}_1(p|\lambda) > \pi$. By the definition of λ_p^1, this implies that λ_p^1 is increasing in p. Next, let $\lambda_{p,\delta} := \min\{\lambda_p^1, \lambda_\delta\}$. It follows that $\lambda_{p,\delta}$ weakly increases in p.

By the definition of $\lambda_{p,\delta}$, if $\lambda < \lambda_{p,\delta}$, then $\tilde{\pi}_1(p) > \pi$. Moreover, since $\tilde{\pi}_1$ is increasing, we have $\tilde{\pi}_1(p') > \pi$ for all $p' \in (p, H)$. For Bayes' rule to hold, there must be some price $p'' \in [L, p)$ such that $\tilde{\pi}_1(p'') < \pi$, implying that $p'' \in [L, p)$ is in the support of type L's strategy. Moreover, since $\lambda < \lambda_\delta$, there is at most one equilibrium price in $(L, H - \delta)$, and since $p \in (L, H - \delta)$ is an equilibrium price, there is no equilibrium price in (L, p); that is $p'' \notin (L, p)$. Combining $p'' \in [L, p)$, we have $p'' = L$, as desired.

In the rest of the proof, we revert to the original notation and write $\tilde{\pi}_1(p|\lambda)$ as $\tilde{\pi}_1(p)$; that is, we omit its dependence on λ.

Step 5. For any $\delta > 0$, let $p \in (L, H - \delta)$. If $\lambda > 0$ is sufficiently small then the PBE $\mathcal{E}_\lambda(p)$, which is constructed in Lemma 13, is Pareto undominated by any mixed-learning PBE with an equilibrium price p.

Proof. Let $\delta > 0$ be sufficiently small, and take any $\lambda < \lambda_{p,\delta}$. By Step 3 and Step 4, any mixed-learning PBE has a unique equilibrium price $p \in (L, H-\delta)$, and type L of S offers price L with a probability $y(p) > 0$. In such a mixed-learning PBE, type L's payoff is L and type H's payoff is $L + \tilde{q}(p)(p-L)$. To show that $\mathcal{E}_\lambda(p)$ is Pareto undominated, it suffices to show that among all such PBEs that S earns those profits, B's payoff is the highest in $\mathcal{E}_\lambda(p)$.

B's payoff in $\mathcal{E}_\lambda(p)$ is

$$U_B(p) = (1-\pi)y(p)(L-p) + \pi(H-p), \tag{21}$$

where $y(p)$ is the probability that type L of S charges price p.

Next, consider another mixed-learning PBE with an equilibrium price p, denoted $\tilde{\mathcal{E}}_\lambda(p)$, where $\tilde{\sigma}$ is S's equilibrium strategy. To ease notation, let $\tilde{y}(p) = \tilde{\sigma}(\{p\} \mid L)$ and $\tilde{x}(p) = \tilde{\sigma}(\{p\} \mid H)$.[22] For PBE $\tilde{\mathcal{E}}_\lambda(p)$, let $\tilde{P} = \operatorname{supp}(\bigcup_v \tilde{\sigma}(\cdot \mid v)) \cap (L, H)$. By Step 3, there is a single price $p \in \operatorname{supp}(\bigcup_v \tilde{\sigma}(\cdot \mid v)) \cap (L, H-\delta)$. Hence, $p' \geq H - \delta$ for all $p' \in \tilde{P} \setminus \{p\}$. B's payoff $\tilde{U}_B(p)$ in $\tilde{\mathcal{E}}_\lambda(p)$ is

$$\begin{aligned}
\tilde{U}_B(p) &= (1-\pi)\mathbb{E}_{\tilde{\sigma}(\cdot|L)}[L-p'] + \pi\mathbb{E}_{\tilde{\sigma}(\cdot|H)}[H-p'] \\
&= (1-\pi)(L-p)\tilde{y}(p) + (1-\pi)\mathbb{E}_{\tilde{\sigma}(\cdot|L)}[(L-p')\mathbf{1}_{\{p'\neq p\}}] \\
&\quad + \pi(H-p)\tilde{x}(p) + \pi\mathbb{E}_{\tilde{\sigma}(\cdot|H)}[(H-p')\mathbf{1}_{\{p'\neq p\}}] \\
&\leq (1-\pi)(L-p)\tilde{y}(p) + \pi(H-p)\tilde{x}(p) \\
&\quad + (1-\pi)(L-H+\delta)\mathbb{E}_{\tilde{\sigma}(\cdot|L)}[\mathbf{1}_{\{p'\neq p\}}] + \pi\delta\mathbb{E}_{\tilde{\sigma}(\cdot|H)}[\mathbf{1}_{\{p'\neq p\}}].
\end{aligned}$$

where the inequality is by $p' \geq H - \delta$ for all $p' \in \tilde{P} \setminus \{p\}$. Here, $\mathbf{1}$ is the indicator function. Since $L - H + \delta < 0$ and $\mathbb{E}_{\tilde{\sigma}(\cdot|H)}[\mathbf{1}_{\{p'\neq p\}}] = 1 - \tilde{x}(p)$, it follows that

$$\tilde{U}_B(p) < \pi\delta(1-\tilde{x}(p)) + (1-\pi)(L-p)\tilde{y}(p) + \pi(H-p)\tilde{x}(p).$$

We consider B's posterior after price p is offered. In both $\mathcal{E}_\lambda(p)$ and $\tilde{\mathcal{E}}_\lambda(p)$, B must assign to type H the same posterior probability $\tilde{\pi}_1(p)$ if price p is offered. Hence, $\tilde{y}(p) = y(p)\tilde{x}(p)$ by Bayes' rule. Using this, we have

$$\tilde{U}_B(p) < \pi\delta(1-\tilde{x}(p)) + [(1-\pi)(L-p)y(p) + \pi(H-p)]\tilde{x}(p).$$

---

[22] Step 3 shows that the set, $\operatorname{supp}(\bigcup_v \tilde{\sigma}(\cdot \mid v)) \cap (L, H-\delta)$, is a singleton. This leaves the possibility that S may offer (multiple) prices greater than or equal to $H - \delta$.

Now we compare B's payoffs $U_B(p)$ and $\tilde{U}_B(p)$:

$$\begin{aligned} U_B(p) - \tilde{U}_B(p) &> (1-\tilde{x}(p))[\pi(H-p-\delta) + (1-\pi)(L-p)y(p)] \\ &= (1-\tilde{x}(p))\pi\left(H-p-\delta-\frac{1-\tilde{\pi}_1(p)}{\tilde{\pi}_1(p)}(p-L)\right), \end{aligned} \tag{22}$$

where $(1-\pi)y(p) = \pi\frac{1-\tilde{\pi}_1(p)}{\tilde{\pi}_1(p)}$ by Bayes' rule. By Lemma 12, $\tilde{\pi}_1(p) \to 1$ as $\lambda \to 0$. For any sufficiently small δ, we have $H-p-\delta > \delta$. For each $p \in (L, H)$, let

$$\lambda_p^2 = \sup\left\{\lambda' \in (0, \lambda_{p,\delta}) : H-p-\delta-\frac{1-\tilde{\pi}_1(p)}{\tilde{\pi}_1(p)}(p-L) \geq 0 \quad \forall \lambda < \lambda'\right\}. \tag{23}$$

That is, λ_p^2 is the highest λ' in $(0, \lambda_{p,\delta})$ such that if $\lambda < \lambda'$, then $U_B(p) \geq \tilde{U}_B(p)$. Therefore, if $\lambda < \lambda_p^2$, then B's payoff in $\mathcal{E}_\lambda(p)$ is weakly higher than in any mixed-learning PBE with an equilibrium price p.

By the definition of λ_p^2, for any $\eta' \in (0, \frac{H-L}{2})$, $\inf\{\lambda_p^2 : p \in (L+\eta', H-\eta')\} > 0$.

Step 6. For any $\delta > 0$, let $p \in (L, H-\delta)$. If λ is sufficiently small then the PBE $\mathcal{E}_\lambda(p)$ is not Pareto dominated by any PBE $\mathcal{E}_\lambda(p')$ for any $p' \in (L, H)$. Recall B's payoff in the PBE $\mathcal{E}_\lambda(p)$ is given by (21), where $y(p) = (\frac{1}{\tilde{\pi}_1(p)} - 1)/(\frac{1}{\pi} - 1)$ by Bayes' rule.

Proof. First, we show that there is some $\epsilon' > 0$ such that $\mathcal{E}_\lambda(p)$ is not Pareto dominated by $\mathcal{E}_\lambda(p')$ for any $p' \in (p, p+\epsilon')$. It suffices to show that that $U_B(p) > U_B(p')$ for any $p' \in (p, p+\epsilon')$. We show this by showing that $U_B'(p) < -\pi/2$ if λ is small enough. Using the expression of $y(p)$ and taking the derivative of both sides of (21), we have

$$U_B'(p) = \pi(p-L)\frac{\tilde{\pi}_1'(p)}{(\tilde{\pi}_1(p))^2} - (1-\pi)y(p) - \pi.$$

By Lemma 12, as $\lambda \to 0$, $\tilde{\pi}_1'(p) \to 0$ and thus $y(p) \to 0$. Hence, $U_B'(p) \to -\pi$. For some $\lambda_p^3 > 0$, we have $U_B'(p) < -\pi/2$ for any $\lambda < \lambda_p^3$. Let $\epsilon' > 0$ be such that $U_B'(p') < 0$ for all $p' \in (p, p+\epsilon')$. Then, $U_B(p) > U_B(p')$ for any $p' \in (p, p+\epsilon')$.

Second, we show that $\mathcal{E}_\lambda(p)$ is not Pareto dominated by $\mathcal{E}_\lambda(p')$ for any $p' \in [p+\epsilon', H)$. By (21), we have $U_B(p) < \pi(H-p)$. As $\lambda \to 0$, we have $U_B(p) \to \pi(H-p)$. Thus, if λ is sufficiently small then for any $p' \geq p+\epsilon'$, we have $U_B(p) > \pi(H-p-\epsilon') \geq \pi(H-p') > U_B(p')$. Thus, $\mathcal{E}_\lambda(p)$ is not Pareto dominated by $\mathcal{E}_\lambda(p')$.

Third, we show that $\mathcal{E}_\lambda(p)$ is not Pareto dominated by $\mathcal{E}_\lambda(p')$ for any $p' \in (L, p)$. Let $\gamma > 0$ be small enough that $p < H-\gamma$. By the proof in Step 3, type H's payoff in $\mathcal{E}_\lambda(p)$ is given by (20) (when type L's payoff equal $u_L = L$), and is strictly increasing on $(L, H-\gamma)$ if $\lambda < \lambda_\gamma$. Since $p < H-\gamma$, for any $p' < p$, type H's payoff in $\mathcal{E}_\lambda(p')$ is strictly less than in

$\mathcal{E}_\lambda(p)$, and thus $\mathcal{E}_\lambda(p)$ is not Pareto dominated by $\mathcal{E}_\lambda(p')$.

Step 7. For any $p \in (L, H)$, if λ is sufficiently small then $\mathcal{E}_\lambda(p)$ is not Pareto dominated by any mixed-learning equilibrium.

Proof. Let ϵ be small enough that $p - L > 2\epsilon$ and $H - p > 2\epsilon$. We divide the set of all mixed-learning PBEs into three sets: Γ^0, Γ^+, and Γ^-, which are the set of PBEs such that the infimum price that type H of S offers is in $[L+\epsilon, H-\epsilon]$, $(H-\epsilon, H]$, and $[L, L+\epsilon)$, respectively.

First, we show that $\mathcal{E}_\lambda(p)$ is not Pareto dominated by any PBE in Γ^0. Let

$$\lambda_p^6 = \inf\Big\{\lambda_p^2 : p \in [L+\epsilon, H-\epsilon]\Big\}.$$

where λ_p^2 is defined in (23). As shown at the end of Step 5, $\lambda_p^6 > 0$. By definition, if $\lambda < \lambda_p^6$ then for any PBE in Γ^0 with an equilibrium price $p' \in [L+\epsilon, H-\epsilon]$, $\mathcal{E}_\lambda(p')$ is not Pareto dominated by any PBE with an equilibrium price p'. Moreover, if $\lambda < \lambda_p^5$, then by Step 6, $\mathcal{E}_\lambda(p)$ is not Pareto dominated by $\mathcal{E}_\lambda(p')$. Therefore, if $\lambda < \lambda_p^5$ and $\lambda < \lambda_p^6$, then $\mathcal{E}_\lambda(p)$ is not Pareto dominated by any PBE in Γ^0.

Second, we show that $\mathcal{E}_\lambda(p)$ is not Pareto dominated by any PBE in Γ^-. Recall that type H's payoff in $\mathcal{E}_\lambda(p)$ is $\tilde{U}_H(p)$, as defined in (20), which converges to p as $\lambda \to 0$ (because $\tilde{q}(p) \to 1$). Since $p > L + \epsilon$, there is a $\lambda_p^7 > 0$ such that if $\lambda < \lambda_p^7$, then $\tilde{U}_H(p) > L + \epsilon$. For any PBE in Γ^-, since type H of S offers a price in $(L, L+\epsilon)$, his payoff is at most $L + \epsilon$, which is strictly lower than his payoff in $\mathcal{E}_\lambda(p)$, $\tilde{U}_H(p)$. Thus, if $\lambda < \lambda_p^7$, then $\mathcal{E}_\lambda(p)$ is not Pareto dominated by any PBE in Γ^-.

Finally, we show that $\mathcal{E}_\lambda(p)$ is not Pareto dominated by any PBE in Γ^+. For any PBE in Γ^+, the prices that type H of S may offer are above $H - \epsilon$. Thus, B's payoff is at most $\pi(H - (H-\epsilon)) = \pi\epsilon$. In $\mathcal{E}_\lambda(p)$, B's payoff $U_B(p)$, as defined in (21), converges to $\pi(H-p)$ as $\lambda \to 0$. Since $p < H - \epsilon$, there is a $\lambda_p^8 > 0$ such that if $\lambda < \lambda_p^8$, then $U_B(p) > \pi\epsilon$. That is, if $\lambda < \lambda_p^8$ then B's payoff in $\mathcal{E}_\lambda(p)$ is higher than in any PBE in Γ^+. Thus, $\mathcal{E}_\lambda(p)$ is not Pareto dominated by any PBE in Γ^+. ■

Step 3 We examine B's (ex-ante) expected payoff in the PBE $\mathcal{E}_\lambda(p^*)$, which is Pareto undominated (Lemma 14). Then, we only need to show that for any $u_B \in (0, \pi(H-L))$, there exists some price p_λ such that B's payoff in the PBE $\mathcal{E}_\lambda(p_\lambda)$ converges to u_B as $\lambda \to 0$.

Fix any $\lambda > 0$ and take any $p^* \in (\underline{p}_\lambda, H)$, where $\underline{p}_\lambda$ is defined in Lemma 13. Consider B's ex-ante payoff in $\mathcal{E}_\lambda(p^*)$. Recall that if price p^* is offered then B randomizes between buying without acquiring information and acquiring information with accuracy $\tilde{q}(p^*)$. This means that B's payoff is the same as the payoff that she obtains from buying without acquiring information. Hence, B's equilibrium payoff is

$$U_B(p^*) = \pi(H - p^*) + (1 - \pi)y^*(L - p^*),$$

where $y^* = \frac{1-\tilde{\pi}_1(p^*)}{\tilde{\pi}_1(p^*)}\frac{\pi}{1-\pi}$ by Bayes' rule.

Let $\lambda \to 0$. Then $y^* \to 0$ since $\tilde{\pi}_1(p^*) \to 1$ by Lemma 12. Hence, $U_B(p^*) \to \pi(H - p^*)$. Since $\underline{p}_\lambda \to L$ as $\lambda \to 0$ by Lemma 13, it follows that for any $p^* \in (L, H)$, there exists a small $\lambda > 0$ such that $p^* > \underline{p}_\lambda$. In particular, let $p^* = H - u_B/\pi \in (L, H)$. Then, B's payoff in $\mathcal{E}_\lambda(p_\lambda)$ converges to u_B. This completes the proof of Proposition 4 in the case of $\lim_{q\to 1} c'(q) = \infty$.

C.2 Case of $\lim_{q\to 1} c'(q) < \infty$

We prove Proposition 4 under the assumption that $c'(1) \equiv \lim_{q\to 1} c'(q) < \infty$.

In the proof of Lemma 11, B's first order condition with respect to q is replaced with

$$q = \begin{cases} 1 & \text{if } \pi_1(p)(H-p) \geq \lambda c'(1) \\ (c')^{-1}\left(\frac{\pi_1(p)(H-p)}{\lambda}\right) & \text{if } \pi_1(p)(H-p) < \lambda c'(1). \end{cases} \tag{24}$$

If there is no $\tilde{q}(p) \leq 1$ that satisfies (13), that is, if p is such that

$$\lambda c'(1)\left(\frac{H-L}{H-p} - 1\right) + \lambda c(1) < p - L, \tag{25}$$

then let $\tilde{q}(p) = 1$ and $\tilde{\pi}_1(p) = 1 - \frac{\lambda c(1)}{p-L}$. This way, both the first-order condition (24) and B's indifference condition (between no information and accuracy $\tilde{q}(p)$):

$$\tilde{\pi}_1(p)(1 - \tilde{q}(p))(H - p) + (1 - \tilde{\pi}_1(p))(L - p) + \lambda c(\tilde{q}(p)) = 0,$$

which is an analog of (15), are satisfied. Moreover, as $\lambda \to 0$, we have $\tilde{q}(p) \to 1$ and $\tilde{\pi}_1(p) \to 1$ in this case.

Lastly, we modify our proofs of Lemma 13 and Proposition 4 to accommodate the present case of $c'(1) < \infty$. If λ is such that there exists no $p \in (L, H)$ satisfying (25) then our proof for Lemma 13 and Proposition 4 is valid without any modification. If λ is such that there exists a $p \in (L, H)$ satisfying (25), then, multiplying $H - p$ on both sides of (25), we have

$$\lambda c'(1)(p - L) + \lambda c(1)(H - p) < (p - L)(H - p). \tag{26}$$

Since there is a $p \in (L, H)$ satisfying (26), there exists an interval $(p_1^\lambda, p_2^\lambda)$ such that (26), or equivalently (25) holds if and only if $p \in (p_1^\lambda, p_2^\lambda)$. For all $p \in (p_1^\lambda, p_2^\lambda)$, we set $\tilde{q}(p) = 1$ and $\tilde{\pi}_1(p) = 1 - \frac{\lambda c(1)}{p-L}$, and Lemma 13 and Proposition 4 hold.