Supplementary Material for Bargaining and Information Acquisition

Kalyan Chatterjee Miaomiao Dong Tetsuya Hoshino

September 5, 2024

1 Proof of Proposition 4

1.1 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 1. *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*

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

   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*

$$\tilde{\pi}_1(p) = \frac{\lambda c'(\tilde{q}(p))}{H - p}. \tag{2}$$

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 (2).

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,

$$\tilde{\pi}_1(p)(1 - \tilde{q}(p))(H - p) + (1 - \tilde{\pi}_1(p))(L - p) + \lambda c(\tilde{q}(p)) = 0. \tag{3}$$

Substituting (2) into (3), we obtain the desired equation (1).

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 (1). Since c is strictly convex and $\frac{H-L}{H-p} - q > 1 - q > 0$, the LHS of (1) is strictly increasing in $\tilde{q}(p)$. It is also continuous in $\tilde{q}(p)$. Moreover,

$$\lambda c'(0)((H-L)/(H-p) - 0) + \lambda c(0) = 0 < p - L,$$
$$\lambda c'(1)((H-L)/(H-p) - 1) + \lambda c(1) = \infty > p - L,$$

which ensures the existence and uniqueness of q that solves (1). □

Lemma 2. *The functions $\tilde{\pi}_1$ and $\tilde{q}$, defined by equations* (1) *and* (2)*, 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 (1). 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 (1) 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 (1),

$$\tilde{q}'(p) = -\frac{\lambda c'(\tilde{q}(p))\frac{H-L}{(H-p)^2} - 1}{\lambda c''(\tilde{q}(p))(\frac{H-L}{H-p} - \tilde{q}(p))}. \tag{4}$$

Substituting it into (2), we have

$$\tilde{\pi}_1'(p) = \frac{1 - \tilde{q}(p)\frac{\lambda c'(\tilde{q}(p))}{H-p}}{H - L - \tilde{q}(p)(H-p)} = \frac{1 - \tilde{q}(p)\tilde{\pi}_1(p)}{H - L - \tilde{q}(p)(H-p)}. \tag{5}$$

Note that $\tilde{\pi}_1'(p) > 0$ for any p such that $\tilde{\pi}_1(p) \leq 1$. This is because both the denominator and the numerator of the RHS of (5) 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 (5), 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 (1) 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 (1) as

$$\lambda c(\tilde{q}(p)) \left[\frac{c'(\tilde{q}(p))}{c(\tilde{q}(p))}\left(\frac{H-L}{H-p} - q(p)\right) + 1\right] = p - L. \tag{6}$$

Since $\lim_{q\to 1} c'(q) = \infty$, we have $\lim_{q\to 1} \frac{c'(q)}{c(q)} = \infty$.[1] 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 (6) goes to infinity. Since the RHS is finite, we have $\lambda c(\tilde{q}(p)) \to 0$. Using (2), we can rewrite (1) as

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

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 (5) and applying $\tilde{q}(p) \to 1$ and $\tilde{\pi}_1(p) \to 1$, we have $\tilde{\pi}_1'(p) \to 1$. □

[1] 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 $M > 0$, $c'(q)/c(q) \leq M$ 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 $\log(c(q)/c(q_0)) \leq M(q - q_0)$ and thus $c(q) \leq c(q_0)\mathrm{e}^{M(q-q_0)}$, but this contradicts the assumption that $\lim_{q\to 1} c(q) = \infty$.

Next, we construct a mixed-learning PBE.

Lemma 3. *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*

$$\tilde{\pi}_1(p^*) = \frac{\pi}{\pi + (1 - \pi)y^*}. \tag{7}$$

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 $\mathbf{x} = N$.*

   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 1.

Second, we derive (7). 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 1, this posterior probability, which we have denoted by $\tilde{\pi}_1(p^*)$, must satisfy (2). Since these two representations must coincide,

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

which is the desired (7).

We show that there exists $\underline{p}_\lambda \in (L, H)$ such that for any $p^* \in (\underline{p}_\lambda, H)$, (7) has a solution y^*. By Lemma 2, $\tilde{\pi}_1$ 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 (7).

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 2. 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 3.

Lemma 4. *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 1, 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 (1), 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 (4). Since the denominator of the RHS in (4) 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 2, $\tilde{q}(p) \to 0$ 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 (2), $\tilde{f}(p) = \tilde{\pi}_1(p)(H-L)/(H-p) - 1$. Since $\tilde{\pi}_1(p) \to 1$ as $\lambda \to 0$ for any $p \in (L, H)$ by Lemma 2, 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, $\text{supp}(\bigcup_v \sigma(\cdot \mid v)) \cap (L, H-\delta)$, 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 1, 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

$$\tilde{U}_H(p) = pz(p) + (1-z(p))\tilde{q}(p)p = u_L + \tilde{q}(p)(p - u_L). \tag{8}$$

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 (4),

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

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 2, $\tilde{q}(p) \to 1$ 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 2, $\tilde{\pi}_1(p|\lambda) \to 1$ 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 2, $\tilde{\pi}_1(p|\lambda)$ 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 3, 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

$$U_B(p) = (1 - \pi)y(p)(L - p) + \pi(H - p), \tag{9}$$

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 $\tilde{x}(p) = \tilde{\sigma}(\{p\} \mid H)$.[2] For PBE $\tilde{\mathcal{E}}_\lambda(p)$, let $\tilde{P} = \text{supp}(\bigcup_v \tilde{\sigma}(\cdot \mid v)) \cap (L, H)$. By Step 3, there is a single price $p \in \text{supp}(\bigcup_v \tilde{\sigma}(\cdot \mid v)) \cap (L, H-\delta)$. 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

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

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

$$\begin{aligned}
U_B(p) - \tilde{U}_B(p) &> (1-\tilde{x}(p))[\pi(H-p-\delta) + (1-\pi)(L-p)y(p)] \\
&= (1-\tilde{x}(p))\pi\left(H - p - \delta - \frac{1-\tilde{\pi}_1(p)}{\tilde{\pi}_1(p)}(p-L)\right),
\end{aligned} \tag{10}$$

where $(1-\pi)y(p) = \pi\frac{1-\tilde{\pi}_1(p)}{\tilde{\pi}_1(p)}$ by Bayes' rule. By Lemma 2, $\tilde{\pi}_1(p) \to 1$ as $\lambda \to 0$. For any sufficiently small δ, we have $H - p - \delta > \delta$. For each $p \in (L, H)$, let

$$\lambda_p^2 = \sup\left\{\lambda' \in (0, \lambda_{p,\delta}) : H - p - \delta - \frac{1-\tilde{\pi}_1(p)}{\tilde{\pi}_1(p)}(p-L) \geq 0 \quad \forall \lambda < \lambda'\right\}. \tag{11}$$

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 (9), 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 (9), 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 2, 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 $p' \in (p, p+\epsilon')$.

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

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 (9), 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 (8) (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

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

where λ_p^2 is defined in (11). 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 (8), 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 (9), 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 4). 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 3. 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 2. Hence, $U_B(p^*) \to \pi(H-p^*)$. Since $\underline{p}_\lambda \to L$ as $\lambda \to 0$ by Lemma 3, 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 $\lim_{q\to 1} c'(q) = \infty$.

1.2 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 1, B's first order condition with respect to q is replaced with

$$q = \begin{cases} 1 & \text{if } \pi_1(p)(H-p) \geq \lambda c'(1) \\ (c')^{-1}\left(\frac{\pi_1(p)(H-p)}{\lambda}\right) & \text{if } \pi_1(p)(H-p) < \lambda c'(1). \end{cases} \tag{12}$$

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

$$\lambda c'(1)\left(\frac{H-L}{H-p}-1\right)+\lambda c(1) < p-L, \tag{13}$$

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 (12) 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 (3), 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 3 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 (13) then our proof for Lemma 3 and Proposition 4 is valid without any modification. If λ is such that there exists a $p \in (L, H)$ satisfying (13), then, multiplying $H-p$ on both sides of (13), we have

$$\lambda c'(1)\,(p-L)+\lambda c(1)(H-p) < (p-L)(H-p). \tag{14}$$

Since there is a $p \in (L, H)$ satisfying (14), there exists an interval $(p_1^\lambda, p_2^\lambda)$ such that (14), or equivalently (13) 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 3 and Proposition 4 hold.