Bargaining and Timing of Information Acquisition*

Kalyan Chatterjee† Miaomiao Dong‡ Tetsuya Hoshino§

May 5, 2022

Abstract

We consider an ultimatum game in which the value of the object being sold to the buyer can be either high or low. The seller knows what the value is but the buyer does not. The value of the object to the seller is zero. We introduce the option for the buyer to acquire information before or after the offer, at a low cost. This information either reveals the value is high or provides no information. As the cost of information vanishes, in all Pareto-undominated equilibria, the buyer gets all the surplus although the option is never used.

*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

We investigate a simple bargaining problem between a seller and a buyer over a single indivisible object. The object's quality is either high or low, drawn according to a common prior. The seller is informed about the quality, but the buyer is initially uninformed. The value of the object to the seller is zero regardless of the quality, while the value to the buyer is equal to the quality. The seller makes a take-it-or-leave-it price offer, and then the buyer either accepts or rejects it.

The new feature of our model is that the buyer has the option to acquire costly information about quality both before and after observing the seller's offer—before accepting or rejecting the offer. The buyer chooses the accuracy of information at each point in time at the cost of that accuracy. At each time of collecting information, the buyer has two possible signal realizations: She is informed of the quality being high or is uninformed. We call this information structure H-focused.[1,2] Neither the choice of accuracy nor the realized signal at each point in time is observed by the seller.

An example we have in mind is selling a used car. The seller knows the quality of his car. If he does not sell the car, he discards it and gains payoff zero regardless of the quality. The buyer may check the car by hiring a mechanic before the seller makes an offer, while the seller cannot observe the nature of the evaluation or its result. The buyer may also test-drive or consult another mechanic to evaluate the quality after observing the offer.

We are interested in the limit equilibrium as the buyer's cost of information vanishes. To put our result in perspective, we consider the following two benchmark cases. In the full-information benchmark in which the object's quality is common knowledge, the buyer must get zero surplus in the unique equilibrium. In the null-information benchmark in which the buyer does not have the option to acquire information, there is a continuum of equilibria, in which the buyer's surplus ranges from zero to the maximum possible.

Our main result—formally stated as Theorem 1—is as follows:[3]

> In *every* Pareto-undominated equilibrium in the limit as the unit 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.

Here, an equilibrium is Pareto-undominated if it is not Pareto-dominated by any other equilibrium. Our result says that the buyer's having cheap and instant access to information—before and after observing the offer—gives her the full surplus. Note that neither the seller's sole power to make a (take-it-or-leave-it) offer nor the existence of his private information about the object's quality gives him any of the surplus.

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

[2] The accuracy of an H-focused signal refers to the probability that the buyer learns a high signal conditional on the object's quality being high.

[3] The same result would be obtained if we restrict ourselves to pure-strategy equilibria.

The main takeaway is that having cheap and instant access to information, while knowing nothing initially, may be significantly better than knowing everything from the start or knowing nothing to the end. Recent information technology has given consumers cheap and instant access to information at any time—before and after being offered the terms of trade. The information accessibility will keep a price low enough that a buyer, in the end, will not have to acquire any information. There are existing studies that look at similar questions but in different models. Their results, not surprisingly, are also different, as discussed below.

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 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 et al. (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 not only before the seller quotes a price but also after that. As a result, as the cost vanishes, the buyer gains the full surplus in our model under the Pareto-undominated or pure-strategy criterion of equilibrium selection.

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.[4] 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 (Table 1). We assume that the buyer can choose the timing of information acquisition. While the existing studies restrict the buyer's timing of information acquisition to either before or after observing the seller's offer, we allow her to acquire information anytime she wants. Another difference is information that the buyer can access. While the existing studies, in the framework of information design or rational inattention, assume that the buyer can flexibly choose any kind of signal, we assume that she chooses an H-focused signal. The cost of flexible information may depend on the buyer's prior belief (e.g., the entropic cost), but the cost of H-focused signals does not.

Our model is not a lemons model because the seller's payoff in the event of no trade does not

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

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

	RS	RRS	R	Ours
Does the seller know the object's quality (its value to the buyer)?	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	both
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[5]	imperfect	perfect	never

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

Table 1: bargaining with information acquisition

depend on the object's quality. As in the original lemons model, similar results would obtain if the object's value to the buyer were to be sufficiently higher than the seller's no agreement payoff. The H-focused information structure that we use has some similarities with that Feinberg and Skrzypacz (2005) use. In their paper, the seller is initially uninformed about the buyer's valuation of the object, which is either high or low. The seller may be (exogenously) informed whether the object is of high value for the buyer or may remain uninformed. Feinberg and Skrzypacz show that the buyer's (second-order) uncertainty about whether the seller is informed of the high valuation or uninformed changes the outcomes. In our model, the buyer who is initially uninformed of the object's quality has the option to acquire information, but, *in equilibrium does not actually acquire any;* thus, she remains uninformed throughout the game and this is known to the seller. This result is due to the buyer exercising her option to acquire information if the proposal is too high, a realistic feature of the result.

Beyond bargaining problems, information acquisition has 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.[6] Information acquisition by agents has been analyzed in other applications, such as auction theory.[7] A common feature of these studies is that a buyer acquires information at exogenously given times, but we depart from the literature by allowing the buyer to endogenously choose when to acquire information.

Our study is also related to early studies on the Coase conjecture (e.g., Fudenberg et al., 1985; Gul et al., 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

[6] Crémer and Khalil (1994) and Crémer et al. (1998) study similar models such that the agent may gather information before the principal offers a contract.

[7] In auction theory, there are studies that discuss bidders' incentives to obtain costly information about their evaluations (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.

gets no surplus even though only the seller can make offers in both models.

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 We consider the following four-stage model:

Time 0: Nature draws the object's quality $\mathbf{v}$ from the set $V \equiv \{H, L\}$, with $H > L > 0$, according to common prior probabilities $\mathbb{P}(\mathbf{v} = H) = \pi \in (0, 1)$ and $\mathbb{P}(\mathbf{v} = L) = 1 - \pi$. By abuse of notation, we denote the prior by the same notation π. S observes a realized quality $v \in V$ (where S is called type v), but B does not. The object's value to S equals zero (regardless of v), while the value to B equals v.

Time 1: B acquires a **pre-offer signal** $\mathbf{x}_1$, which is H-focused. That is, the value of the signal is either H (which conclusively reveals that $\mathbf{v} = H$) or N (which we interpret it as B's observing nothing).[8] She chooses accuracy $q_1 \in [0, 1]$, which yields the following conditional distributions:

$$\begin{aligned} \mathbb{P}(\mathbf{x}_1 = H \mid \mathbf{v} = H) &= q_1, & \mathbb{P}(\mathbf{x}_1 = H \mid \mathbf{v} = L) &= 0 \\ \mathbb{P}(\mathbf{x}_1 = N \mid \mathbf{v} = H) &= 1 - q_1, & \mathbb{P}(\mathbf{x}_1 = N \mid \mathbf{v} = L) &= 1. \end{aligned}$$

Simultaneously, S charges a price p conditional on her type v. Then, B observes a realized signal x_1 as well as the price p.

Time 2: B acquires a **post-offer signal** $\mathbf{x}_2$, which is also H-focused. If she chooses accuracy $q_2 \in [0, 1]$, it yields the following conditional distributions:

$$\begin{aligned} \mathbb{P}(\mathbf{x}_2 = H \mid \mathbf{v} = H) &= q_2, & \mathbb{P}(\mathbf{x}_2 = H \mid \mathbf{v} = L) &= 0, \\ \mathbb{P}(\mathbf{x}_2 = N \mid \mathbf{v} = H) &= 1 - q_2, & \mathbb{P}(\mathbf{x}_2 = N \mid \mathbf{v} = L) &= 1. \end{aligned}$$

Then, B observes a realized signal x_2.

Time 3: B decides whether to buy the object or not at the price p. Payoffs are then realized, and the game ends.

[8] Feinberg and Skrzypacz (2005) also consider the H-focused signals. In their model, a buyer knows her own valuation of the object but a seller does not; the seller is either informed of the buyer's valuation being high or uninformed. Unlike our model, their model assumes that the seller's information is exogenous.

Information Cost Pre- and post-offer signals are costly, and the cost depends on accuracy. Let $c : [0,1] \to \mathbb{R}_+$ be a twice continuously differentiable function such that $c'(q) > 0$ and $c''(q) > 0$ for each $q \in (0,1)$. Assume that $c(0) = 0$ and $c'(0) \equiv \lim_{q \downarrow 0} c'(q) = 0$. Given a unit cost parameter $\lambda > 0$, we assume that the costs of pre- and post-offer signals with accuracy q_1 and q_2, are $\lambda c(q_1)$ and $\lambda c(q_2)$, respectively.

Payoffs For a quality v and a price p and for B's accuracy q_1 and q_2, if B buys then S receives payoff p and B receives payoff $v - p - \lambda c(q_1) - \lambda c(q_2)$; otherwise, S receives payoff 0 and B receives payoff $-\lambda c(q_1) - \lambda c(q_2)$.

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_t at time $t = 1, 2$ and makes a purchase decision at time 3. At time 1, when acquiring a pre-offer signal, B has a (trivial) history $h_1 = \varnothing$. At time 2, when acquiring a post-offer signal, B knows the charged price p, the chosen accuracy q_1, and the realized signal x_1, thereby having a history $h_2 = (p, q_1, x_1)$. Let $H_1 = \{\varnothing\}$ be the singleton of the time-1 history, and let $H_2 = \mathbb{R}_+ \times [0,1] \times \{H, N\}$ be the set of time-2 histories. At time 3, when making a purchase decision, B has a history $h_3 = (h_2, q_2, x_2)$ with accuracy q_2 and a realized signal x_2. Let $H_3 = H_2 \times [0,1] \times \{H, N\}$ be the set of time-3 histories. B's strategy is a function β, where (i) $\beta(\cdot \mid h_1) \in \Delta([0,1])$ is the distribution of pre-offer accuracy at time-1 history h_1, (ii) $\beta(\cdot \mid h_2) \in \Delta([0,1])$ is the distribution of post-offer accuracy at time-2 history h_2, and (iii) $\beta(h_3) \in [0,1]$ is the probability of buying at time-3 history h_3.

Equilibrium Our equilibrium concept is perfect Bayesian equilibrium (Fudenberg and Tirole, 1991). 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 quality $\mathbf{v}$). Since the type space V is binary, let μ denote the probability of high quality $\mathbf{v} = H$.

Definition 1. An assessment $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$ is an **equilibrium** if it satisfies all of 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. $\mu^*(p, q_1) = \mu^*(p, q_1')$ for all p, q_1, q_1', and $\mu^*(h_2, q_2) = \mu^*(h_2, q_2')$ for all h_2, q_2, q_2'.

The third condition—often called "no signaling what you don't know"—requires that B's own deviation does not change her belief.

In the following sections, we will focus on the set of **Pareto-undominated** equilibria, in which each equilibrium payoff vector is not Pareto-dominated by the payoff vector of any other equilibrium.

Benchmarks We refer to the benchmark cases described in Section 1. In the full-information benchmark, where the object's quality is common knowledge, there exists a unique equilibrium in which S of type H offers price H, S of type L offers price L, and B accepts both offers. In this equilibrium, B gets zero surplus. In the null-information benchmark, where B does not have the option to acquire information, there is a continuum of equilibria, in which S of both types 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 zero to the maximum $\pi(H - L)$.

3 Main Result

We are interested in the case where B can acquire information cheaply before and after S makes an offer. Our main result states that in the limit as unit cost λ tends to zero, S offers the lowest price L and B acquires neither pre- nor post- information on any equilibrium path; hence, S has the smallest surplus L, and B has the greatest surplus $\pi(H - L)$.[9]

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

1. *S of each type $v \in V$ offers a price $p^* \in [L, L + \epsilon)$ with probability 1 (i.e., $\sigma^*(\{p^*\} \mid v) = 1$).*
2. *On path, B acquires no pre- or post-offer information and buys at price p^* with probability 1.*

Compared with the benchmarks, B's option of acquiring cheap and instant information deters S from offering a high price, although she never exercises the option.

3.1 Proof Sketch

Step 1 We show that on any equilibrium path, B does not acquire any pre- or post-offer information. That is, $\beta^*(\{0\} \mid h_1) = \beta^*(\{0\} \mid h_2) = 1$ for all on-path h_1, h_2.

To begin with, we note that B does not acquire any post-offer information after observing either price L or H. Since the object's value $\mathbf{v}$ is between L and H, B must get a non-negative (gross) payoff at price L and a non-positive (gross) payoff at price H, and therefore B need not know the exact value (at a cost) in either case. Hence, we can assume that price $p \in (L, H)$ is offered.

We show that with information acquisition, as the unit cost λ vanishes, the limit equilibrium outcome is unique and B gains full surplus. Our proof consists of four substeps.

1. At any on-path history $h_2 = (p, q_1, x_1)$ with a price $p \in (L, H)$ and a signal realization $x_1 \in \{H, N\}$, B acquires no post-offer information with positive probability: $\beta^*(\{0\} \mid h_2) > 0$.

   This claim is trivial if B is certain about S's type $\mathbf{v}$ (when it is revealed by a price p or a pre-offer signal realization H). Hence, we assume that B is uncertain about S's type $\mathbf{v}$ at history $h_2 = (p, q_1, N)$. Suppose, on the contrary, that at this history h_2, B acquires some post-offer information for sure: $\beta^*(\{0\} \mid h_2) = 0$. Then, signal $\mathbf{x}_2$ must be relevant for her purchasing decision. That is, she buys if $\mathbf{x}_2 = H$ and does not if $\mathbf{x}_2 = N$. This means that

[9] The same result obtains for every pure-strategy equilibrium.

S of type L will make no sale because he must generate signal $\mathbf{x}_2 = N$ for sure, but then, he has a profitable deviation to charging price L (to secure a sale and profit L). This is a contradiction.

2. B never acquires any pre-offer information : $\beta^*(\{0\} \mid h_1) = 1$.

   As shown above, B weakly prefers $q_2 = 0$ at any on-path history h_2, and thus her equilibrium payoff is equal to her payoff when she chooses $q_2 = 0$ (with probability 1) and then takes an optimal action at the subsequent history $h_3 = (h_2, 0, N)$. To evaluate this payoff, suppose that she takes $q_2 = 0$ at history h_2. Then, she is not indifferent between buying and not buying at history h_3.[10] This lack of indifference implies that B buys with probability 1, because otherwise, S should have cut the price to secure a sale and profit. B's equilibrium payoff is, therefore, equal to the payoff when she takes $q_2 = 0$ at history h_2 and buys with probability 1 at history $h_3 = (h_2, 0, N)$.

   We then find that regardless of her pre-offer accuracy q_1 and the realized signal x_1, B buys with probability 1 in the end. That is, pre-offer signal $\mathbf{x}_1$ is useless (because it does not affect her purchase decision) and thus B optimally chooses pre-offer accuracy $q_1 = 0$ with probability 1.

3. At any on-path history $h_2 = (p, q_1, x_1)$ with a price $p \in (L, H)$ and a signal $x_1 \in \{H, N\}$, B acquires no post-offer information with probability 1: $\beta^*(\{0\} \mid h_2) = 1$.

   Suppose, by negation, that in some equilibrium, there is an on-path history h_2 such that B might acquire post-offer information: $\beta^*(\{0\} \mid h_2) < 1$. Recall that B acquires no pre-offer information, but if she "splits" information acquisition by allocating positive accuracy to pre-offer acquisition, then while keeping a gross expected payoff, she can reduce the cost of information because the cost function c is strictly convex. This is a profitable deviation for B, which establishes the claim.

4. At any on-path history h_3, with pre- and post-offer accuracy $q_1 = q_2 = 0$ and realized signals $x_1 = x_2 = N$, B buys with probability 1 at any price $p \neq H$.

   At the on-path history h_3 with price $p \neq H$, B is not indifferent between buying and not buying (because otherwise, B should have acquired a possibly small amount of information). Then, she buys with either probability 1 or 0. Hence, she must buy with probability 1 because otherwise, S would have not offered price p.

Step 2 We show that every Pareto-undominated equilibrium is a pure-strategy pooling equilibrium with some price $p^* \in [L, H)$.

[10] To see this, suppose that B is indifferent between buying and not buying. Then, she should have strictly preferred negligible but positive post-offer accuracy. The marginal cost $c'(0)$ at accuracy $q_2 = 0$ is zero, but any small amount of information breaks the indifference and allows B to make a better purchase decision and bring a positive payoff.

1. In any equilibrium, there is only one equilibrium price that is less than H.[11]

   Suppose that there are at least two equilibrium prices, denoted $p' < p'' < H$. Recall that B never acquires any information. As argued in Step 1, she then buys (with probability 1) at both prices, but this means that S earns more profit from the higher price p'', which contradicts the assumption that S could offer the lower price p' in the equilibrium.

2. In every Pareto-undominated equilibrium, price H is not charged.

   Take any equilibrium with price H in the support of S's strategy σ^*. As shown above, there is a (unique) price $p^* \neq H$ such that $\{p^*, H\} = \bigcup_v \operatorname{supp}(\sigma^*(\cdot \mid v))$. Price H must come from S of type H. Recall, from Step 1, that B never acquires any information. Now suppose that B buys with probability y at price H. For S of type L not to offer price H, we must have $y \leq p^*/H$. Moreover, if $y < p^*/H$ then S of type H would not offer price H, which contradicts the assumption that price H is in the support of $\sigma^*(\cdot \mid H)$. Hence, $y = p^*/H$, which implies that S of each type gains profit p^*. However, this equilibrium is Pareto-dominated by the pooling equilibrium in which S of each type offers price p^* and B buys with probability 1 (without information acquisition).

Step 3 For any (pure-strategy) pooling equilibrium derived in Step 1, we show that the equilibrium price p^* tends to L as a unit cost λ vanishes. From Step 1, we recall that B never acquires any information. That is, the equilibrium price must be such that B has no incentive to acquire information. If a (pooling) price p were sufficiently larger than L, then B would have a strict incentive to acquire post-offer information because otherwise, she might end up buying the object worth payoff $\mathbf{v} = L$ at the "high" price $p \gg L$. Hence, the equilibrium price must be close enough to L. In the proof, we give an upper bound of the equilibrium price so that B does not acquire information and then show that the upper bound tends to L as λ vanishes.

4 Conclusion

We have considered an ultimatum bargaining problem with a seller informed about the quality of his good (as the proposer) and a buyer who is initially uninformed about the quality (as the responder). We introduce an option for the buyer to acquire information both before and after the proposal at a low cost, and use a refinement among perfect Bayesian equilibria, selecting only a Pareto-undominated one. Compared to the benchmark cases, we obtain striking conclusions. If the buyer is exogenously supplied the information about the object's quality and this is commonly known (the full information case), the seller uses his proposal power to obtain full surplus from the bargaining. If the buyer has no opportunity to obtain information (the null information case), there is a continuum of equilibria, which are all Pareto-undominated. In these equilibria, the buyer obtains different shares of the surplus depending on off-the-equilibrium path beliefs. However, in

[11] There may exist an equilibrium in which price H is in the support of S's strategy.

our case, with cheap and instant access to information about the object's quality, in all Pareto-undominated or pure-strategy equilibria, the *buyer obtains the full surplus* although she never acquires any information.

A Proof of Theorem 1

The proof consists of four steps. Section A.1 shows that B acquires no post-offer information in any equilibrium. Section A.2 shows that every equilibrium in which S never offers price H is a pure-strategy pooling equilibrium. Section A.3 shows that B extracts (almost) all surplus in the pooling equilibrium. Section A.4 shows that every equilibrium in which S offers price H with positive probability is a Pareto-dominated, mixed-strategy equilibrium.

A.1 B Acquires No Post-Offer Information

Fix an equilibrium $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$ otherwise stated. Note that B will not acquire any post-offer information when she observes $\mathbf{x}_1 = H$, which reveals S's type $\mathbf{v} = H$. Hence, B might acquire post-offer information only when she observes $\mathbf{x}_1 = N$. Given any $q_1 \in [0,1]$, we then define two sets of equilibrium prices:

- $P^{\mathcal{E}}_{\text{acq}}(q_1)$ denotes the set of equilibrium prices after which B acquires post-offer information with non-zero probability:

$$P^{\mathcal{E}}_{\text{acq}}(q_1) \equiv \left\{ p \in \bigcup_{v \in V} \operatorname{supp}(\sigma^*(\cdot \mid v)) : \operatorname{supp}(\beta^*(\cdot \mid h_2)) \neq \{0\} \text{ at history } h_2 = (p, q_1, N) \right\}.$$

- $P^{\mathcal{E}}_{\text{not}}(q_1)$ denotes the set of equilibrium prices after which B acquires post-offer information with zero probability:

$$P^{\mathcal{E}}_{\text{not}}(q_1) \equiv \left\{ p \in \bigcup_{v \in V} \operatorname{supp}(\sigma^*(\cdot \mid v)) : \operatorname{supp}(\beta^*(\cdot \mid h_2)) = \{0\} \text{ at history } h_2 = (p, q_1, N) \right\}.$$

For each $q_1^* \in \operatorname{supp}(\beta^*(h_1))$, every equilibrium price is in either set. That is,

$$\bigcup_{v \in V} \operatorname{supp}(\sigma^*(\cdot \mid v)) \subseteq P^{\mathcal{E}}_{\text{acq}}(q_1^*) \cup P^{\mathcal{E}}_{\text{not}}(q_1^*).$$

B does not acquire any post-offer information after observing either price L or H. Since the object's value $\mathbf{v}$ is between L and H, B must get a non-negative (gross) payoff at price L and a non-positive (gross) payoff at price H, and therefore B need not know the exact value (at a cost) in either case. This is formalized below.

Lemma 1. *It holds that $L, H \notin P^{\mathcal{E}}_{\text{acq}}(q_1^*)$ for any $q_1^* \in \operatorname{supp}(\beta^*(\cdot \mid h_1))$.*

Next, we decompose $P^{\mathcal{E}}_{\text{acq}}(q_1)$ into two subsets according to whether B acquires any post-offer information at all. Because B will acquire no post-offer information at signal realization $\mathbf{x}_1 = H$ since it reveals S's type $\mathbf{v} = H$, B might acquire post-offer information only at signal realization $\mathbf{x}_1 = N$.

- $P^{\mathcal{E}}_a(q_1)$ denotes the set of equilibrium prices after which B acquires post-offer information with probability 1:

$$P^{\mathcal{E}}_a(q_1) \equiv \Big\{p \in P^{\mathcal{E}}_{\text{acq}}(q_1) : 0 \notin \operatorname{supp}(\beta^*(\cdot \mid h_2)) \text{ at history } h_2 = (p, q_1, N)\Big\}.$$

- $P^{\mathcal{E}}_0(q_1)$ denotes the set of equilibrium prices after which B mixes between acquiring and not acquiring information:

$$P^{\mathcal{E}}_0(q_1) \equiv \Big\{p \in P^{\mathcal{E}}_{\text{acq}}(q_1) : 0 \in \operatorname{supp}(\beta^*(\cdot \mid h_2)) \text{ at history } h_2 = (p, q_1, N)\Big\}.$$

By definition, we have for each q_1,

$$P^{\mathcal{E}}_{\text{acq}}(q_1) = P^{\mathcal{E}}_a(q_1) \cup P^{\mathcal{E}}_0(q_1).$$

Our goal is to show that B will not acquire any post-offer information—that is, $P^{\mathcal{E}}_{\text{acq}}(q_1)$ is empty for any on-path q_1^*.

Lemma 2. *The set $P^{\mathcal{E}}_a(q_1^*)$ is empty for any $q_1^* \in \operatorname{supp}(\beta^*(\cdot \mid h_1))$.*

Proof. Suppose, by negation, that $P^{\mathcal{E}}_a(q_1^*)$ is non-empty. Fix any $p^* \in P^{\mathcal{E}}_a(q_1^*)$. By Lemma 1, $p^* \in (L, H)$. At history $h_2 = (p^*, q_1^*, N)$, any $q_2 \in \operatorname{supp}(\beta^*(\cdot \mid h_2))$ must be such that $q_2 > 0$. Then, B buys if $\mathbf{x}_2 = H$ but does not if $\mathbf{x}_2 = N$.[12] This means that S of type L makes zero profit because conditional on type L, B observes $\mathbf{x}_2 = N$ with probability 1 and thus does not buy. S of type L has a profitable deviation of offering price L.[13] ■

To show that $P^{\mathcal{E}}_0(q_1^*)$ is empty, we introduce a convenient notation. Given any belief system μ, let $b_\mu(h_2, q_2)$ be B's optimal payoff when she chooses accuracy q_2 at history $h_2 = (p, q_1, x_1)$, which does not include the (sunk) pre-offer cost $c(q_1)$. Formally,

$$b_\mu(h_2, q_2) \equiv \mu(h_2)q_2 \max\{H - p, 0\} + (1 - \mu(h_2)q_2) \max\{\mathbb{E}_\mu[\mathbf{v} \mid h_3 = (h_2, q_2, N)] - p, 0\} - \lambda c(q_2),$$

where $\mu(h_2)q_2$ and $1 - \mu(h_2)q_2$ are the probabilities of $\mathbf{x}_2 = H$ and $\mathbf{x}_2 = N$ respectively. Then, we say that the support of B's post-offer accuracy $\beta^*(\cdot \mid h_2)$ at history h_2 must be a subset of

[12] This is because if B's purchase decision were independent of signal $\mathbf{x}_2$, she would profitably deviate to choosing $q_2 = 0$, by which she can save the cost without changing her purchase decision.

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

$\text{argmax}_{q_2} b_{\mu^*}(h_2, q_2)$. That is,

$$\text{supp}(\beta^*(\cdot \mid h_2)) \subseteq \underset{q_2 \in [0,1]}{\text{argmax}}\, b_{\mu^*}(h_2, q_2).$$

Lemma 3. *Let $h_2 = (p, q_1, x_1)$ be a history with any $p \in [L, H)$, any $q_1 \in [0, 1)$, and any $x_1 \in \{H, N\}$. If $0 \in \text{argmax}_{q_2} b_{\mu^*}(h_2, q_2)$ then at history $h_3 = (h_2, 0, N)$, B buys with probability $\beta^*(h_3) = 1$. That is, $\max_{q_2} b_{\mu^*}(h_2, q_2)$ equals B's payoff when she chooses accuracy $q_2 = 0$ and buys with probability 1.*

Proof. Suppose, for a contradiction, that there exists history $h_2 = (p, q_1, x_1)$ such that $0 \in \text{argmax}_{q_2} b_{\mu^*}(h_2, q_2)$ but B buys with probability $\beta^*(h_3) < 1$ at history $h_3 = (h_2, 0, N)$. Since S does not offer any price at which B buys with probability 0, it follows that $\beta^*(h_3) \in (0, 1)$, which means that B is indifferent between buying and not buying at history h_3.

However, we show that $0 \notin \text{argmax}_{q_2} b_{\mu^*}(h_2, q_2)$. Since not buying yields payoff 0 and B is indifferent between buying and not buying at history h_3, it follows that $b_{\mu^*}(h_2, 0) = 0$. Now suppose that B chooses $q_2 > 0$ at history h_2. Then, signal $\mathbf{x}_2$ must be relevant for B's purchasing decision: If $\mathbf{x}_2 = H$ then she buys to get (gross) payoff $H - p$, while if $\mathbf{x}_2 = N$ then she does not buy to get (gross) payoff 0. Since $\mathbf{x}_2 = H$ and $\mathbf{x}_2 = N$ have probabilities $\mu^*(h_2)$ and $1 - \mu^*(h_2)$ respectively, it follows that $b_{\mu^*}(h_2, q_2) = \mu^*(h_2) q_2 (H - p) - \lambda c(q_2)$. Since $\mu^*(h_2) > 0$, it must be that $0 \notin \text{argmax}_{q_2} b_{\mu^*}(h_2, q_2)$.[14] It is a contradiction. ■

Lemma 4. *The set $P_0^{\mathcal{E}}(q_1^*)$ is empty for any $q_1^* \in \text{supp}(\beta^*(\cdot \mid h_1))$.*

Proof. Suppose, by negation, that $P_0^{\mathcal{E}}(q_1^*)$ is non-empty. Fix any $p^* \in P_0^{\mathcal{E}}(q_1^*)$. By Lemma 1, $p^* \in (L, H)$. By definition, $0 \in \text{argmax}_{q_2} b_{\mu^*}(h_2^*, q_2)$ at history $h_2^* = (p^*, q_1^*, N)$. By Lemma 3, B's payoff at history h_2^* is equal to the payoff when she chooses accuracy $q_2 = 0$ and buys with probability 1.

We introduce a convenient notation. When B chooses pre-offer accuracy q_1 and observes a pre-offer signal realization x_1 (but not an offered price), we denote B's belief about a price $\mathbf{p}$ by

$$\rho^*(\cdot \mid q_1, x_1) \equiv \mathbb{E}_{\mathbf{v} \sim \mu^*(q_1, x_1)}[\sigma^*(\cdot \mid \mathbf{v})],$$

where σ^* is S's equilibrium strategy and μ^* is the equilibrium a belief system.

First, we show that the optimal pre-offer accuracy is $q_1^* = 0$. Let $q_1 \in [0, 1]$ be B's choice of pre-offer accuracy. When B observes $\mathbf{x}_1 = H$ (with probability πq_1), she expects the gross payoff $H - \mathbb{E}_{\mathbf{p}^* \sim \sigma^*(\cdot|H)}[\mathbf{p}^*]$ because her belief is $\rho^*(\cdot \mid q_1, H) = \sigma^*(\cdot \mid H)$. When B observes $\mathbf{x}_1 = N$ (with probability $1 - \pi q_1$), she expects the gross payoff $\mathbb{E}_{\mathbf{p}^* \sim \rho^*(\cdot|q_1, N)}[b_{\mu^*}(\mathbf{h}_2^*, \beta^*(\mathbf{h}_2^*))]$ with (random) history $\mathbf{h}_2^* = (\mathbf{p}^*, q_1, N)$. Hence, her payoff is

$$\pi q_1 \Big(H - \mathbb{E}_{\mathbf{p}^* \sim \sigma^*(\cdot|H)}[\mathbf{p}^*] \Big) + (1 - \pi q_1) \mathbb{E}_{\mathbf{p}^* \sim \rho^*(\cdot|q_1, N)} \Big[b_{\mu^*}(\mathbf{h}_2^*, \beta^*(\mathbf{h}_2^*)) \Big] - \lambda c(q_1). \tag{1}$$

[14] The marginal benefit of choosing $q_2 = 0$ is $\mu^*(h_2)(H - p) - \lambda c'(0) > 0$ for $c'(0) = 0$. Thus, $0 \notin \text{argmax}_{q_2} b_{\mu^*}(h_2, q_2)$.

Since $p^* \in P_0^{\mathcal{E}}(q_1^*)$, it holds that at history $h_2^* = (p^*, q_1^*, N)$, B's payoff $b_{\mu^*}(h_2^*, q_2)$ is the payoff when she chooses $q_2 = 0$ and buys with probability 1. Since $b_{\mu^*}(\mathbf{h}_2^*, \beta^*(\mathbf{h}_2^*)) = \mathbb{E}_{\mathbf{v}\sim\mu^*(\mathbf{h}_2^*)}[\mathbf{v}] - \mathbf{p}^*$, it follows that (1) is equal to

$$\pi q_1\Big(H - \mathbb{E}_{\mathbf{p}^*\sim\sigma^*(\cdot|H)}[\mathbf{p}^*]\Big) + (1-\pi q_1)\mathbb{E}_{\mathbf{p}^*\sim\rho^*(\cdot|q_1,N)}\Big[\mathbb{E}_{\mathbf{v}\sim\mu^*(\mathbf{h}_2^*)}[\mathbf{v}] - \mathbf{p}^*\Big] - \lambda c(q_1). \tag{2}$$

Since the posteriors are a martingale split of the prior, it follows that

$$(2) = \mathbb{E}_{\mathbf{v}\sim\pi}[\mathbf{v}] - \mathbb{E}_{\mathbf{v}\sim\pi}[\mathbb{E}_{\mathbf{p}^*\sim\sigma^*(\cdot|\mathbf{v})}[\mathbf{p}^*]] - \lambda c(q_1).$$

B's accuracy q_1^* maximizes this payoff. Since the first two terms are independent of the choice of accuracy q_1, we have a unique maximizer $q_1^* = 0$.

Second, we derive a contradiction by finding B's profitable deviation. For a small $\eta > 0$, let β^η be a deviating strategy such that:

- At history h_1, $\beta^\eta(\{\eta\} \mid h_1) = 1$
- At history h_2, if $x_1 = H$ then $\beta^\eta(\{0\} \mid h_2) = 1$, and if $x_1 = N$ then $\beta^\eta(\cdot \mid h_2) = \beta^*(\cdot \mid p, 0, N)$.
- At history h_3, B buys with probability 1 if either $\mathbf{x}_1 = H$ or $\mathbf{x}_2 = H$ and never buys otherwise.

B's gross payoff is weakly greater under this strategy β^η than under her equilibrium strategy β^*. Hence, it suffices to show that the cost of β^η is strictly less than the cost of β^*.

Let $\varphi : [0,1] \to \mathbb{R}$ be a function such that

$$\varphi(q) \equiv \underbrace{c(q)}_{\text{pre-offer cost}} + (1-\pi q)\underbrace{\mathbb{E}_{\mathbf{p}^*\sim\rho^*(\cdot|q,N)}[c(\beta^*(\cdot \mid \mathbf{p}^*, 0, N))]}_{\text{post-offer cost}},$$

where $c(\beta^*(\cdot \mid h_2))$ denotes $\mathbb{E}_{\mathbf{q}_2\sim\beta^*(\cdot|h_2)}[c(\mathbf{q}_2)]$. Then, $\varphi(\eta)$ is the cost of β^η, while $\varphi(0)$ is the cost of β^*. To show that $\varphi(\eta) < \varphi(0)$ for some $\eta > 0$, it suffices to show that $\varphi'(0) < 0$. By Bayes's rule, $\mu^*(q, N) = \frac{\pi - \pi q}{1-\pi q}$ and then

$$\rho^*(\cdot \mid q, N) = \frac{\pi - \pi q}{1-\pi q}\sigma^*(\cdot \mid H) + \frac{1-\pi}{1-\pi q}\sigma^*(\cdot \mid L).$$

Since

$$(1-\pi q)\mathbb{E}_{\mathbf{p}^*\sim\rho^*(\cdot|q,N)}\Big[c(\beta^*(\cdot \mid \mathbf{p}^*, 0, N))\Big] = \int \Big((\pi - \pi q)\sigma^*(\mathrm{d}p \mid H) + (1-\pi)\sigma^*(\mathrm{d}p \mid L)\Big)c(\beta^*(\cdot \mid p, 0, N)),$$

it follows that

$$\varphi'(0) = c'(0) + \int \frac{\partial}{\partial q}\Big((\pi - \pi q)\sigma^*(\mathrm{d}p \mid H) + (1-\pi)\sigma^*(\mathrm{d}p \mid L)\Big)c(\beta^*(\cdot \mid p, 0, N))\Big|_{q=0}$$

$$= -\pi \int \sigma^*(\mathrm{d}p \mid H) c(\beta^*(\cdot \mid p, 0, N)) < 0,$$

where we may interchange the partial differentiation and integral by the Leibniz rule.[15] ■

From the above lemmas, we conclude that the set $P^{\mathcal{E}}_{\text{acq}}(q_1) = P^{\mathcal{E}}_a(q_1^*) \cup P^{\mathcal{E}}_0(q_1^*)$ is empty for any $q_1^* \in \text{supp}(\beta^*(\cdot \mid h_1))$. That is, B will not acquire any post-offer information in any equilibrium.

A.2 Equilibria in Which S Never Offers Price H Are Pooling

We will show that the set $P^{\mathcal{E}}_n(q_1^*)$ is a singleton for any on-path q_1^*.

Lemma 5. *For any $q_1^* \in \text{supp}(\beta^*(\cdot \mid h_1))$, the set $P^{\mathcal{E}}_{\text{not}}(q_1^*)$ is non-empty; moreover, if $H \notin P^{\mathcal{E}}_{\text{not}}(q_1^*)$ then the set $P^{\mathcal{E}}_{\text{not}}(q_1^*)$ is a singleton.*

Proof. Recall that $\bigcup_{v \in V} \text{supp}(\sigma^*(\cdot \mid v)) \subseteq P^{\mathcal{E}}_{\text{acq}}(q_1^*) \cup P^{\mathcal{E}}_{\text{not}}(q_1^*)$. Since $P^{\mathcal{E}}_{\text{acq}}(q_1^*) = \varnothing$, it follows that $P^{\mathcal{E}}_{\text{not}}(q_1^*) \neq \varnothing$. Moreover, it is a singleton because otherwise, it would have at least two prices but then S of type L would choose the higher one only. ■

In any equilibrium $\mathcal{E}$ in which S never offers price H, B will acquire neither pre- nor post-offer information; moreover, S of each type v offers the (pooling) price p^*.

Lemma 6. *Let $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$ be any equilibrium in which S of each type v never offer price H. Then, it is a pure-strategy pooling equilibrium such that:*

- *S of each type v offers price p^* with probability 1.*
- *B does not acquire pre- or post-offer information on path: $\beta^*(h_t) = 0$ for any on-path histories $h_t \in H_1 \cup H_2$.*

Proof. Fix any $q_1^* \in \text{supp}(\beta^*(\cdot \mid h_1))$. If $\mathbf{x}_1 = H$ then B chooses $q_2^* = 0$ whatever price S charges. If $\mathbf{x}_1 = N$ then since the equilibrium price p^* is such that $P^{\mathcal{E}}_{\text{acq}}(q_1^*) = \varnothing$ and $P^{\mathcal{E}}_{\text{not}}(q_1^*) = \{p^*\}$ (Lemma 5), B chooses $q_2^* = 0$ at history $h_2 = (p^*, q_1^*, N)$. Then, B's equilibrium payoff is equal to $\mathbb{E}_{\mathbf{h}_3}[\mathbb{E}_{\mathbf{v} \sim \mu(\cdot \mid \mathbf{h}_3)}[\mathbf{v}]] - p^* - \lambda c(q_1)$, with history $\mathbf{h}_3 = (p^*, q_1^*, \mathbf{x}_1, 0, N)$. Because $\mathbb{E}_{\mathbf{h}_3}[\mathbb{E}_{\mathbf{v} \sim \mu(\cdot \mid \mathbf{h}_3)}[\mathbf{v}]] = \mathbb{E}_{\mathbf{v} \sim \pi}[\mathbf{v}]$ as the belief process is a martingale, B's equilibrium payoff is equal to $\mathbb{E}_{\mathbf{v} \sim \pi}[\mathbf{v}] - p^* - \lambda c(q_1^*)$. Now consider B's strategy of choosing $q_1 = q_2 = 0$ and buying at price p^* for sure. B's payoff from this strategy is $\mathbb{E}_{\mathbf{v} \sim \pi}[\mathbf{v}] - p^*$. Since q_1^* is optimal, B's payoff from choosing q_1^* cannot be lower than the payoff from this strategy: $\mathbb{E}_{\mathbf{v} \sim \pi}[\mathbf{v}] - p^* - \lambda c(q_1^*) \geq \mathbb{E}_{\mathbf{v} \sim \pi}[\mathbf{v}] - p^*$. Since this inequality holds with equality if and only if $q_1^* = 0$, it follows that $\text{supp}(\beta^*(\cdot \mid h_1)) = \{0\}$. ■

A.3 B Extracts Almost Full Surplus in Pooling Equilibria

We characterize an upper bound $\bar{p}$ of all (pooling) equilibrium prices $p^* > L$ and then show that for each $\epsilon > 0$, there exists some $\bar{\lambda} > 0$ such that for each $\lambda < \bar{\lambda}$, $\bar{p} \in [L, L + \epsilon)$. In a (pooling) equilibrium with price $p^* > L$, B acquires no information at any on-path histories (Lemma 6) and

[15] The derivative of the integrand is $-\pi c(\beta^*(p, 0, N))$, which is dominated by an integrable function because for each $p \in \text{supp}(\sigma^*(\cdot \mid H))$, B's accuracy would be at most $\bar{q} \in [0, 1]$ such that $c(\bar{q}) \leq H$. Hence, $|-\pi c(\beta^*(\cdot \mid p, 0, N))| \leq \pi H$.

then buys at any on-path histories. Hence, her equilibrium payoff is $\mathbb{E}_{\mathbf{v}\sim\pi}[\mathbf{v}] - p^*$. Next, consider a deviating strategy such that B chooses accuracy $q_1 = 0$ at history h_1 but accuracy $q_2 > 0$ at history $h_2 = (p^*, 0, N)$ and then she buys (with probability 1) if $\mathbf{x}_2 = H$ but does not (with probability 1) if $\mathbf{x}_2 = N$. This deviation yields B's payoff $\pi q_2(H - p^*) - \lambda c(q_2)$. Since the equilibrium payoff $\mathbb{E}_{\mathbf{v}\sim\pi}[\mathbf{v}] - p^*$ must be at least the deviating payoff, it follows that $\mathbb{E}_{\mathbf{v}\sim\pi}[\mathbf{v}] - p^* \geq \pi q_2(H - p^*) - \lambda c(q_2)$ for any $q_2 > 0$. Since $\mathbb{E}_{\mathbf{v}\sim\pi}[\mathbf{v}] = \pi H + (1-\pi)L$, it follows that for any $q_2 > 0$,

$$p^* \leq \bar{p} \equiv L + \frac{\pi(1-q_2)}{1-\pi q_2}(H - L) + \frac{\lambda c(q_2)}{1-\pi q_2}.$$

For each $\epsilon > 0$, there exists some $\underline{q} < 1$ such that any $q_2 \in (\underline{q}, 1)$, $\frac{\pi(1-q_2)}{1-\pi q_2}(H-L) < \frac{\epsilon}{2}$. Moreover, for each $\epsilon > 0$ and each $q_2 \in (\underline{q}, 1)$, there exists some $\bar{\lambda} > 0$ such that for any $\lambda < \bar{\lambda}$, $\frac{\lambda c(q_2)}{1-\pi q_2} < \frac{\epsilon}{2}$. Hence, $\bar{p} < L + \epsilon$.

A.4 Equilibria in Which S May Offer Price H Are Pareto-Dominated

In the rest of the proof, we assume that S of either type offers price H with non-zero probability in an equilibrium $\mathcal{E} = (\beta^*, \sigma^*, \mu^*)$. That is, $H \in \mathrm{supp}(\sigma^*(\cdot \mid H)) \cup \mathrm{supp}(\sigma^*(\cdot \mid L))$.

We make three preliminary observations. First, $H \in \mathrm{supp}(\sigma^*(\cdot \mid H))$ and $H \notin \mathrm{supp}(\sigma^*(\cdot \mid L))$. The reason is as follows. At price H, B will not acquire any information since her surplus from buying is at most 0 and thus her payoff will be negative if she acquires information. Since B does not acquire any information and S of type L charges price H with positive probability, B's expected value after seeing price H must be below H and she will not buy. Thus, S of type L gets zero profit from charging price H. But then, S of type L would obtain a higher profit by charging price L instead. This is a contradiction.

Second, $\{H\} = \mathrm{supp}(\sigma^*(\cdot \mid H)) \setminus \mathrm{supp}(\sigma^*(\cdot \mid L))$. To see this, suppose, by negation, that there is some $p < H$ such that $p \in \mathrm{supp}(\sigma^*(\cdot \mid H)) \setminus \mathrm{supp}(\sigma^*(\cdot \mid L))$. After seeing this price p, B believes that it comes from type H and thus will not acquire any information; and since p must be below her valuation, she buys with probability 1. We have two cases to consider. First, suppose that there is some $p' \in \mathrm{supp}(\sigma^*(\cdot \mid L))$ after which B does not acquire any post-offer information. We must have (i) $p \geq p'$ and (ii) $p \leq p'$. For (i), we note that if $p < p'$, S of type H would charge p' rather than p, which contradicts $p \in \mathrm{supp}(\sigma^*(\cdot \mid H))$. For (ii), we note that if $p > p'$, S of type L would charge p rather than p', which contradicts $p' \in \mathrm{supp}(\sigma^*(\cdot \mid L))$. From (i) and (ii), it follows that $p = p'$, which contradicts our assumption that $p \notin \mathrm{supp}(\sigma^*(\cdot \mid L))$ and $p' \in \mathrm{supp}(\sigma^*(\cdot \mid L))$. Second, suppose that, for all $p' \in \mathrm{supp}(\sigma^*(\cdot \mid L))$, with probability $U(p)$, B does not acquire any post-offer information, but with probability $1 - U(p)$ B acquires post-offer information with accuracy $q_2(p) > 0$. Note that if she acquires post-offer information, her information must be useful for her purchasing decision; that is, B buys if and only if $x_2 = H$. This implies that the profit for type L in the equilibrium is $p'U(p')$ for all $p' \in \mathrm{supp}(\sigma^*(\cdot \mid L))$. For S of type L to be willing to charge such a price, we must have $p'U(p') \geq p$. On the other hand, if type H charges any price $p' \in \mathrm{supp}(\sigma^*(\cdot \mid L))$, the profit for type H is $p'(U(p') + (1 - U(p'))q_2(p))$, which is strictly higher than the profit $p'U(p')$ for type L. Since $p'U(p') \geq p$, the profit for type H when charging price p' is strictly higher than his profit when charging price p. This contradicts the assumption that p is in the support of type H's strategy. Therefore, $\{H\} = \text{supp}(\sigma^*(\cdot \mid H)) \setminus \text{supp}(\sigma^*(\cdot \mid L))$.

Third, $\text{supp}(\sigma^*(\cdot \mid L))$ is a singleton. Recall that B does not acquire any post-offer information in any equilibrium history h_2 (Lemma 4). Hence, if $\text{supp}(\sigma^*(\cdot \mid L))$ were not a singleton, then S of type L would only charge the highest price in the support.

It is helpful to summarize these observations. If there is an equilibrium $\mathcal{E}$ in which S of either type offers price H with non-zero probability, then it is either a separating equilibrium or a semi-separating equilibrium such that:

- There is a price $p^* < H$ such that $\text{supp}(\sigma^*(\cdot \mid L)) = \{p^*\}$, and $\text{supp}(\sigma^*(\cdot \mid H)) = \{H\}$ (in the case of a separating equilibrium) or $\text{supp}(\sigma^*(\cdot \mid H)) = \{H, p^*\}$ (in the case of a semi-separating equilibrium).
- B does not acquire any post-offer information in any equilibrium history h_2 (Lemma 4), and buys in any equilibrium history h_3.

In what follows, we show that the above equilibrium is Pareto-dominated. In the equilibrium, S of type L has profit p^*. B believes that price H must be charged by type H, and thus will not acquire any information. For type L to be willing to charge price p^* rather than H, it must be that, for any $p \in \text{supp}(\sigma^*(\cdot \mid H))$, B must buy with probability at most p^*/p; in addition, for type H to be willing to charge $p \in \text{supp}(\sigma^*(\cdot \mid H))$ rather than p^*, B must buy with probability at least p^*/p. That is, for any $p \in \text{supp}(\sigma^*(\cdot \mid H))$, B must buy with probability p^*/p. That is, the equilibrium must satisfy the following conditions:

1. S of type H charges price $p \in \text{supp}(\sigma^*(\cdot \mid H))$, and S of type L charges price p^*.
2. B acquires no information on the equilibrium path:
   (a) At any history h_1, B acquires no pre-offer information.[16]
   (b) At any history h_2, B acquires no post-offer information.
   (c) At history h_3, B buys with probability 1 if $p = p^*$; B buys with probability p^*/H if $p = H$; otherwise, B buys with probability 0.

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

1. S of each type v offers price p^*.
2. B acquires no information on the equilibrium path.
   (a) At any history h_1, B acquires no pre-offer information.
   (b) At any history h_2, B acquires no post-offer information.
   (c) At history h_3, B buys with probability 1 if $p = p^*$; otherwise, B buys with probability 0.

Note that if the former assessment is indeed an equilibrium, so is the latter. This is because the conditions for the latter to be an equilibrium—that is, S of type L is willing to offer price p rather

[16]This follows from a similar argument as in the proof of Lemma 5. Specifically, since B's equilibrium payoff is the same as that from buying with probability 1 at any equilibrium prices (i.e., p^* or H), information is not useful in changing her purchasing decision. That is, if B had acquired pre-offer information in an equilibrium, this information had no value. But since it is costly to acquire information, it is not optimal to choose $q_1^* > 0$.

than L, 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. This is because S of each type v has payoff p and thus they are the same in both equilibria, but B's payoff is strictly higher in the latter equilibrium.

References

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.

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, 1991.

D. Fudenberg, D. K. Levine, and J. Tirole. *Infinite-Horizon Models of Bargaining with One-Sided Incomplete Information*, 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.

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.

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