# Learning Before Trading: On the Inefficiency of Ignoring Free Information*

Doron Ravid[†]
*University of Chicago*

Anne-Katrin Roesler[‡]
*University of Toronto*

Balázs Szentes[§]
*London School of Economics*

June 30, 2021

**Abstract**

This paper analyzes a bilateral trade model in which the buyer's valuation for the object is uncertain and she can privately purchase any signal about her valuation. The seller makes a take-it-or-leave-it offer to the buyer. The cost of a signal is smooth and increasing in informativeness. We characterize the set of equilibria when learning is free, and show they are strongly Pareto ranked. Our main result is that when learning is costly but the cost of information goes to zero, equilibria converge to the worst free-learning equilibrium.

## 1 Introduction

Recent developments in information technology have given consumers access to new information sources that allow them to learn about products prior to trading. For example,

*We would like to thank Ben Brooks, Yeon-Koo Che, Tommaso Denti, Piotr Dworczak, Benjamin Hebert, Navin Kartik, John Leahy, Elliot Lipnowski, R. Vijay Krishna, George Mailath, Daniel Martin, Stephen Morris, Philip Reny, Andrzej Skrzypacz, Rakesh Vohra, and Ming Yang for extremely useful comments and suggestions. Elena Istomina provided superb research assistance.

†Kenneth C. Griffin Department of Economics, University of Chicago, Chicago, IL 60637. Email: dravid@uchicago.edu, website: http://www.doronravid.com

‡Department of Economics, University of Toronto, Toronto, ON M5S 3G7, Canada. Email: ak.roesler@utoronto.ca, website: http://www.akroesler.com

§Department of Economics, London School of Economics, London, WC2A 2AE, UK, Email: b.szentes@lse.ac.uk, website: http://personal.lse.ac.uk/szentes

online resources enable buyers to learn about a mechanic's reputation, a contractor's reliability, or an over-the-counter (OTC) asset's value. This information acquisition often takes place before the buyers learn the terms of trade. Indeed, to get a price quote, customers may need to bring their cars to the mechanic, have a contractor over, or waste their first contact with an OTC dealer.[1] Because the buyer's willingness-to-pay depends on her information about the product, the seller's price depends on what he expects the buyer to learn. Conversely, the seller's pricing strategy determines what information is worth learning for the buyer. For example, there may be no point in knowing more about the value of an asset if the buyer is already sure it is below its price. Therefore, the buyer's information acquisition depends on the seller's expected prices. The goal of this paper is to study this mutual dependency between the buyer's learning strategy and the seller's pricing policy.

We consider a stylized model in which the seller has a single object for sale and full bargaining power. Initially, the buyer does not know anything about the value of the object except its prior distribution, which is assumed to be regular.[2] We model the buyer's learning as *flexible* information acquisition; that is, she can purchase *any* signal about her valuation privately. Then, the seller, without observing the buyer's learning strategy and her signal realization, sets a price. Signals are costly and we assume this cost is a *smooth and strongly increasing* function of the signal's informativeness. Below, we explain these assumptions in detail. Our aim is to characterize the set of equilibria of this game. We are especially interested in the limit where the buyer's cost vanishes. This limit appears to be particularly relevant in a world where information is becoming cheaper and more accessible to consumers. To this end, we parameterize the cost by a multiplicative constant and consider the limit when this parameter converges to zero.

We now describe the buyer's action space and the cost of information. The demand of the buyer, which is the probability of trade occurring at a given price, is fully determined by the distribution of her posterior value estimate. In turn, the seller's profit from any given price is pinned down by the buyer's demand. As a consequence, trade outcomes are fully determined by the distribution of the buyer's posterior estimate. The prior distribution is a mean-preserving spread of any such distribution because each signal

[1] A stylized feature of OTC markets is that prices quoted on a second call can be dramatically higher than the first one; see Bessembinder and Maxwell (2008) or Zhu (2012).

[2] A distribution is said to be regular if the corresponding virtual valuation is well-defined and strictly increasing.

contains less information than the valuation itself. Since the buyer can choose any signal, we identify her action space with the set of these distributions and define the cost of information acquisition on this set. To require this function to be smooth, we appeal to a generalized notion of differentiability, because the domain is a set of cumulative distribution functions (CDFs). In particular, we postulate that the cost function is Fréchet differentiable.

We now turn to our main assumption on the cost of information. A signal is more informative than another if its induced distribution over posterior-value estimates is a mean-preserving spread of that of the other. Thus, a cost function is said to be monotonic in the signal's informativeness if mean-preserving spreads cost more. As we argue, a cost function is monotonic whenever its Fréchet derivative, which is a function itself, is convex. Our main assumption is somewhat stronger than monotonicity: in addition to requiring the Fréchet derivative to be convex, we assume this derivative at a given CDF is strictly convex on the convex hull of the CDF's support. Imposing this assumption in addition to monotonicity resembles stipulating that a differentiable increasing function has a strictly positive derivative. Therefore, one can interpret our assumption as supposing that acquiring more information has a strictly positive cost at the margin.

Monotonicity of the learning cost implies the seller randomizes in every equilibrium in which the buyer learns. To see why, suppose an equilibrium exists in which the seller sets a fixed price and the buyer receives an informative signal about her valuation. Then, this signal must be binary, indicating whether the buyer should trade or not. The reason is that any other signal can be made less informative, and hence cheaper, while still leading to the same trading decisions. The seller's best response to such a binary signal is to charge the expected valuation of the buyer conditional on one of the two signal realizations. To get a contradiction, notice the buyer is strictly better off by not learning, irrespective of which of these prices is set. If the price is the lower signal realization, the buyer always trades, so learning yields no benefit. If the price is the higher signal realization, the buyer's surplus from trade is zero, so she could again profitably deviate by saving the cost of learning and not trading.

Our aforementioned strong monotonicity assumption also has important implications for the buyer's equilibrium learning strategy. We show the support of the buyer's equilibrium signal is an interval and the buyer's demand generated by this signal makes the seller indifferent between setting any price on its support. This indifference condition implies the buyer's equilibrium CDF is a *truncated Pareto distribution*, and hence her

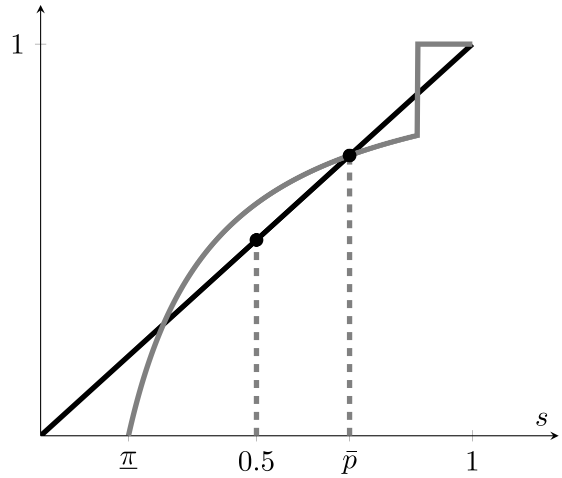


Figure 1: An Illustration of the best and worst equilibria in the uniform case.

equilibrium demand is unit elastic.

As mentioned above, our main objective is to characterize equilibrium outcomes as the buyer's cost vanishes. To this end, we first consider the case in which learning is free. We show this case admits multiple equilibria, all of which can be Pareto ranked. In the Pareto-best equilibrium, which maximizes both players' payoffs across all free-learning equilibria, the buyer learns her valuation perfectly. The equilibrium price in this equilibrium is smaller than in any other free-learning equilibrium. The Pareto-worst equilibrium turns out to be the unique equilibrium in which the buyer's posterior estimate is distributed according to a truncated Pareto distribution.

Figure 1 illustrates the best and worst free-learning equilibria when the prior is uniform on $[0, 1]$. In the Pareto-best equilibrium, the buyer learns her valuation perfectly; thus, the distribution of her value estimate is also uniform and therefore is represented by the 45-degree line. In this case, the seller's equilibrium price is .5, his profit is .25, and the buyer's payoff is .125. The buyer's CDF in the Pareto-worst equilibrium is depicted as a gray curve in Figure 1. In this worst equilibrium, the seller's profit, $\underline{\pi}$, is approximately .2, the price is $\bar{p}\,(\approx .715)$, and the buyer's payoff is only slightly above .04. Therefore, the buyer's payoff is less than one-third of her payoff in the perfect-learning equilibrium.

At first, it may appear counter-intuitive that there are equilibria in which the buyer does not learn perfectly although information is free. In the Pareto-worst equilibrium described above, the seller's price, $\bar{p}$, is defined by the highest intersection of the Pareto curve and the prior-value distribution. At this point, the mean-preserving spread constraint binds; that is, the integral of the Pareto curve and the prior CDF on $[0, \bar{p}]$ coincide. We call such a price *separating*. The important property of a separating price is that the buyer never confuses a value below such a price with a value above it. That is, a value below $\bar{p}$ never generates the same signal realization as a value above $\bar{p}$. Hence, the buyer would not gain anything by learning more, because this Pareto signal already reveals if her valuation is above or below $\bar{p}$, which is the only information she needs to know in order to trade ex-post efficiently.

Our main result is that as the buyer's learning cost vanishes, equilibria converge to a Pareto-worst free-learning equilibrium. For an explanation, recall that when learning is costly, the buyer's equilibrium CDF is a truncated Pareto. The limit of truncated Pareto distributions is also a truncated Pareto, so the same must hold for the costless limit, which is a free-learning equilibrium. Hence, as costs shrink, we obtain a free-learning equilibrium in which the buyer's demand is unit elastic. All that remains is to recall the fact mentioned above, namely, that the unique such equilibrium yields the Pareto-worst free-learning equilibrium outcome.

The main takeaway from our paper is that possessing information might be significantly better than having cheap access to it. When information is costly, buyers must have incentives to acquire it. In equilibrium, prices fail to provide these incentives, so buyers choose to ignore large amounts of information even when costs are minuscule. In turn, this ignorance triggers prices that are too high relative to those in a full-information environment, leading to considerable welfare losses.

More broadly, our analysis warns about the danger of approximating environments with freely available data by models in which learning is costless. The reason is that, even with easy-to-access information, learning may not be literally free, because decision makers might still need to incur attention or cognitive-processing costs when presented with new information. Hence, assuming learning has tiny but positive costs may be more accurate than assuming learning is free. Our results show that the difference between these two assumptions is not merely a manner of minutiae: equilibria of the free-learning model may be significantly better for both contracting parties than the equilibrium of the costly learning model even if learning costs are minuscule. In fact, since obtaining full information is a weakly dominant strategy, one might even assume that, under free learning, traders will coordinate on the Pareto-dominant, full-information equilibrium, whereas the Pareto-*worst* equilibrium is selected when costs are vanishingly small. As such, our work suggests that reasoning about scenarios in which information is cheap using a free-learning model may result in misleading conclusions.

In terms of applications, our model may offer a different angle for looking at markets where there is heterogeneity in traders' information. A large literature documents that the performance of traders depends on how knowledgeable they are; see, for example, Grubb (2015) and Handel and Schwartzstein (2018) for extensive surveys. Our theory might speak to those environments where such a knowledge differential is due to heterogeneity of the traders' information acquisition. To be more specific, consider an asset market that is populated by two kinds of buyers: experienced traders and novices. Experienced traders may already possess a significant amount of information related to the asset while novices are yet to learn about it. Our theory then suggests that, even if the relevant information is easy to obtain, veteran traders can get significantly better terms of trade than novices.[3]

This paper contributes to the large and growing literature on information design (e.g., Kamenica and Gentzkow (2011), Bergemann and Morris (2013)). In particular, our results serve as a cautionary tale about interpreting recent papers characterizing consumer and producer surplus pairs that can arise as an equilibrium outcome for *some* information structure.[4] Of particular relevance is Roesler and Szentes (2017), who consider a setting similar to ours in which the seller observes the buyer's signal before setting a price. Their key result identifies the signal-equilibrium pair that maximizes the buyer's payoff. The buyer-optimal signal turns out to be the same Pareto signal as in our worst free-learning equilibrium. At first glance, their result might seem surprising given that the worst free-information equilibrium minimizes the buyer's payoff. However, because the seller in Roesler and Szentes's (2017) model sets a price only after observing the buyer's signal, he can set any profit-maximizing price, and in the buyer-optimal equilibrium, he chooses the lowest such price. By contrast, in our model, the seller's price must also justify the buyer's signal choice, forcing the seller to choose a separating point. Thus, our analysis suggests the same information structure can lead to two drastically different outcomes. Which outcome is selected depends on the mechanism through which trade occurs.

Costly consumer learning is extensively studied in the context of rational inattention. This literature was initiated by Sims (1998, 2003, 2006); for more recent contributions,

[3]Consistent with this implication, O'Hara et al. (2018) report that traders get a better price for the same bond on the same day if they traded more frequently in the past.

[4]See for example, Bergemann et al. (2015), Roesler and Szentes (2017), Kartik and Zhong (2019) and Yang (2019).

see, for example, Matějka (2015) and Ravid (2019). In these models, information cost is proportional to the resulting expected reduction in entropy. In contrast to this literature, we treat the cost of information in an abstract way and do not assume such a particular form. Still, one can show our results go through even when the buyer's information costs are given by expected entropy reduction.[5]

In the context of auctions, many papers conclude that the buyers' incentives for acquiring costly information about their valuations depend on the selling mechanism. Persico (2000) shows buyers acquire more information in a first-price auction than in a second-price one, provided that their signals are affiliated. Compte and Jehiel (2007) argue dynamic auctions generate higher revenue than simultaneous ones. Shi (2012) characterizes the revenue-maximizing auction in private-value settings. In all of these models, the seller commits to a mechanism before the buyers decide how much information to acquire. By contrast, we consider environments where the monopolist cannot commit and best responds to the buyer's signal structure.[6]

Yang (2020) studies a security-design problem related to our model. In his model, a seller offers an asset-backed security to a buyer who can then flexibly learn about the asset's returns. Unlike in our model, the buyer observes the seller's offer before deciding how much to learn. The author shows the optimal security is a debt contract.

Finally, our work also adds to the recent literature on the relationship between free-learning equilibria and the vanishing-cost limits of equilibria, see, for example, Yang (2015), Hoshino (2018), Morris and Yang (2019), and Denti (2019). This literature primarily focuses on static flexible-learning models in which all players have access to the same information. In these models, free information always yields a perfect-learning outcome, and the vanishing-cost limit can be viewed as an equilibrium-selection device in a symmetric-information game. By contrast, learning is asymmetric in our model because the seller cannot acquire information about the buyer's valuation. Consequently, perfect-learning corresponds to an asymmetric-information game with a substantially smaller equilibrium set than our free-learning game. And, indeed, our vanishing-cost limit selects a free-learning outcome that is simply not an equilibrium under full information.

[5]We note entropy reduction costs are not Fréchet differentiable. Thus, to show our results hold for entropy reduction, we generalize our strong monotonicity assumption to non-differentiable costs, see the Online Appendix.

[6]Another strand of the literature analyzes the seller's incentives to reveal information about the buyers' valuations prior to participating in an auction; see, for example, Milgrom and Weber (1982), Ganuza (2004), Bergemann and Pesendorfer (2007), Ganuza and Penalva (2010), and Li and Shi (2017).

## 2 The Model

A seller, S (he), has an object to sell to a single buyer, B (she). B's valuation, $\mathbf{v}$, takes values in $[0,1]$ according to the CDF $F_0$ whose expected value is $\bar{v} = \int v \, \mathrm{d}F_0(v) > 0$. We assume $F_0$ is **regular**, meaning it has a strictly positive density, $f_0$, on $[0,1]$ and that $v - (1 - F_0(v))/f_0(v)$ is strictly increasing in $v$. B does not see $\mathbf{v}$ but can choose to observe *any* signal, $\mathbf{s}$, at a cost that depends on the signal's informativeness. Below, we describe the set of signals available to the buyer and the associated cost in detail. Then S, *without observing* B's information-acquisition strategy and signal realization, makes a take-it-or-leave-it price offer, $p \in [0,1]$. S wants to maximize revenue, whereas B's utility from trade when her valuation is $v$ and the price is $p$ is $v-p$ if she accepts, and 0 if she does not. We assume B purchases the good if and only if her expected valuation conditional on her signal weakly exceeds $p$, which is without loss.[7] Both players are risk-neutral expected-payoff maximizers.

*Signal structures and B's action space.* We allow B to choose *any* signal, $\mathbf{s}$, to learn about $\mathbf{v}$. Now, B's expected payoff from trade at any given price depends only on her posterior mean, $\mathbb{E}[\mathbf{v}|\mathbf{s}]$. Therefore, the marginal distribution of $\mathbb{E}[\mathbf{v}|\mathbf{s}]$ determines both B's trade surplus and the probability she purchases the good at any given price, which, in turn, is sufficient for calculating S's profits and optimal prices. In other words, the game's trade outcome depends only on the marginal distribution of B's posterior mean valuation, and so we identify each signal with the CDF of this marginal.[8] More precisely, taking $\mathcal{F}$ to be the space of all CDFs over $[0,1]$, we let B choose any element of $\mathcal{F}$ that can arise as the marginal CDF of $\mathbb{E}[\mathbf{v}|\mathbf{s}]$ for some $\mathbf{s}$. We denote this set by $\mathcal{A}$ and describe it formally below.

As observed by Gentzkow and Kamenica (2016), $F$ is the CDF of the marginal distribution of B's posterior mean for some signal if and only if it is a mean-preserving contraction of the prior, $F_0$. Recall that $F \in \mathcal{F}$ is a **mean-preserving spread** of $F' \in \mathcal{F}$ (denoted by $F \succeq F'$) if and only if

$$\int_0^x (F - F') \, \mathrm{d}s \geq 0 \text{ for all } x \text{ with equality for } x = 1. \tag{1}$$

[7] Standard arguments deliver that, in equilibrium, B must trade if indifferent. This assumption therefore has no effect on our results but makes the analysis simpler.

[8] This method of modeling flexible information is common in the information-design literature—see, for example, Gentzkow and Kamenica (2016), Roesler and Szentes (2017), and Dworczak and Martini (2019).

The CDF $F$ is a **strict mean-preserving spread** of $F'$ (denoted by $F \succ F'$) if both $F \succeq F'$ and $F' \neq F$.[9] Letting

$$I_F(x) := \int_0^x (F_0 - F)\ \mathrm{d}s,$$

B's actions space, $\mathcal{A}$, can be defined as the set $\{F \in \mathcal{F} : I_F(x) \geq 0 \text{ for all } x \text{ and } I_F(1) = 0\}$. In what follows, we refer to CDFs in $\mathcal{A}$ as signals.

*The cost of information acquisition.* Information acquisition is costly. In general, different information structures generating the same distribution of posterior expectations might come at different costs. However, because B's expected payoff from trading depends only on the distribution of this posterior expectation, $F$, she would always use the least expensive signal structure that leads to $F$. In fact, B may even randomize to get $F$. Thus, we can evaluate the cost of $F$ by the expected cost of the cheapest randomization that generates it, resulting in a *convex* cost function,

$$C : \mathcal{A} \to \mathbb{R}_+.$$

We require the function $C$ to be sufficiently smooth. More precisely, we endow $\mathcal{F}$ with the $\mathcal{L}_1$-norm (denoted by $\|\cdot\|$)[10] and assume $C$ is **Fréchet differentiable**; that is, each $F \in \mathcal{A}$ admits a Lipschitz function,[11] $c_F : [0,1] \to \mathbb{R}$, such that for every $F' \in \mathcal{A}$,

$$C(F') - C(F) = \int c_F\ \mathrm{d}(F' - F) + o\left(\|F' - F\|\right), \tag{2}$$

where $o$ is a function that equals 0 at zero and $\lim_{x \searrow 0} [o(x)/x] = 0$. We refer to $c_F$ as $C$'s **derivative at** $F$.[12]

To discuss our next assumption, ranking signals in terms of their informativeness is useful. We say that $\mathbf{s}$ is more informative than $\mathbf{s}'$ if observing $\mathbf{s}$ is the same as observing $\mathbf{s}'$ and another informative variable, $\mathbf{t}$. One can easy verify that, if $F$ and $F'$ are the CDFs of $\mathbb{E}[\mathbf{v}|\mathbf{s}]$ and $\mathbb{E}[\mathbf{v}|\mathbf{s}']$, respectively, then $F$ is a mean-preserving spread of $F'$ whenever $\mathbf{s}$

[9] Notice $\succeq$ is reflexive and anti-symmetric, meaning $F \succeq F'$ and $F' \succeq F$ if and only if $F = F'$.

[10] That is, the norm that maps any Borel measurable $\phi : [0,1] \to \mathbb{R}$ to $\|\phi\| = \int_0^1 |\phi(x)| \mathrm{d}x$. Restricted to the set of CDFs over $[0,1]$, this norm metrizes weak* convergence; see, for example, Machina (1982).

[11] The standard definition of Fréchet differentiability requires $c_F$ to define a continuous linear function from $\mathcal{L}_1[0,1]$ to $\mathbb{R}$. This requirement is equivalent to $c_F$ having an almost-everywhere derivative in $L^\infty[0,1]$; that is, $c_F$ must be Lipschitz. Our results still hold, however, if we require $c_F$ to be continuous (rather than Lipschitz).

[12] Although irrelevant to our analysis, we note $c_F$ is unique only up to the addition of an affine function.

is more informative than $\mathbf{s}'$.[13] Conversely, every mean-preserving contraction of $F$ can be attained as the marginal distribution of $\mathbb{E}[\mathbf{v}|\mathbf{s}'']$ for some $\mathbf{s}''$ that is less informative than $\mathbf{s}$.[14]

The above implies mean-preserving spreads have higher costs whenever learning more information is more expensive. For an explanation, note that if $F$ is a mean-preserving spread of $F'$, then every $\mathbf{s}$ that induces $F$ admits a less-informative, and therefore cheaper, $\mathbf{s}'$ that yields $F'$ as the distribution of B's posterior valuation. Therefore, the *cheapest* $\mathbf{s}'$ generating $F'$ cannot cost more than the least expensive information structure attaining $F$.

Given the above, we say $F$ is **(strictly) more informative** than $F'$ if and only if $F$ is a (strict) mean-preserving spread of $F'$, and that $C$ is **monotone** if $C(F) \geq C(F')$ whenever $F$ is more informative than $F'$. Next, we show that $C$ is monotone if and only if its Fréchet derivative is convex.[15]

**Claim 1** *Let $C$ be Fréchet differentiable. Then, $C$ is monotone if and only if $c_F$ is convex for each $F \in \mathcal{A}$.*

For the intuition behind the claim and for better understanding the concept of Fréchet differentiability, we restrict attention to signals whose support lies in a finite set, say, $\{s_1, \ldots, s_N\}$. Then, each $F \in \mathcal{F}$ can be represented by a vector in the $n$-dimensional simplex $(\alpha_1, \ldots, \alpha_N) \in \Delta^n$ for which $F = \sum_{n=1}^{N} \alpha_n \mathbf{1}_{[s_n,1]}$.[16] In this case, the function $C$ is a mapping from $\Delta^n$ to $\mathbb{R}$, and the Fréchet derivative at $F$ evaluated at $s_n$, $c_F(s_n)$, is $C$'s partial derivative with respect to the probability of $s_n$; that is, $\partial C(F)/\partial \alpha_n = c_F(s_n)$. Thus, the marginal cost of a small shift from $F$ to $F'$ is the sum of the marginal cost at each signal realization times the change in each realization's probability, that is, $\int c_F \, \mathrm{d}(F' - F)$. Of course, if $F' \succeq F$, this quantity is positive whenever $c_F$ is convex.

[13] To see why, let $\mathbf{x} = \mathbb{E}[\mathbf{v}|\mathbf{s}']$ and $\mathbf{y} = \mathbb{E}[\mathbf{v}|\mathbf{s}]$. Showing $\mathbf{y}$ is a mean-preserving spread of $\mathbf{x}$ is equivalent to proving $\mathbb{E}[\mathbf{y}|\mathbf{x}] = \mathbf{x}$. Towards this goal, observe $\mathbf{x}$ is $\mathbf{s}'$-measurable, and so $\mathbb{E}[\mathbf{v}|\mathbf{x}] = \mathbb{E}[\mathbb{E}[\mathbf{v}|\mathbf{x}, \mathbf{s}']|\mathbf{x}] = \mathbb{E}[\mathbb{E}[\mathbf{v}|\mathbf{s}']|\mathbf{x}] = \mathbf{x}$ by the Law of Iterated Expectations. Therefore, $\mathbb{E}[\mathbf{y}|\mathbf{x}] = \mathbb{E}\left[\mathbb{E}[\mathbf{v}|\mathbf{s}]|\mathbf{x}\right] = \mathbb{E}\left[\mathbb{E}[\mathbf{v}|\mathbf{s}', \mathbf{t}, \mathbf{x}]|\mathbf{x}\right] = \mathbb{E}\left[\mathbf{v}|\mathbf{x}\right] = \mathbf{x}$,where the second equality follows from $\mathbf{s}$ being more informative than $\mathbf{s}'$ and $\mathbf{x}$ being $\mathbf{s}'$-measurable, and the third equality from the Law of Iterated Expectations.

[14] This claim follows from Proposition 1 of Gentzkow and Kamenica (2016).

[15] See Machina (1982), Hong et al. (1987), Chatterjee and Krishna (2011), and Cerreia-Vioglio et al. (2017) for related results for functions whose domain is $\mathcal{F}$ rather than $\mathcal{A}$.

[16] For any $A \subseteq [0, 1]$, $\mathbf{1}_A$ is the indicator function that is equal to 1 on $A$, and 0 otherwise. So, $\mathbf{1}_{[x,1]}$ is the CDF corresponding to a unit atom at $x$.

Our main assumption requires $c_F$ to be not only convex but also strictly convex on the support of $F$.

**Assumption 1** *For each $F \in \mathcal{A}$, $c_F$ is convex and strictly convex on* co(supp $F$)*.*

Next, we explain the restriction that Assumption 1 imposes on the cost function $C$ in addition to strict monotonicity, which requires that $C(F) > C(F')$ if $F \succ F'$. Assumption 1 requires that, whenever $F \succ F'$, the marginal cost of a small shift from $F$ to $F'$ to be strictly positive. By contrast, strict monotonicity only requires that these shifts have a strictly positive marginal cost almost everywhere.[17] Thus, imposing Assumption 1 in addition to smoothness and strict monotonicity resembles stipulating that a smooth, strictly increasing function has a strictly positive derivative everywhere.[18]

*Strategies and payoffs.* A mixed strategy for S is a random price, represented by a CDF over prices, $H \in \mathcal{F}$, whereas a strategy for B is a signal, $F \in \mathcal{A}$.[19] If B's signal is $F$, S's expected payoff from the random price $H$ is given by

$$\Pi(H, F) = \int p(1 - F(p-)) \, \mathrm{d}H(p),$$

where $F(p-)$ denotes $\sup_{s<p} F(s)$. We denote S's maximal profit by $\pi_F := \max_{p\in[0,1]} \Pi(p, F)$ and the set of profit-maximizing prices by $P(F) = \arg\max_{p\in[0,1]} \Pi(p, F)$.[20] In Appendix B, we establish continuity of S's maximal profit and upper hemicontinuity of the profit-maximizing prices, $P(\cdot)$.[21]

If S's randomization over prices is $H$, B's expected payoff from the signal $F$ is

$$U_\kappa(H, F) = \int \int_0^s (s - p) \, \mathrm{d}H(p) \, \mathrm{d}F(s) - \kappa C(F),$$

[17]Formally, Assumption 1 requires $\int c_{F''} \, \mathrm{d}(F - F') > 0$ whenever $F'' = F + \alpha(F' - F)$ for some $\alpha \in [0, 1)$, whereas strict monotonicity allows this quantity to be zero over a zero-Lebesgue measure set of $\alpha$'s. For an example of a strictly monotone function violating Assumption 1, consider $C(F) = 1 - 0.5 \int F^2(x) \, \mathrm{d}x$. One can show $C$ is strictly monotone, and Fréchet differentiable with $c_F(x) = \int_0^x F(s) \, \mathrm{d}s$. Observe that $\mathrm{d}c_F/\mathrm{d}x = F(x)$, meaning $c_F$ is convex, but not strictly convex over co(supp $F$) if supp $F$ is not convex. We note this $C$ is concave and we do not know whether a convex counterexample exists.

[18]For example, the function $x^3$ is smooth and strictly increasing but has a zero derivative $x = 0$.

[19]Since B's objective is concave (due to the convexity of $C$), S's objective is affine, and $\mathcal{A}$ is convex, any mixed strategy of $B$ can be replaced by its its average marginal without influencing payoffs and trade outcomes. Hence, without loss off generality, we assume B uses a pure strategy.

[20]We slightly abuse notation and let $\Pi(p, F)$ denote $\Pi(\mathbf{1}_{[p,1]}, F)$.

[21]Notice S's profit is only upper semicontinuous; thus, said properties do not follow from Berge's Maximum Theorem.

where $\kappa \in \mathbb{R}_+$ is a constant parameterizing B's cost of information.

*Equilibrium Definition and Existence.* An equilibrium is a pair, $(H, F) \in \mathcal{F} \times \mathcal{A}$, such that

1. $H$ maximizes $\Pi(\cdot, F)$ over $\mathcal{F}$;

2. $F$ maximizes $U_\kappa(H, \cdot)$ over $\mathcal{A}$.

Because B's best response and S's (mixed) best response are upper hemicontinuous, and non-empty-convex-compact valued, an equilibrium exists by Kakutani's Fixed-Point Theorem.[22]

*Truncated Pareto Distributions.* As mentioned in the introduction, the set of truncated Pareto distributions plays an important role in our analysis. To formally define this set, for each $\pi \in [0, 1]$ and $t \in [\pi, 1]$, let

$$G_{\pi,t}(s) = \mathbf{1}_{[\pi,t)}(s) \left(1 - \frac{\pi}{s}\right) + \mathbf{1}_{[t,1]}(s). \tag{3}$$

We refer to the set $\{G_{\pi,t}\}$ as the set of truncated Pareto distributions and an element of $\{G_{\pi,t}\} \cap \mathcal{A}$ as a **Pareto signal**.

## 2.1 Examples of Cost of Learning

This section provides three examples of cost functions that satisfy our assumptions and describes their Fréchet derivatives.

**Example 1.** (Constant Marginal Cost) Fix some strictly convex Lipschitz function $c : [0, 1] \to \mathbb{R}_+$. Define

$$C(F) = \int c \, \mathrm{d}F.$$

[22] Convexity (compactness) of the best response follows from concavity and linearity (continuity and upper semicontinuity) of B's and S's objectives, respectively. Upper hemicontinuity of B's best response follows from Berge's Maximum Theorem. To see S's mixed best response, $F \mapsto \arg\max_{H \in \mathcal{F}} \Pi(H, F)$, viewed as a correspondence, is upper hemicontinuous, consider a convergent sequence of signals, $F_n \to F_\infty$, and suppose $H_n \in \arg\max_{H \in \mathcal{F}} \Pi(H, F_n)$ converges to $H_\infty$. Because $\mathcal{F}$ is compact, it is enough to show $H_\infty$ is an S best response to $F_\infty$, that is, supp $H_\infty \subseteq P(F_\infty)$. Now, on the one hand, supp$(\cdot)$ is lower hemicontinuous, and so $p_\infty \in$ supp $H_\infty$ only if a sequence $p_n \in$ supp $H_n$ exists that attains $p_\infty$ as its limit. On the other hand, $P(\cdot)$ is upper hemicontinuous (see Appendix B), and so the limit of any convergent sequence $p_n \in$ supp $H_n \subseteq P(F_n)$ is in $P(F_\infty)$. Therefore, $p_\infty \in$ supp $H_\infty$, only if $p_\infty \in P(F_\infty)$, that is, supp $H_\infty \subseteq P(F_\infty)$.

Then, $C$'s Fréchet derivative equals $c$ for all $F$.[23] When $c(s) = s^2 - \bar{v}^2$, this example assigns each $F$ a cost equal to its variance. We explain how our results specialize to this cost function when B's value is uniformly distributed throughout our analysis.

**Example 2.** (Increasing Marginal Cost) Fix some strictly Lipschitz convex $c : [0,1] \to \mathbb{R}_+$ and a strictly increasing, convex, Lipschitz and differentiable $\psi : \mathbb{R} \to \mathbb{R}$. Then, the function

$$C(F) = \psi\left(\int c \, \mathrm{d}F\right)$$

satisfies our assumptions. Indeed, by the chain rule, the above cost function is Fréchet differentiable, with the derivative being given by

$$c_F(\cdot) = \psi'\left(\int c \, \mathrm{d}F\right) c(\cdot),$$

which is strictly convex for all $F$.

**Example 3.** (Quadratic Costs) Let $c : [0,1] \times [0,1] \to \mathbb{R}_+$ be some strictly convex, Lipschitz, symmetric function; that is, $c(s_1, s_2) = c(s_2, s_1)$ for all $s_1, s_2 \in [0,1]$. Assume further that $c$ is positive semidefinite, that is, $\int\int c \, \mathrm{d}(F - F')\mathrm{d}(F - F') \geq 0$ for all $F, F' \in \mathcal{F}$. Then, the cost function[24]

$$C(F) = \frac{1}{2}\int\int c(s_1, s_2) \, \mathrm{d}F(s_1)\mathrm{d}F(s_2)$$

is convex and Fréchet differentiable, with the derivative being given by the strictly convex Lipschitz function,

$$c_F(\cdot) = \int c(\cdot, s_2) \, \mathrm{d}F(s_2).$$

# 3 Costless Learning

In this section, we analyze the set of equilibria when learning is free, that is, when $\kappa = 0$. We first provide geometric characterizations of the best responses of S and B, respectively. Then, we use these characterizations to identify the set of payoff profiles that arise in equilibrium. We also show the free-learning equilibrium set can be strongly Pareto ranked, with the best equilibrium being the one given by perfect learning, that is, $F = F_0$.

[23] Example 1 describes the set of posterior separable costs (see Gentzkow and Kamenica, 2014, and Caplin et al., 2017) that are also mean-measurable; that is, that depend only on a signal's induced posterior-mean distribution.

[24] Example 3 belongs to the Fréchet differentiable subset of the quadratic functional form introduced by Machina (1982) and studied in Chew et al. (1991).

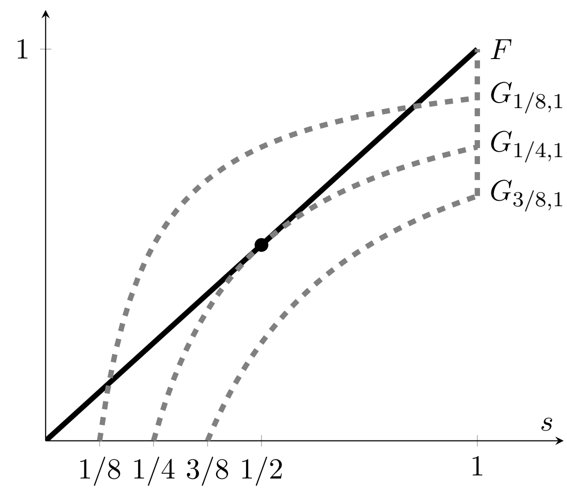


Figure 2: The seller's best response against the uniform distribution.

### 3.1 The Seller's Best Responses

We begin by characterizing the set of profit-maximizing prices. To this end, we first describe S's iso-profit curves on the price-cumulative probability space. Note that if the price is $p$ and the probability that B's valuation is strictly less than $p$ is $y$, S's profit is $p(1-y)$. Hence, the iso-profit curve in this space corresponding to a given profit, say, $\pi\,(>0)$, is defined by

$$\{(p, y) : y \in [0, 1], p(1-y) = \pi\}.$$

Of course, if $p < \pi$, the profit cannot exceed $p$ and no $y \in [0, 1]$ exists that generates $\pi$. Otherwise, for each $p \in [\pi, 1]$, the cumulative probability, $y$, that guarantees profit $\pi$ is $1-\pi/p$. Observe that $1-\pi/p$ is the CDF corresponding to the Pareto distribution parameterized by $\pi$. Because $p \leq 1$, we conclude the iso-profit curve of the seller corresponding to profit $\pi$ is essentially identical to the truncated Pareto distribution, $G_{\pi,1}$.

These iso-profit curves can be used to analyze S's best response against B's signal distribution as illustrated in Figure 2 for the case of a uniform $F$. Note that lower iso-profit curves correspond to larger profits. In addition, the set of feasible outcomes is $\{(p, F(p-)) : p \in [0, 1]\}$. Therefore, S's profit is defined by the smallest $\pi$ such that the curve $G_{\pi,1}(s)$ is weakly below that of $F(s-)$. In Figure 2, three iso-profit curves are depicted as the gray dashed contours, and the middle one, $G_{1/4,1}$, is the highest iso-profit curve below $F$, so the profit of S is 1/4. Furthermore, the set of optimal prices, $P(F)$, are those values at which $F$ is tangent to the largest iso-profit curve below it. In Figure 2,

$p = 1/2$ is the only point of tangency. Since iso-profit curves are strictly increasing, the signal $F$ must also be strictly increasing at any point of tangency, and hence any such points must lie in the support of $F$. The following lemma summarizes these observations.

**Lemma 1** *Fix any $F \in \mathcal{A}$. Then,*

*(i) for all $s \in [0,1]$, $F(s-) \geq G_{\pi_F,1}(s-)$; and*

*(ii) $P(F) = \{p \geq \pi_F : F(p-) = G_{\pi_F,1}(p-)\} \subseteq \text{supp } F$.*

Part (i) states that B's CDF is first-order stochastically dominated by the Pareto distribution parameterized by S's profit, $\pi_F$. Part (ii) says the set of profit-maximizing prices are those signals at which B's CDF essentially coincides with this Pareto distribution.

### 3.2 The Buyer's Best Responses

If $S$ sets price $p$ and B learns her valuation perfectly, she makes an ex-post efficient trading decision. To make such decisions, B's signal must reveal whether the true valuation is above or below $p$. In what follows, we characterize the set of such signal distributions.

Note that if B chooses $F$ and the price is $p$, her expected payoff from trade is

$$\int_p^1 (s-p) \ \mathrm{d}F(s) = (1-p) - \int_p^1 F(s) \ \mathrm{d}s,$$

where the equality follows from integration by parts. Of course, when information is free, perfect learning is a best response to any pricing strategy of S. In fact, using the previous equation, the increase in B's payoff from switching from $F$ to perfect learning can be expressed as

$$\int_p^1 (F-F_0)(s) \ \mathrm{d}s = \int_0^1 (F-F_0)(s) \ \mathrm{d}s - \int_0^p (F-F_0)(s) \ \mathrm{d}s = I_F(p) \geq 0, \tag{4}$$

where the inequality follows from (1). Thus, the slackness in the signal's information constraint at $p$, $I_F(p)$, is the benefit of obtaining all remaining information. Whenever this benefit is zero, B cannot gain from learning more. Because B can only lose from learning less, $F$ is optimal for B if and only if $I_F(p) = 0$. Intuitively, $I_F(p) = 0$ means $p$ separates $F$'s realizations: Either B's true valuation and the signal realization generated by $F$ are smaller or both of them are larger than $p$. In what follows, we refer to such a price as $F$-**separating** and we denote the collection of such prices by $S(F)$, that is,

$$S(F) = \{p \in [0,1] : I_F(p) = 0\}.$$

In summary, if S sets price $p$, the signal $F$ is B's best response if and only if $p \in S(F)$. The next lemma extends this argument to the case where S randomizes over prices: the signal $F$ generates the same payoff to B as perfect learning if and only if S only charges $F$-separating prices.

**Lemma 2** *The signal $F$ is a best response against $H$ if and only if* $\operatorname{supp} H \subseteq S(F)$.

Next, we show two geometric properties of separating prices. First, the graph of $F$ and $F_0$ must intersect at every $F$-separating price. Second, given an $F$-separating $p$, one can find a larger price arbitrarily close to $p$ at which $F$ lies below $F_0$.

**Lemma 3** *Suppose $F \in \mathcal{A}$ and $p \in S(F)$. Then,*

*(i) $F$ is continuous at $p$ and $F(p) = F_0(p)$.*

*(ii) For every $p' > p$, a $p'' \in (p, p')$ exists such that $F(p'') \leq F_0(p'')$.*

Let us describe the intuition behind the lemma. To see why part (i) holds, recall that $p$ is $F$-separating if the signal reveals whether the valuation is above or below $p$. Hence, the probability that B observes a signal realization below $p$ must be the same as the probability that her valuation is below $p$; that is, the CDFs $F$ and $F_0$ must coincide at $p$. To explain part (ii), note it is enough to show no $p' > p$ exists such that for each $p'' \in (p, p')$, $F(p'') > F_0(p'')$. Since the valuation is larger than $p$ if and only if the signal is larger than $p$, the mean-preserving spread relationship between the valuation and the signal remains even when conditioning on them being larger than $p$. It follows that the conditional value distribution has more mass closer to the boundaries of its support, $p$ and 1, than the conditional signal distribution does. Therefore, the CDF of the conditional signal cannot be above that of the conditional valuation just to the right of $p$. In other words, no $p' > p$ exists such that for each $p'' \in (p, p')$,

$$\left(F\left(p''\right) - F\left(p\right)\right) / \left(1 - F\left(p\right)\right) > \left(F_0\left(p''\right) - F_0\left(p\right)\right) / \left(1 - F_0\left(p\right)\right).$$

By part (i), $F(p) = F_0(p)$, so the desired conclusion follows.

### 3.3 Free-Learning Equilibrium Characterization

We now turn to characterizing the set of free-learning equilibrium payoffs. We begin by showing that an equilibrium price is never below the full-information monopoly price.

Before stating this result, note that because $F_0$ is regular, the function $\Pi(\cdot, F_0)$ is strictly quasiconcave; see Figure 3 below for an illustration. Therefore, a unique profit-maximizing price exists under $F_0$, denoted by $p_0^*$. The next lemma states that $p_0^*$ is below the support of S's randomization in every free-learning equilibrium.

**Lemma 4** *Let $(H, F)$ be a free-learning equilibrium. Then,* $\operatorname{supp} H \subseteq [p_0^*, 1]$.

For an explanation, consider a free-learning equilibrium in which B learns signal $F$ and S charges price $p$. We need to demonstrate that $p$ is weakly higher than the full-information monopoly price, $p_0^*$. Since $\Pi(\cdot, F_0)$ is strictly quasiconcave, this function is increasing below $p_0^*$ and decreasing above it; see Figure 3 below. So, to conclude that $p \geq p_0^*$, it is enough to argue S's marginal profit at $p$ is weakly negative under $F_0$. To make this argument, observe first that the marginal profit at $p$ is weakly negative under $F$ because $p$ is profit-maximizing when B acquires signal $F$. Further observe that Lemma 2 and part (i) of Lemma 3 imply B's demand as well as S's profit from setting price $p$ are the same under full information and under $F$. In addition, part (ii) of Lemma 3 implies a small increase in $p$ results in a larger reduction in B's demand under perfect learning than under $F$. Consequently, the decrease in profit due to a small increase in $p$ under full information is even larger than under $F$. Thus, it follows that the profit at $p$ is locally decreasing under full information, as required.

Next, we show S never randomizes in equilibrium. More specifically, we prove that if $(H, F)$ is a free-learning equilibrium, $H$ specifies an atom of size one at a price that would generate profit $\pi_F$ even if B learns perfectly instead of getting signal $F$. To state this result precisely, for each $\pi$, let $X_\pi$ be the set of prices that yield profit $\pi$ under $F_0$, namely,

$$X_\pi := \{p : \Pi(p, F_0) = \pi\}.$$

The next lemma states that S's equilibrium price is the largest element of $X_{\pi_F}$. Before we state this result, note that because the function $\Pi(\cdot, F_0)$ is strictly quasiconcave, $X_\pi$ contains at most two such prices for every $\pi$. Moreover, $\Pi(\cdot, F_0)$ attains any value between 0 and $\pi_{F_0}$ because it is continuous.[25] Therefore, for each $\pi \in [0, \pi_{F_0}]$, $X_\pi$ is non-empty and contains at most two prices. Let $\bar{p}_\pi$ be the higher of those prices, that is, $\bar{p}_\pi = \max X_\pi$. The following lemma says that if B's free-learning equilibrium signal is $F$, then S charges $\bar{p}_{\pi_F}$ for sure.

[25] This claim follows from the Intermediate Value Theorem and that charging zero generates zero profit.

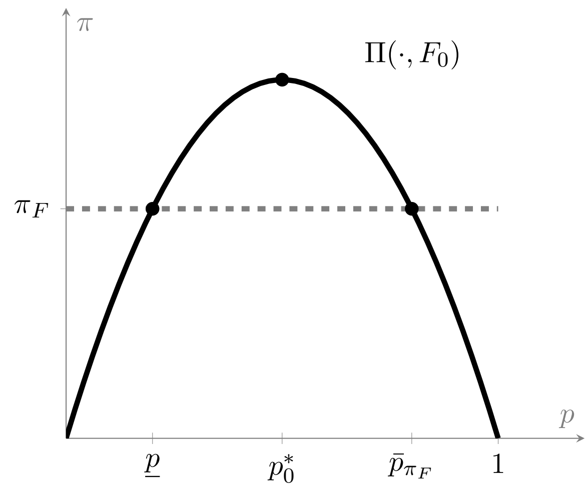


Figure 3: The seller's profit under perfect learning when $F_0$ is uniform on $[0,1]$, along with the seller's profit maximizing price, $p_0^*$, and the set $X_{\pi_F} = \{\underline{p}, \bar{p}_{\pi_F}\}$ when $\pi_F = 4/25$. In this case, $p_0^* = 0.5$, $\underline{p} = 0.2$, and $\bar{p}_{\pi_F} = 0.8$.

**Lemma 5** *Let $(H, F)$ be a free-learning equilibrium. Then,* supp $H = \{\bar{p}_{\pi_F}\}$.

Let us explain the argument of this lemma's proof. Consider any price, $p$, in the support of $H$. Since $p$ is profit-maximizing under $F$, $\Pi(p, F) = \pi_F$. By Lemma 2 and part (i) of Lemma 3, the demand as well as the profit at $p$ are the same under $F$ and under perfect learning. Thus, charging $p$ under perfect learning also generates profit $\pi_F$; that is, $p \in X_{\pi_F}$. Recall that $X_{\pi_F}$ has at most two elements, with $\bar{p}_{\pi_F}$ defined as the larger one. To complete the proof, it is enough to show that if $X_{\pi_F}$ has two elements, $p$ cannot be the smaller one, say, $\underline{p}$. To see why, note $\underline{p} < p_0^*$, as illustrated by Figure 3, and $p_0^* \le p$, by Lemma 4, so

$$\underline{p} < p_0^* \le p.$$

We now turn to the main result of this section, which characterizes the set of payoff profiles that can arise in equilibrium. Before stating this result, we introduce an additional piece of notation. Let $\underline{\pi}$ denote S's minmax profit, that is, the smallest possible profit that can be generated by some learning strategy when S responds optimally. Formally,[26]

$$\underline{\pi} = \min_{F \in \mathcal{A}} \max_{p \in [0,1]} \Pi(p, F) = \min_{F \in \mathcal{A}} \pi_F.$$

[26] Roesler and Szentes (2017) attain the minmax profit with a Pareto signal, so $\underline{\pi}$ is well defined. Alternatively, one can show $\underline{\pi}$ is well defined by observing that $\mathcal{A}$ is compact, $\Pi(\cdot, F)$ is upper semicontinuous, and $F \mapsto \pi_F$ is continuous (see Appendix B for a proof of the last fact).

Theorem 1 shows S's minimal and maximal equilibrium profits are $\underline{\pi}$ and $\pi_{F_0}$, respectively, and that S can attain any profit in between.[27] If S's equilibrium profit is $\pi$, B's equilibrium payoff is given by her expected utility under full information when S's price is $\bar{p}_\pi$.

**Theorem 1** *A free-learning equilibrium* $(H, F)$ *exists such that* $\pi_F = \pi$ *and* $U_0(H, F) = u$ *if and only if* $\pi \in [\underline{\pi}, \pi_{F_0}]$ *and* $u = U_0(\bar{p}_\pi, F_0) = \int_{\overline{p}_\pi}^{1} (s - \overline{p}_\pi)\,\mathrm{d}F_0(s)$.

The "only if" part of this theorem implies that in a free-learning equilibrium, S can never attain a profit above his maximal profit when B collects full information. This result is a straightforward consequence of Lemma 5. Recall that this lemma states that if B's signal is $F$, the equilibrium price is the largest price that generates profit $\pi_F$ under perfect learning, $\overline{p}_{\pi_F}$. But if learning is perfect, S can achieve $\pi_{F_0}$ by setting the optimal price instead of $\overline{p}_{\pi_F}$, showing $\pi_F \leq \pi_{F_0}$. The theorem also states that if B's signal is $F$, her equilibrium payoff is the same as if she learns perfectly and S charges a price of $\overline{p}_{\pi_F}$. This conclusion follows from the facts that S sets a price of $\overline{p}_{\pi_F}$ in every equilibrium where his profit is $\pi_F$ (see Lemma 5) and that perfect learning is always a best response when information is free.

The "if" part of the theorem's proof is constructive. Specifically, we find an equilibrium for each $\pi \in (\underline{\pi}, \pi_{F_0})$ such that S's profit is $\pi$. Existence of an equilibrium with profit $\underline{\pi}$ follows from the equilibrium payoff set being closed.[28] Figure 4 illustrates our construction, which obtains an equilibrium by applying two modifications to the $\pi$-iso profit curve, $G_{\pi,1}$. The first modification creates a CDF with separating and profit-maximizing price $p$ that gives S a profit of $\pi$. To get this CDF, we replace the realizations in the lowest $q$ quantiles of $G_{\pi,1}$ with realizations from the same quantiles of $F_0$. The resulting CDF is equal to $F_0$ at any $x$ such that $F_0(x) \leq q$, to $G_{\pi,1}$ when $G_{\pi,1}(x) \geq q$, and to $q$ otherwise.

[27]Note $\underline{\pi} < \pi_{F_0}$. This inequality follows from two facts. First, each of the two profits is associated with a unique Pareto signal, $G_{\underline{\pi},\bar{t}}$ and $G_{\pi_{F_0},t}$ (see Roesler and Szentes, 2017). And second, $S(G_{\underline{\pi},\bar{t}}) \cap \mathrm{supp}\ G_{\underline{\pi},\bar{t}} \neq \varnothing$ (by Theorem 2), whereas regularity of $F_0$ delivers $S(G_{\pi_{F_0},t}) \cap \mathrm{supp}\ G_{\pi_{F_0},t} = \varnothing$.

[28]To see why the equilibrium payoff set is closed, note first that upper hemicontinuity of the players' best-response correspondences implies closedness of the set of equilibrium strategy profiles. Because both players' strategies live in a compact set, the set of equilibrium strategy profiles is closed only if it is compact. As such, every convergent sequence of equilibrium payoffs is associated with a convergent sequence of equilibria. Because both players' maximal value is continuous in the other player's strategy, the payoffs generated by the limit equilibrium equal the limit of the equilibrium payoff sequence. Hence, the limit of every converging sequence of equilibrium payoffs is itself an equilibrium payoff; that is, the equilibrium payoff set is closed.

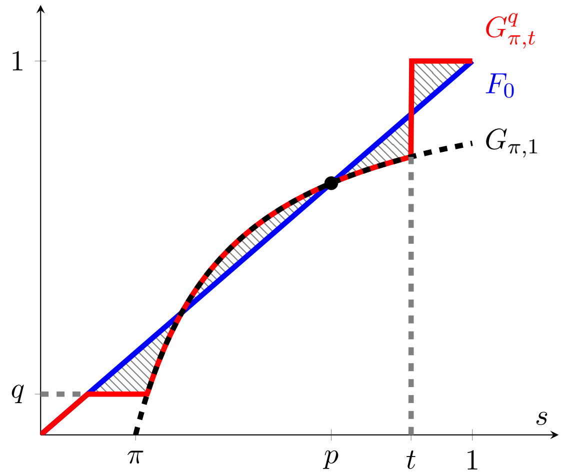


Figure 4: A constructed free-learning equilibrium, $(\mathbf{1}_{[p,1]}, G^q_{\pi,t})$.

The value of $q$ is determined so that the maximal profit $\pi$ is attained by setting a separating price, $p$; that is, the two shaded areas to the left from $p$ in Figure 4 have the same size. The Intermediate Value Theorem guarantees the existence of such a $q$. This CDF, however, fails to be a signal, due to having too large a mean. To change this CDF into a signal, we reduce the mean through the second modification: the distribution is truncated at some value $t \in (p, 1)$, resulting in the CDF $G^q_{\pi,t}$. To see that a $t$ exists for which $G^q_{\pi,t}$ is a signal, note first that the integral of $F_0 - G^q_{\pi,t}$ on the interval $[0, p]$ is zero for all $t \in (p, 1)$ because $p$ was a separating price. Thus, to guarantee $F_0$ is a mean-preserving spread of $G^q_{\pi,t}$, the truncation value, $t$, must be chosen so that the integral of $F_0 - G^q_{\pi,t}$ on $[p, 1]$ is also zero. To obtain such a $t$, one can again use the Intermediate Value Theorem, and hence, the CDF $G^q_{\pi,t}$ is indeed a signal; see Figure 4. Since $t \in (p, 1)$, the price $p$ still yields profit $\pi$, and remains separating and profit-maximizing. Thus, having S offer $p$ and B use $G^q_{\pi,t}$ gives a free-learning equilibrium.

Using Theorem 1, we can deduce that free-learning equilibria are strongly Pareto ranked; that is, B prefers one free-learning equilibrium to another if and only if S does as well.

**Corollary 1** *All free-learning equilibria are strongly Pareto ranked. That is, for any two free-learning equilibria,* $(H, F)$ *and* $(H', F')$*,*

$$\Pi(H, F) \geq \Pi(H', F') \text{ if and only if } U_0(H, F) \geq U_0(H', F').$$

For an explanation, consider two equilibria with corresponding profits $\pi_1$ and $\pi_2$ such that $\pi_1 < \pi_2$. To conclude that these equilibria are Pareto ranked, we need to demonstrate that the consumer surplus is also larger in the second one. To this end, we first argue the price is smaller in the second equilibrium. Recall that, in any equilibrium featuring profit $\pi$, S charges the higher price attaining $\pi$ under perfect learning, $\bar{p}_\pi$. Further recall that the function $\Pi(\cdot, F_0)$ is strictly quasi-concave, maximized at $p_0^*$, and $p_0^* < \bar{p}_\pi$, as illustrated in Figure 3. So, when $\pi$ increases, the price $\bar{p}_\pi$ decreases and moves toward $p_0^*$. Therefore, since $\pi_1 < \pi_2$, the equilibrium price in the second equilibrium is indeed smaller, that is, $\bar{p}_{\pi_2} < \bar{p}_{\pi_1}$. It remains to argue that B's payoff is decreasing in the equilibrium price. As explained above, B's surplus from trade remains unchanged if she were to choose to observe her value, that is, B's payoffs in the two equilibria are $U(\bar{p}_{\pi_1}, F_0)$ and $U(\bar{p}_{\pi_2}, F_0)$, respectively. Since $\bar{p}_{\pi_2} < \bar{p}_{\pi_1}$ it immediately follows that $U(\bar{p}_{\pi_1}, F_0) < U(\bar{p}_{\pi_2}, F_0)$.

To conclude this section, we characterize the set of free-learning equilibrium payoffs in an example. We then carry this example through the paper and use it to illustrate our results even when learning is costly.

**Example**. Suppose the value distribution, $F_0$, is uniform on the interval $[0, 1]$. By Theorem 1, S can attain any payoff between his perfect-learning profit, $\pi_{F_0} = 0.25$, and his minmax-profit, $\underline{\pi} \approx 0.2$ (see Roesler and Szentes, 2017). Next, we compute S's equilibrium price corresponding to any equilibrium profit level $\pi \in [\underline{\pi}, \pi_{F_0}]$. Since the value distribution is uniform, B's demand is $1 - p$ if the price is $p$. Therefore, the set $X_\pi$ consists of the solutions to

$$p(1-p) = \pi.$$

This quadratic equation admits at most two roots, $p = 0.5\left(1 \pm \sqrt{1-4\pi}\right)$. By Lemma 5, S charges the larger of these roots, $\bar{p}_\pi = 0.5\left(1 + \sqrt{1-4\pi}\right)$, in any equilibrium in which his profit is $\pi$. B's utility in such an equilibrium is

$$u = \int_{\bar{p}_\pi}^{1} (s - \bar{p}_\pi)\, \mathrm{d}s = 0.25 - 0.5\pi - 0.25\sqrt{1-4\pi}.$$

To summarize, if the value distribution is uniform, the set of free-learning equilibrium payoff profiles is given by

$$\left\{\left(\pi,\, 0.25 - 0.5\pi - 0.25\sqrt{1-4\pi}\right) :\, \pi \in [\underline{\pi},\, \pi_{F_0}]\right\}.$$

# 4 Costly Learning

This section accomplishes two goals. First, we provide an equilibrium characterization in our model of costly learning. In particular, B's equilibrium signal is shown to belong to the family of Pareto signals. Second, we prove the main result of this paper: as the cost of learning vanishes, equilibria converge to the worst free-learning equilibrium.

## 4.1 Equilibrium Characterization

The next result provides a partial characterization of the equilibrium when B's learning cost satisfies Assumption 1.

**Proposition 1** *Suppose $(H, F)$ is an equilibrium in the $\kappa > 0$ game. Then,*

*(i)* $\text{supp } H = \text{supp } F = \text{co}(\text{supp } F)$*, and*

*(ii) $F$ is a Pareto signal.*

Part (i) of this proposition states that the supports of B's signal and S's randomization coincide. Furthermore, this support is an interval. From these two observations, it is straightforward to conclude part (ii). The reason is that S must be indifferent on supp $H$, so each price in supp $H$ must generate the same profit. Therefore, part (i) implies B's equilibrium signal, $F$, must coincide with an iso-profit curve over its support. Because the iso-profit curve is a Pareto distribution truncated at 1, $F$ must be a Pareto signal.[29]

Next, we explain how to establish part (i). The key step is to show S charges every price between any two possible signal realizations; that is,

$$\text{co}(\text{supp } F) \subseteq \text{supp } H. \tag{5}$$

To prove this inclusion, it is enough to show that S offers a price between any two possible realizations of B's signal; that is, supp $H \cap (x, y) \neq \emptyset$ if $x, y \in \text{supp } F$.[30] Suppose first that $F$ places atoms at both $x$ and $y$. Then, B can profitably deviate by bunching together the signals $x$ and $y$; that is, instead of observing these signals, she only learns that the

[29]The seller's randomization is also behind the emergence of Pareto distributions in the context of robust pricing, see Bergemann and Schlag (2008), Carrasco et al. (2018), and Du (2018).

[30]To see why this claim is sufficient, note that if (5) does not hold, co(supp $F$) includes a non-empty interval $(x, y)$ that never contains S's price, that is, supp $H \cap (x, y) = \emptyset$. In the Appendix, we show that if $(x, y)$ is maximal among such intervals, $x, y$ must both lie in supp $F$.

signal is in $\{x, y\}$. By Assumption 1, this bunching strictly reduces B's learning cost. Moreover, because S never sets a price in $(x, y)$, such a bunching leaves B's trade surplus unchanged. To understand why, note that conditional on the original signal being $x$, the buyer trades if and only if the price is weakly less than $x$, irrespective of whether the signals are bunched together. The only difference in trading decisions is that if the original signal is $y$, B trades if the price is $y$ but rejects this price after the bunching. Because the buyer breaks even in both cases, this difference does not change her payoff. We conclude that when $F$ has atoms at both $x$ and $y$, it cannot be a best response against $H$ if supp $H \cap (x, y) = \emptyset$. If either $x$ or $y$ have zero mass according to $F$, one can construct a profitable deviation in a similar fashion by pooling together small neighborhoods of $x$ and $y$. Finally, notice that

$$\text{co}(\text{supp } F) \subseteq \text{supp } H \subseteq \text{supp } F \subseteq \text{co}(\text{supp } F),$$

where the first inclusion is just (5), the second follows from the observation that S never sets a price that is not a possible signal realization (see part (ii) of Lemma 1). This chain of inclusion implies part (i) of the proposition.

Next, we introduce costly information acquisition in the example of Section 3 and illustrate how to characterize an equilibrium in this example.

**Example** (continued). Recall that the value distribution, $F_0$, is uniform on $[0, 1]$. Suppose now that the cost of each signal, $F$, is given by its variance, $C(F) = \int (s^2 - \bar{v}^2)\, \mathrm{d}F$ and $\kappa = 1$.

An important feature of this example is that B's problem of finding a best response can be reduced to maximizing the value of an integral with respect to a measure. Indeed, given a randomization of S over prices, $H$, B chooses $F \in \mathcal{A}$ to maximize $\int u_H(s)\mathrm{d}F(s)$, where

$$u_H(s) = \int_{p \leq s} (s - p)\mathrm{d}H(p) - (s^2 - \bar{v}^2).$$

A useful consequence of this observation is that if the integrand $u_H$ is concave, applying a mean-preserving contraction to B's signal, $F$, increases her payoff. Moreover, if $u_H$ is affine on the convex hull of $F$'s support, B is indifferent between choosing $F$ and the uninformative signal, $\mathbf{1}_{[0.5,1]}$, which specifies an atom of size one at the prior expectation, $\bar{v} = 0.5$.[31]

[31] Since $\int s\ \mathrm{d}F(s) = \int v\ \mathrm{d}F_0(v) = \overline{v}$ for each $F \in \mathcal{F}$, $\overline{v} \in \text{co}(\text{supp } F)$.

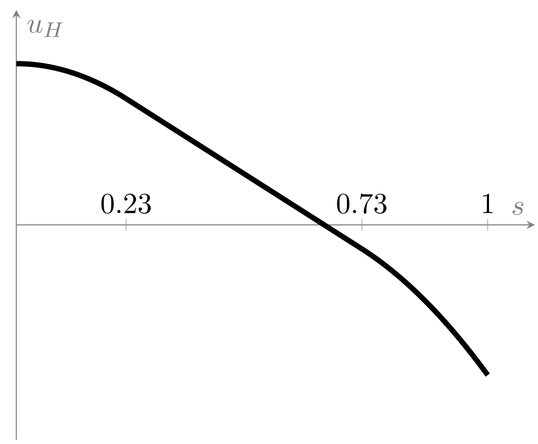


Figure 5: The function $u_{H^*}$ for the example's unique equilibrium.

In the unique equilibrium of this example, B chooses the Pareto signal, $G_{\pi^*,t^*}$ with $\pi^* \approx 0.23$ and $t^* \approx 0.73$, and S's randomization, $H^*$, is uniform over the support of $G_{\pi^*,t^*}$. We now confirm this strategy profile is indeed an equilibrium. First note the uniform randomization of S is profit maximizing because each price in the support of $G_{\pi^*,t^*}$ generates revenue $\pi^*$, and prices outside of the support generate smaller revenues. It remains to argue that $G_{\pi^*,t^*}$ is a best response against $H^*$. Figure 5 plots the integrand, $u_{H^*}$. As illustrated in the figure, the function $u_{H^*}$ is globally concave and affine on $[\pi^*, t^*]$. As mentioned above, since $u_{H^*}$ is concave, B's objective increases in mean-preserving contractions, so acquiring the uninformative signal, $\mathbf{1}_{[0.5,1]}$, is a best response to $H^*$. Since $u_{H^*}$ is affine on supp $G_{\pi^*,t^*} = [\pi^*, t^*]$, $G_{\pi^*,t^*}$ generates the same utility for B as $\mathbf{1}_{[0.5,1]}$, and hence, it is also a best response.

## 4.2 Vanishing Learning Cost

We are now ready to state and prove the main result of the paper: As the cost of learning vanishes, equilibria converge to a free-learning equilibrium that minimizes both players' payoffs. In this equilibrium, S achieves only his minmax profit, $\underline{\pi} = \min_{F\in\mathcal{A}} \pi_F$, and B uses the Pareto signal associated with this profit, $G_{\underline{\pi},\bar{t}}$.[32]

**Theorem 2** *For $\kappa > 0$, let $(H_\kappa, F_\kappa)$ be any equilibrium of the $\kappa$-game. Then,*

$$\lim_{\kappa\to 0}(H_\kappa, F_\kappa) = (\mathbf{1}_{[\bar{p}_{\underline{\pi}},1]}, G_{\underline{\pi},\bar{t}}).$$

[32] Roesler and Szentes (2017) establish the existence and uniqueness of such a Pareto signal.

Recall that $\bar{p}_{\underline{\pi}}$ is the largest price that generates profit $\underline{\pi}$ when B learns perfectly. Therefore, this theorem says that in the limit as learning becomes free, B uses a Pareto signal that generates S's minmax profit, and S charges the higher of the two prices yielding this profit when B collects full information. By Corollary 1, this limit is the worst free-learning equilibrium for both players.

Let us explain the argument of the proof. The main idea is to connect our analysis of costly learning with our observations regarding free-learning equilibria. When costs are positive, B uses a Pareto signal (see Proposition 1). Since the set of Pareto signals is closed, she must also be using a Pareto signal in the limit, say, $G_{\pi,t}$. In turn, equilibrium dictates that S sets a price in the support of $G_{\pi,t}$ (see part (ii) of Lemma 1). The key step in the proof, which we explain in detail in the next paragraph, is to show that against such a price, the Pareto signal associated with the minmax profit is strictly better for B than any other Pareto signal. Therefore, B cannot possibly use a Pareto signal other than $G_{\underline{\pi},\bar{t}}$ in a free-learning equilibrium. To conclude the theorem, we note that if S's profit is $\underline{\pi}$, he must charge $\bar{p}_{\underline{\pi}}$ by Lemma 5.

We now return to the key step of the proof, and explain why $G_{\underline{\pi},\bar{t}}$ is a profitable deviation from any other Pareto signal whose support contains S's price. Thus, fix any $G_{\pi,t}$ that differs from $G_{\underline{\pi},\bar{t}}$, and consider any price $p$ in supp $G_{\pi,t}$. By equation (4), $G_{\underline{\pi},\bar{t}}$ is strictly better for B against $p$ than $G_{\pi,t}$ if and only if $I_{G_{\underline{\pi},\bar{t}}}(p) < I_{G_{\pi,t}}(p)$. This inequality is equivalent to $\int_0^p (G_{\underline{\pi},\bar{t}} - G_{\pi,t})\ \mathrm{d}s$ being strictly positive. In other words, we need to show that the area to the left of any $p \in [\pi, t]$ between $G_{\pi,t}$ and $G_{\underline{\pi},\bar{t}}$ is strictly positive. As can be seen in Figure 6, the integral $\int_0^p (G_{\underline{\pi},\bar{t}} - G_{\pi,t})\ \mathrm{d}s$ is zero for $p \leq \underline{\pi}$ and strictly increasing over $[\underline{\pi}, t]$. Since $\pi > \underline{\pi}$, it follows that the integral is strictly positive for any $p \in [\pi, t]$ and so against any such price, $G_{\underline{\pi},\bar{t}}$ generates a strictly higher surplus for B than $G_{\pi,t}$.

Theorem 2 characterizes the unique equilibrium in the limit as $\kappa$ converges to zero but it provides little information regarding the equilibrium strategies along such sequences. In what follows, we explain how the randomization of S and the signal of B changes as $\kappa$ becomes smaller and smaller in the context of our example.

**Example** (continued). Recall that the prior, $F_0$, is uniform on $[0, 1]$ and that the cost of each signal, $F$, is given by its variance, $\kappa C(F) = \kappa \int (s^2 - \bar{v}^2)\ \mathrm{d}F$.

By Proposition 1, for each $\kappa$, the equilibrium signal of B is a Pareto one, $G_{\pi^*,t^*}$, and the support of S's equilibrium randomization, $H^*$, is $[\pi^*, t^*]$. One can show that, in this

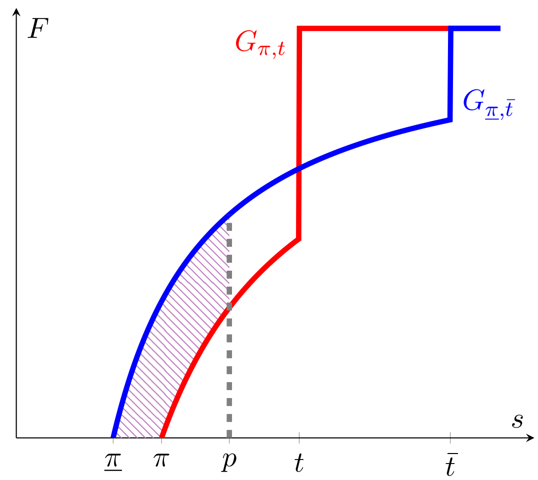


Figure 6: The minmax Pareto signal, $G_{\underline{\pi},\bar{t}}$, and another Pareto signal, $G_{\pi,t}$, with $\pi > \underline{\pi}$.

example, the equilibrium strategies are unique for each $\kappa$.[33] Next, we describe $G_{\pi^*,t^*}$ and $H^*$ as a function of $\kappa$ and illustrate them in Figure 7.

It turns out that there is a threshold value of $\kappa$, $\bar{\kappa} \approx 0.75$, above which $H^*$ is uniform on $[\pi^*, t^*]$. Moreover, when $\kappa > \bar{\kappa}$, a decrease in $\kappa$ results in an increase of the support of B's signal, that is, a smaller $\pi^*$ and a larger $t^*$. Panel A depicts the equilibrium strategies for the case when $\kappa = 1$, discussed at the end of section 4.1. At the threshold, $\bar{\kappa}$, the equilibrium profit is the minmax profit, $\underline{\pi} \approx 0.2$. As illustrated in Panel B, if $\kappa = \bar{\kappa}$, B's signal is $G_{\underline{\pi},\bar{t}}$ and $H^*$ is still uniform on $[\underline{\pi}, \bar{t}]$. When $\kappa$ is below the threshold, $\bar{\kappa}$, B's signal remains $G_{\underline{\pi},\bar{t}}$, but the randomization of S, $H^*$, is a combination of the uniform distribution on $[\underline{\pi}, \bar{t}]$ (with probability $\kappa/\bar{\kappa}$) and an atom at $\bar{p}_{\underline{\pi}} \approx 0.71$ (with probability $1 - \kappa/\bar{\kappa}$). Such an equilibrium is plotted on Panel C for $\kappa = 0.37$. As $\kappa$ converges to zero, the probability of S offering $\bar{p}_{\underline{\pi}}$ goes to one, and equilibria converge to the Pareto-worst

[33] To sketch the argument for uniqueness, suppose $(H, G_{\pi,t})$ is a $\kappa$-equilibrium. Since $G_{\pi,t}(x) \neq F_0(x)$ for any $x \in \{\pi, t\}$, Lemma 3 implies $\pi, t \notin S(G_{\pi,t})$. Using Dworczak and Martini's (2019) Theorem 2, one can show $G_{\pi,t}$ is optimal for B only if $u_H(s) = \int_{p \leq s}(s-p)\, \mathrm{d}H(p) - \kappa(s^2 - v^2)$ is locally concave at $\pi$ and $t$, which implies $H$ is continuous at both points. The same theorem also implies $u_H$ must be affine on any interval $[x,y] \subseteq [\pi, t]$ over which $I_{G_{\pi,t}}$ is strictly positive. A simple derivation reveals $u_H$ is affine on $[x,y]$ only if $H$ admits $h(z) = 2\kappa z$ as a density on $(x,y)$. Since $P(G_{\pi,t}) = [\pi, t]$, Lemmas 2 and 5 imply $S(G_{\pi,t}) \cap [\pi, t] \subseteq \{\bar{p}_\pi\}$, and so $1 = H(t) - H(\pi) \geq H(t) - (H(\bar{p}_\pi) - H(\bar{p}_\pi -)) - H(\pi) = 2\kappa(t - \pi)$, where the inequality holds with equality whenever $S(G_{\pi,t}) \cap [\pi, t] = \varnothing$—which is true for all $\pi > \underline{\pi}$ (see proof of Theorem 2). One can then prove at most one $\kappa$-equilibrium exists for any $\kappa > 0$ by using the fact that $G_{\pi',t'}, G_{\pi'',t''} \in \mathcal{A}$ for $\pi' < \pi''$ only if $t' > t''$.

free-learning equilibrium, $(\mathbf{1}_{[\bar{p}_{\underline{\pi}},1]}, G_{\underline{\pi},\bar{t}})$, depicted in Figure 7's Panel D.

## 5 Discussion

To conclude, we discuss some of our assumptions and how they can be relaxed.

*Production costs.* We assumed S's production cost is zero. We now discuss how our results generalize to the case in which S has to incur a positive production cost upon trade. Thus, suppose S's payoff when trading is $p - m$, where $m \in (0, 1)$. For $m \in (0, \bar{v})$, our analysis goes through with the $m$-**shifted truncated Pareto** signal,

$$\hat{G}^m_{\pi,t}(s) = \mathbf{1}_{[\pi+m,t)}\left(1 - \frac{\pi}{s-m}\right) + \mathbf{1}_{[t,1]} \qquad t \geq \pi + m,\ \pi \geq 0,$$

replacing the truncated Pareto, $G_{\pi,t}$. Other than this replacement, all results hold as stated.

For $m \geq \bar{v}$, our analysis implies trade breaks down: in the costless limit, B collects no information and no trade occurs. To see why, note that even when $m > 0$, Proposition 1's part (i) continues to hold for any costly learning equilibrium in which B acquires information. In other words, in any costly learning equilibrium in which B learns, the support of S's price and of B's signal must equal the same interval. As such, if B's signal is non-degenerate, its CDF is an $m$-shifted truncated Pareto. But when $m \geq \bar{v}$, no informative signal can have an $m$-shifted truncated Pareto distribution.[34] Hence, B acquires no information when learning is costly, and so the same must hold in the costless limit. However, if $p < 1$ and learning is free, full information strictly benefits B over no information. Thus, the vanishing-cost limit is autarky with no learning.

*Robustness and purification: random production costs.* Our main result appears to rely on the observation that if information is free, B learns whether her valuation is above or below the equilibrium price but chooses to ignore large amounts of information. If many equilibrium prices were possible, B may need to learn more and compare her valuation with any of these prices. Therefore, one may wonder whether our results extend to environments where the price is stochastic. Another concern is that when learning is costly, S randomizes in equilibrium, and it is not obvious that S's strategy can be purified without affecting our main conclusion. To address these issues, we now describe what

[34] For an explanation, suppose $F = \hat{G}^m_{\pi,t}$ for some signal $F \in \mathcal{A}$. Then, supp $F$ = supp $\hat{G}^m_{\pi,t} \subseteq [m, 1] \subseteq [\bar{v}, 1]$. Therefore, $\int s \, \mathrm{d}F \geq \bar{v}$, with equality only if supp $F = \{\bar{v}\}$, that is, if $F$ is uninformative.

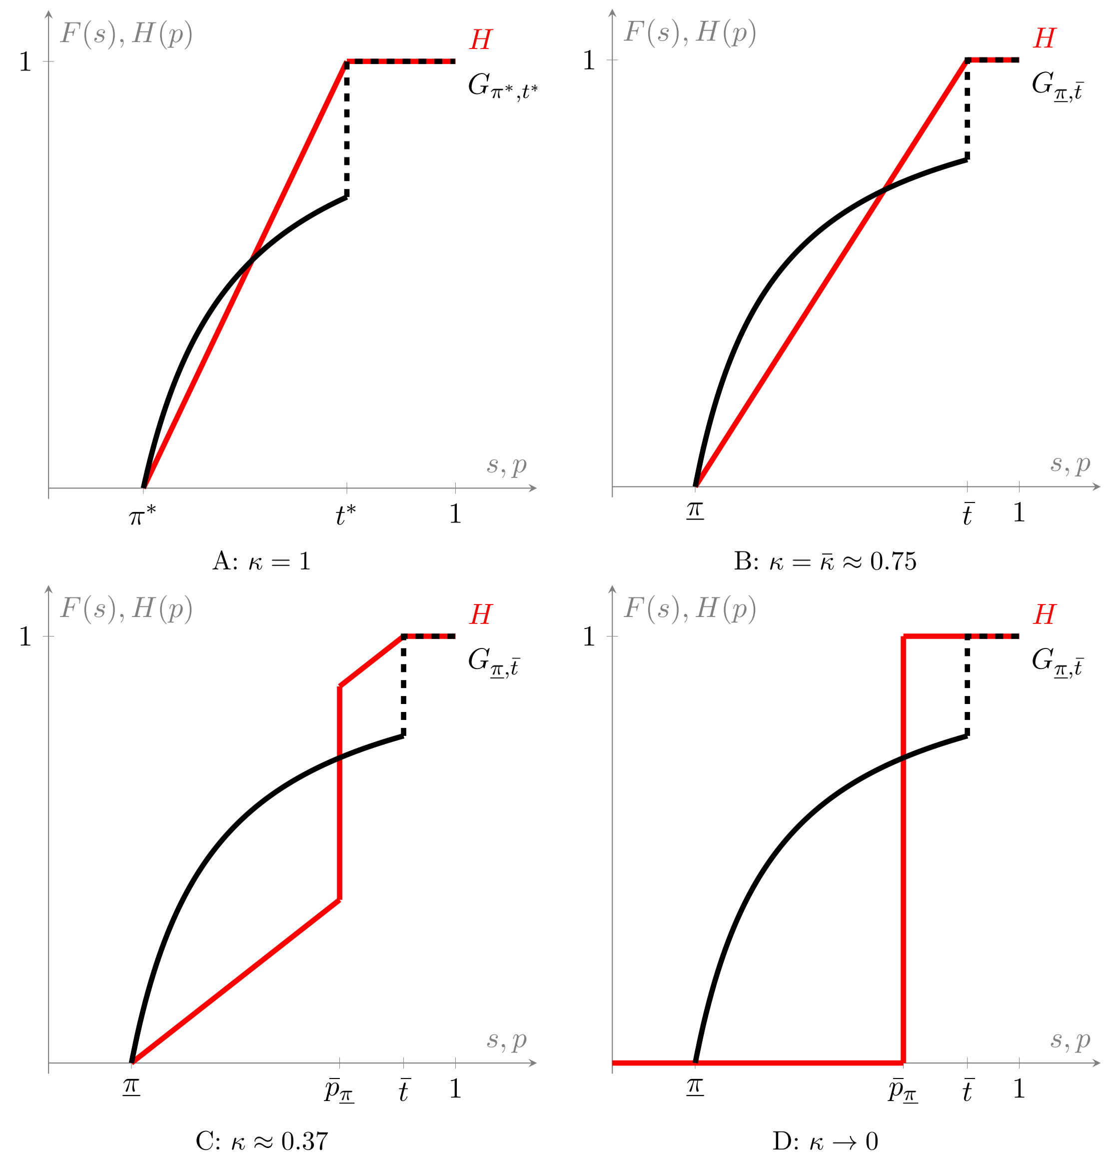


Figure 7: The example's equilibrium $(G_{\pi^*,t^*}, H^*)$ for $\kappa = 1$ (Panel A), $\bar{\kappa} \approx 0.75$ (Panel B), $\kappa \approx 0.37$ (Panel C) and $\kappa \approx 0$ (Panel D).

happens if S has a random production cost with full support in $[0, 1]$ that is independent of B's valuation[35]—see the online appendix for the formal details. S privately observes the cost realization, $m$, before setting a price. Then, his utility from trade at price $p$ is $p - m$, where $m$ is the production-cost realization. In this case, free-learning equilibria are still strongly Pareto ranked and are indexed by the price S charges when $m = 0$. This price is offered for all values of $m$ for which S would set a lower price under perfect learning, and B's signal distribution above this price agrees with the CDF of her prior. For higher values of $m$, S sets the same price as he would under perfect learning. Both players turn out to strictly prefer equilibria in which the price is lower conditional on $m = 0$. Because this price must be separating in equilibrium, its maximum across all B signals is $\bar{p}_{\underline{\pi}}$, whereas its minimum is attained when B learns perfectly. As such, perfect learning is still a Pareto-best equilibrium. In the Pareto-worst equilibrium, the CDF of B's signal coincides with the truncated Pareto, $G_{\underline{\pi},\bar{t}}$, for all values below $\bar{p}_{\underline{\pi}}$. One can show this free-learning equilibrium is the only one in which B uses this CDF, and that the same CDF is attained at the vanishing-cost limit.[36] Hence, even when the production cost is stochastic, our main result is valid and the costless limit still selects the Pareto-worst free-learning equilibrium.

*Random prices as general mechanisms.* We argue that it is without loss of generality for S to set a price instead of a more general mechanism. Consider a more general model, where S and B simultaneously choose a mechanism and a signal, respectively. Then, B observes her signal's realization and decides whether to participate in S's mechanism. A mechanism constitutes a set of messages for B, and each message is associated with a transfer and a probability of trade. Note B's interim expected payoff from any of the messages is fully determined by her posterior-value estimate. Hence, by the Revelation Principle, restricting attention to individually rational and incentive-compatible mechanisms in which B truthfully reports her posterior-value estimate is without loss. Then, standard arguments imply any mechanism is equivalent to setting a random price; see,

[35]A related extension is one in which S's production costs are correlated with B's valuation. In this case, S's price may provide B with information about her own value. Such signaling complicates the analysis in two ways. First, it introduces a large equilibrium multiplicity due to the indeterminacy of off-path beliefs. And second, B's inference from S's price will generally depend on other moments of her posterior in addition to the mean, and so the marginal distribution of B's posterior mean is no longer sufficient for characterizing trade outcomes. Hence, our analysis is not applicable to such extensions.

[36]In the vanishing cost limit, prices have a full support distribution over the interval $[\bar{p}_{\underline{\pi}}, 1]$, with an atom on $\bar{p}_{\underline{\pi}}$.

for example, Börgers (2015), Proposition 2.5.

*Non-regular prior.* Most of our results generalize to the case in which B's prior-value distribution is not regular (including the possibility that $F_0$ has atoms).[37] When learning is free, equilibrium requires S's price to be separating, and the full-information outcome remains profit maximizing regardless of the prior. Our construction of free-learning equilibrium for each profit level between $\underline{\pi}$ and $\pi_{F_0}$ is also valid for non-regular priors. Regularity of the prior also plays no role in showing B uses a Pareto signal when learning is costly, and the same holds in the costless limit. Because the costless limit is a free-learning equilibrium, the Pareto signal in the limiting case still has a separating price in its support, so this signal is still profit minimizing. Therefore, even without regularity, the costless limit still minimizes S's profits across all signal structures and generates the lowest profit across all free-learning equilibria.

However, a non-regular prior does affect the conclusion that the costless limit minimizes B's payoff, for two reasons. First, a non-regular prior can result in Pareto-incomparable free-learning equilibria, and so the profit-minimizing equilibrium may not minimize B's payoff. Second, when the prior is non-regular, the profit-minimizing Pareto signal may have more than one separating price in its support, so many free-learning equilibria may exist in which B uses the profit-minimizing Pareto signal. In fact, one can show that under Assumption 1, each such equilibrium is a limit of some equilibrium sequence with vanishing costs. As a consequence, without regularity, B may obtain different outcomes in the vanishing-cost limit depending on the fine details of the prior and the converging equilibrium sequence.

*Non-smooth learning costs.* To simplify exposition, we assumed B's cost is a Fréchet differentiable function satisfying Assumption 1. These assumptions, however, are stronger than necessary. Specifically, we show in the online appendix that our results extend to any convex cost function as long as it satisfies three properties. First, the function is lower semicontinuous. Second, full information can be approximated at finite cost. Finally, one can approximate the pooling of any two signal realizations in a way that strictly reduces costs at the margin.[38] We verify these properties for any continuous posterior-separable

[37]Our results hold without change when the buyer's prior is absolutely continuous and supported on a subinterval, $[\underline{x}, \bar{x}] \subseteq [0, 1]$, over which $f_0$ is strictly positive and $v - (1 - F_0(v))/f_0(v)$ is strictly increasing. Whenever $\underline{x} > 0$, it is possible, however, that $\underline{\pi} = \pi_{F_0}$, meaning a unique free-learning equilibrium exists. Such uniqueness arises if and only if $\pi_{F_0} = \underline{x}$, which is impossible when $\underline{x} = 0$.

[38]The described property corresponds to Condition 1 stated in the online appendix. The appendix also

cost function (Caplin et al., 2017) in the Appendix. Thus, our results apply to a much wider class of information-cost functions than might appear at first.

## References

Aliprantis, Charalambos D and Kim Border (2006), *Infinite Dimensional Analysis: A Hitchhiker's Guide*. Springer Science & Business Media.

Aumann, Robert J and M Maschler (1966), "Game Theoretic Aspects of Gradual Disarmament." *Report of the US Arms Control and Disarmament Agency*, 80, 1–55.

Benoît, Jean-Pierre and Juan Dubra (2011), "Apparent Overconfidence." *Econometrica*, 79, 1591–1625.

Bergemann, Dirk, Benjamin Brooks, and Stephen Morris (2015), "The Limits of Price Discrimination." *American Economic Review*, 105, 921–57.

Bergemann, Dirk and Stephen Morris (2013), "Robust Predictions in Games with Incomplete Information." *Econometrica*, 81, 1251–1308.

Bergemann, Dirk and Martin Pesendorfer (2007), "Information Structures in Optimal Auctions." *Journal of Economic Theory*, 137, 580–609.

Bergemann, Dirk and Karl H Schlag (2008), "Pricing without Priors." *Journal of the European Economic Association*, 6, 560–569.

Bessembinder, Hendrik and William Maxwell (2008), "Markets: Transparency and the Corporate Bond Market." *Journal of Economic Perspectives*, 22, 217–234.

Börgers, Tilman (2015), *An Introduction to the Theory of Mechanism Design*. Oxford University Press, USA.

Brøndsted, Arne and Ralph Tyrrell Rockafellar (1965), "On the Subdifferentiability of Convex Functions." *Proceedings of the American Mathematical Society*, 16, 605–611.

Caplin, Andrew, Mark Dean, and John Leahy (2017), "Rationally Inattentive Behavior: Characterizing and Generalizing Shannon Entropy." Technical report, National Bureau of Economic Research.

contains a more general sufficient condition, namely that the subdifferential of $C$ only contains functions satisfying a convex inequality.

Carrasco, Vinicius, Vitor Farinha Luz, Nenad Kos, Matthias Messner, Paulo Monteiro, and Humberto Moreira (2018), "Optimal Selling Mechanisms under Moment Conditions." *Journal of Economic Theory*, 177, 245–279.

Cerreia-Vioglio, Simone, Fabio Maccheroni, and Massimo Marinacci (2017), "Stochastic Dominance Analysis Without the Independence Axiom." *Management Science*, 63, 1097–1109.

Chatterjee, Kalyan and R Vijay Krishna (2011), "A Nonsmooth Approach to Nonexpected Utility Theory Under Risk." *Mathematical Social Sciences*, 62, 166–175.

Chew, Soo Hong, Larry G Epstein, and Uzi Segal (1991), "Mixture Symmetry and Quadratic Utility." *Econometrica: Journal of the Econometric Society*, 139–163.

Compte, Olivier and Philippe Jehiel (2007), "Auctions and Information Acquisition: Sealed Bid or Dynamic Formats?" *The Rand Journal of Economics*, 38, 355–372.

Denti, Tommaso (2019), "Unrestricted Information Acquisition." *Working Paper*.

Du, Songzi (2018), "Robust Mechanisms under Common Valuation." *Econometrica*, 86, 1569–1588.

Dworczak, Piotr and Giorgio Martini (2019), "The Simple Economics of Optimal Persuasion." *Journal of Political Economy*, 127, 1993–2048.

Ganuza, Juan-José (2004), "Ignorance Promotes Competition: An Auction Model with Endogenous Private Valuations." *Rand Journal of Economics*, 583–598.

Ganuza, Juan-José and Jose S Penalva (2010), "Signal Orderings Based on Dispersion and the Supply of Private Information in Auctions." *Econometrica*, 78, 1007–1030.

Gentzkow, Matthew and Emir Kamenica (2014), "Costly Persuasion." *American Economic Review*, 104, 457–62.

Gentzkow, Matthew and Emir Kamenica (2016), "A Rothschild-Stiglitz Approach to Bayesian Persuasion." *American Economic Review*, 106, 597–601.

Grubb, Michael D (2015), "Failing to Choose the Best Price: Theory, Evidence, and Policy." *Review of Industrial Organization*, 47, 303–340.

Handel, Benjamin and Joshua Schwartzstein (2018), "Frictions or Mental Gaps: What's Behind the Information We (Don't) Use and When Do We Care?" *Journal of Economic Perspectives*, 32, 155–78, URL `https://www.aeaweb.org/articles?id=10.1257/jep.32.1.155`.

Hong, Chew Soo, Edi Karni, and Zvi Safra (1987), "Risk Aversion in the Theory of Expected Utility with Rank Dependent Probabilities." *Journal of Economic theory*, 42, 370–381.

Hoshino, Tetsuya (2018), "Rational Inattention in Games." *Working Paper*.

Kamenica, Emir and Matthew Gentzkow (2011), "Bayesian Persuasion." *American Economic Review*, 101, 2590–2615.

Kartik, Navin and Weijie Zhong (2019), "Lemonade from Lemons: Information Design and Adverse Selection."

Li, Hao and Xianwen Shi (2017), "Discriminatory Information Disclosure." *American Economic Review*, 107, 3363–85.

Machina, Mark J (1982), ""Expected Utility" Analysis without the Independence Axiom." *Econometrica: Journal of the Econometric Society*, 277–323.

Matějka, Filip (2015), "Rigid Pricing and Rationally Inattentive Consumer." *Journal of Economic Theory*, 158, 656–678.

Milgrom, Paul R and Robert J Weber (1982), "A Theory of Auctions and Competitive Bidding." *Econometrica: Journal of the Econometric Society*, 1089–1122.

Milgrom, Paul R and Robert J Weber (1985), "Distributional Strategies for Games with Incomplete Information." *Mathematics of operations research*, 10, 619–632.

Morris, Stephen and Ming Yang (2019), "Coordination and Continuous Stochastic Choice." *Working Paper*.

O'Hara, Maureen, Yihui Wang, and Xing Alex Zhou (2018), "The Execution Quality of Corporate Bonds." *Journal of Financial Economics*, 130, 308–326.

Persico, Nicola (2000), "Information Acquisition in Auctions." *Econometrica*, 68, 135–148.

Ravid, Doron (2019), "Bargaining with Rational Inattention." *Available at SSRN 2957890*.

Reny, Philip J (1999), "On the Existence of Pure and Mixed Strategy Nash Equilibria in Discontinuous Games." *Econometrica*, 67, 1029–1056.

Roesler, Anne-Katrin and Balazs Szentes (2017), "Buyer-Optimal Learning and Monopoly Pricing." *American Economic Review*, 107, 2072–2080.

Shi, Xianwen (2012), "Optimal Auctions with Information Acquisition." *Games and Economic Behavior*, 74, 666–686.

Sims, Christopher A (1998), "Stickiness." In *Carnegie-Rochester Conference Series on Public Policy*, volume 49, 317–356, Elsevier.

Sims, Christopher A (2003), "Implications of Rational Inattention." *Journal of Monetary Economics*, 50, 665–690.

Sims, Christopher A (2006), "Rational Inattention: Beyond the Linear-Quadratic Case." *American Economic Review*, 96, 158–163.

Yang, Kai Hao (2019), "Buyer-Optimal Information with Nonlinear Technology." *Available at SSRN 3306455*.

Yang, Ming (2015), "Coordination with Flexible Information Acquisition." *Journal of Economic Theory*, 158, 721–738.

Yang, Ming (2020), "Optimality of Debt Under Flexible Information Acquisition." *The Review of Economic Studies*, 87, 487–536.

Zhu, Haoxiang (2012), "Finding a Good Price in Opaque Over-the-Counter Markets." *The Review of Financial Studies*, 25, 1255–1285.

# Appendix

## A Proof of Claim 1

We begin by proving the following useful lemma, which shows that for every $F$, $w, z \in \text{int}\,(\text{co}(\text{supp } F))$, and $\alpha \in (0,1)$, two distributions, $F'$, $F''$, exist such that $F \succeq F' \succ F''$ and

$$F' - F'' = \gamma\left(\alpha \mathbf{1}_{[w,1]} + (1-\alpha)\mathbf{1}_{[z,1]} - \mathbf{1}_{[\alpha w + (1-\alpha)z,1]}\right)$$

for some $\gamma > 0$.

**Lemma 6** *Fix some $F \in \mathcal{F} \setminus \{\mathbf{1}_{[x,1]} : x \in [0,1]\}$, let $[x, x'] = \text{co}(\text{supp } F)$, and take $\bar{w} = \int s \, \text{d}F$. Take any $w, y, z \in (x, x')$, and $\alpha \in (0,1)$ such that $y = \alpha w + (1-\alpha)z$. For $\lambda, \beta \in [0,1)$, define $x_\lambda = \frac{\bar{w} - \lambda y}{1-\lambda}$, and*

$$F_{\lambda,\beta} := (1-\lambda)\mathbf{1}_{[x_\lambda,1]} + \lambda(1-\beta)\mathbf{1}_{[y,1]} + \lambda\beta\left[\alpha\mathbf{1}_{[w,1]} + (1-\alpha)\mathbf{1}_{[z,1]}\right].$$

*Then, $\beta, \lambda \in (0,1)$ exists such that $F \succeq F_{\lambda,\beta} \succ F_{\lambda,0}$.*

**Proof.** Suppose without loss that $z > w$. Note $F_{\lambda,0} \succeq \mathbf{1}_{[\bar{w},1]}$ for all $\lambda > 0$ because $\lambda y + (1-\lambda)x_\lambda = \bar{w}$. We now show $F_{\lambda,\beta} \succeq F_{\lambda,0}$ for every $\beta \geq 0$. For this purpose, notice that

$$F_{\lambda,\beta} - F_{\lambda,0} = \lambda\beta[\alpha\mathbf{1}_{[w,1]} + (1-\alpha)\mathbf{1}_{[z,1]}] - \lambda\beta\mathbf{1}_{[y,1]}.$$

Therefore, for all $\bar{s} \in [0,1]$,

$$\int_0^{\bar{s}} (F_{\lambda,\beta} - F_{\lambda,0}) \, \text{d}s = \lambda\beta \int_0^{\bar{s}} (\alpha\mathbf{1}_{[w,1]} + (1-\alpha)\mathbf{1}_{[z,1]} - \mathbf{1}_{[y,1]}) \, \text{d}s \geq 0,$$

in view of $\left(\alpha\mathbf{1}_{[w,1]} + (1-\alpha)\mathbf{1}_{[z,1]}\right) \succeq \mathbf{1}_{[y,1]}$. Because $\bar{s}$ was arbitrary, we have $F_{\lambda,\beta} \succeq F_{\lambda,0}$.

Let us introduce some helpful definitions, which rely on $x_\lambda$ being continuous in $\lambda$ and $x_0 = \bar{w}$. Fixing some $\epsilon > 0$ for which $(\bar{w} - \epsilon, \bar{w} + \epsilon) \subseteq (x, x')$, choose a $\bar{\lambda}$ to be such that $\{x_\lambda\}_{\lambda \in [0,\bar{\lambda}]} \subseteq (\bar{w} - \epsilon, \bar{w} + \epsilon) \subseteq (x, x')$. Let $x^* = \max\left(\{z\} \cup \{x_\lambda\}_{\lambda \in [0,\bar{\lambda}]}\right)$ and $x_* = \min\left(\{w\} \cup \{x_\lambda\}_{\lambda \in [0,\bar{\lambda}]}\right)$, and define the function

$$\begin{aligned} \varphi : [x_*, x^*] \times [0, \bar{\lambda}]^2 &\to \mathbb{R} \\ (\bar{s}, \lambda, \beta) &\mapsto \int_0^{\bar{s}} (F - F_{\lambda,\beta}) \, \text{d}s. \end{aligned}$$

Taking $(\cdot)_+ := \max\{\cdot, 0\}$, we can write

$$\varphi(\bar{s}, \lambda, \beta) = \int_0^{\bar{s}} F \, ds - (1-\lambda)(\bar{s} - x_\lambda)_+ - \lambda(1-\beta)(\bar{s} - y)_+ \\ - \lambda\beta\alpha(\bar{s} - w)_+ - \lambda\beta(1-\alpha)(\bar{s} - z)_+,$$

and so $\varphi$ is continuous in the product topology. Therefore,

$$\varphi^* : [0, \bar{\lambda}]^2 \to \mathbb{R} \\ (\lambda, \beta) \mapsto \min_{s \in [x_*, x^*]} \varphi(s, \lambda, \beta)$$

is also continuous by Berge's Maximum Theorem.

We now show $\varphi(\bar{s}, 0, 0) > 0$ for all $\bar{s} \in [x_*, x^*]$. To do so, notice $x_0 = \bar{w}$, and therefore, $F_{0,0} = \mathbf{1}_{[x_0,1]} = \mathbf{1}_{[\bar{w},1]}$. Because $\bar{w} > x_* > x$ (by choice of $F$), we also have $F(s) > 0 = \mathbf{1}_{[\bar{w},1]}(s)$ for all $s \in [x, \bar{w})$. As such, if $\bar{s} \in [x_*, \bar{w}]$ then $\int_0^{\bar{s}}(F - \mathbf{1}_{[\bar{w},1]})(s) \, ds = \int_x^{\bar{s}} F(s) \, ds > 0$. Similarly, for all $s \in [\bar{w}, x')$, $F(s) < 1 = \mathbf{1}_{[\bar{w},1]}(s)$. As such, if $\bar{s} \in [\bar{w}, x^*]$, $\int_{\bar{s}}^1 (1 - F(s)) \, ds > 0 = \int_{\bar{s}}^1 (1 - \mathbf{1}_{[\bar{w},1]}(s)) \, ds$, and so $\int_{\bar{s}}^1 (F - \mathbf{1}_{[\bar{w},1]})(s) \, ds < 0$. Since $\int_0^1 (F - \mathbf{1}_{[\bar{w},1]})(s) \, ds = 0$, we obtain $\int_0^{\bar{s}}(F - \mathbf{1}_{[\bar{w},1]})(s) \, ds > 0$ for all $\bar{s} \in [\bar{w}, x^*]$ as well.

We are now in a position to complete the proof; that is, we show $F \succeq F_{\lambda,\beta}$ for all small $\lambda, \beta > 0$. By the previous paragraph, $\varphi(\bar{s}, 0, 0) > 0$ for all $\bar{s} \in [x_*, x^*]$. As such, $\varphi^*(0,0) = \min_{s \in [x_*, x^*]} \varphi(s, 0, 0) > 0$, and so by continuity of $\varphi^*$, one must then have $\varphi^*(\lambda, \beta) > 0$ for all $\lambda, \beta > 0$ small enough. Fixing any such $\lambda$ and $\beta$, we now show $\int_0^{\bar{s}}(F - F_{\lambda,\beta}) \, ds \geq 0$ for all $\bar{s}$ by considering three cases. First, if $\bar{s} \in [x_*, x^*]$,

$$\int_0^{\bar{s}} (F - F_{\lambda,\beta}) \, ds \geq \varphi^*(\lambda, \beta) > 0.$$

Second, if $\bar{s} \in [x, x_*)$, $F(x) \geq 0 = F_{\lambda,\beta}(x)$, and so $\int_0^{\bar{s}}(F - F_{\lambda,\beta}) \, ds = \int_0^{\bar{s}} F \, ds \geq 0$. Third, if $\bar{s} \in (x^*, 1]$,

$$\int_0^{\bar{s}} (F - F_{\lambda,\beta}) \, ds = \int_0^{x^*} (F - F_{\lambda,\beta}) \, ds + \int_{x^*}^{\bar{s}} (F - 1) \, ds \\ \geq \int_0^{x^*} (F - F_{\lambda,\beta}) \, ds + \int_{x^*}^{1} (F - 1) \, ds = \int_0^1 (F - F_{\lambda,\beta}) \, ds = 0,$$

in view of supp $F_{\lambda,\beta} \subseteq [x_*, x^*]$ and $F_{\lambda,\beta} \succeq \mathbf{1}_{[\bar{w},1]}$. We have therefore shown that for all sufficiently small $\lambda$ and $\beta$, $\int_0^{\bar{s}}(F - F_{\lambda,\beta}) \, ds \geq 0$ for all $\bar{s} \in [0,1]$, with equality holding at $\bar{s} = 1$ (because $F_{\lambda,\beta} \succeq \mathbf{1}_{[\bar{w},1]}$). Therefore, $F \succeq F_{\lambda,\beta}$, thereby completing the proof. ■

We are now ready to prove Claim 1.

**Proof of Claim 1.** Suppose first that $c_F$ is convex for all $F$. Fix some $F' \succeq F$. Define the function

$$\begin{aligned} \psi : [0,1] &\to \mathbb{R} \\ \alpha &\mapsto C(F + \alpha(F' - F)). \end{aligned}$$

Because $C$ is Fréchet differentiable, $\psi$ is a differentiable function whose derivative is given by

$$\psi'(\alpha) = \int c_{F+\alpha(F'-F)} \, \mathrm{d}(F' - F) \geq 0,$$

where the inequality follows from convexity of $c_{F+\alpha(F'-F)}$ and $F' \succeq F$. Applying the fundamental theorem of calculus then gives

$$C(F') = \psi(1) = C(F) + \int_0^1 \psi'(s) \, \mathrm{d}s \geq C(F).$$

Hence, $C$ is monotone.

Suppose now that $C$ is monotone. Fix any $w, y, z \in \text{co}(\text{supp } F_0)$ such that $y = \alpha w + (1-\alpha)z$ for some $\alpha \in (0,1)$. By continuity of $c_F$, it is sufficient to show $c_F(y) \leq \alpha c_F(w) + (1-\alpha)c_F(z)$ whenever $w, y, z \in \text{int}\,(\text{co}(\text{supp } F_0))$.

By Lemma 6, an $F'$ and $F''$ exist such that $F_0 \succeq F' \succ F''$ and

$$F' - F'' = \left(\alpha \mathbf{1}_{[w,1]} + (1-\alpha)\mathbf{1}_{[z,1]} - \mathbf{1}_{[\alpha w+(1-\alpha)z,1]}\right),$$

for some $\gamma > 0$. Because $\succeq$ respects convex combinations,

$$F + \epsilon(F' - F) \succeq F + \epsilon(F'' - F)$$

must hold for all $\epsilon \in [0,1]$. Appealing to monotonicity of $C$ then yields that, for all $\epsilon \in (0,1)$,

$$\begin{aligned} 0 &\leq C(F + \epsilon(F' - F)) - C(F + \epsilon(F'' - F)) \\ &= \left[C(F + \epsilon(F' - F)) - C(F)\right] - \left[C(F + \epsilon(F'' - F)) - C(F)\right]. \end{aligned}$$

Dividing by $\epsilon > 0$, taking $\epsilon \searrow 0$ and substituting for $F'$ and $F''$ then yields

$$\begin{aligned} 0 &\leq \frac{1}{\epsilon}\left[C(F + \epsilon(F' - F)) - C(F)\right] - \frac{1}{\epsilon}\left[C(F + \epsilon(F'' - F)) - C(F)\right] \\ &\to \int c_F \, \mathrm{d}(F' - F) - \int c_F \, \mathrm{d}(F'' - F) = \alpha c_F(w) + (1-\alpha)c_F(z) - c_F(y), \end{aligned}$$

thereby concluding the proof. ■

## B Upper Hemicontinuity of S's Best Response

In this section, we prove the following lemma about S's best-response correspondence and maximal value.

**Lemma 7** *S's maximal profit, $F \mapsto \pi_F$, is continuous, and $P(\cdot)$ is upper hemicontinuous.*

**Proof.** Let $\{F_n\}_{n\geq 0}$ be some sequence attaining $F_\infty$ as its limit. We show $\lim_{n\to\infty} \pi_{F_n} = \pi_{F_\infty}$. Because $\Pi$ is upper semicontinuous, $F \mapsto \pi_F$ is also upper semicontinuous.[39] As such, it suffices to show that $\liminf_{n\to\infty} \pi_{F_n} \geq \pi_\infty$. To do so, take any $p \in P(F_\infty)$. Then, for all $\epsilon > 0$,

$$\pi_{F_n} \geq \Pi(p-\epsilon, F_n) \geq (p-\epsilon)(1-F_n(p-\epsilon)).$$

Thus,

$$\liminf_n \pi_{F_n} \geq \liminf_n (p-\epsilon)(1-F_n(p-\epsilon)) \geq (p-\epsilon)(1-F_\infty(p-\epsilon)) \geq p(1-F_\infty(p-)) - \epsilon,$$

where the second inequality follows from the Portmanteau theorem. Because $\epsilon$ above is arbitrary, the result follows.

To see that $P(\cdot)$ is upper hemicontinuous, take any convergent sequence $p_n \in P(F_n)$ attaining $p_\infty$ as its limit. Because $\Pi$ is upper semicontinuous and $F \mapsto \pi_F$ is continuous,

$$\pi_{F_\infty} = \lim \pi_{F_n} = \limsup_n \Pi(p_n, F_n) \leq \Pi(p_\infty, F_\infty) \leq \pi_{F_\infty}.$$

Thus, $\Pi(p_\infty, F_\infty) = \pi_{F_\infty}$; that is, $p_\infty \in P(F_\infty)$. ■

## C Proof of Lemma 1

To prove part (i), note S's profit from setting a certain price cannot exceed $\pi_F$; that is, for all $s \in [0,1]$, $s\,(1 - F\,(s-)) \leq \pi_F$. Rearranging this inequality yields

$$G_{\pi_F,1}\,(s-) = 1 - \frac{\pi_F}{s} \leq F\,(s-)\,,$$

which proves part (i).

To see part (ii), note $s \in P\,(F)$ if and only if the inequality in the previous displayed chain is an equality. Hence, $P\,(F) = \{p \geq \pi_F : F(p-) = G_{\pi_F}(p-)\}$. It remains to show

[39] See Aliprantis and Border (2006), Lemma 17.30, for example.

that $P(F) \subseteq \text{supp } F$. Suppose, by contradiction, a $p$ exists such that $p \in P(F) \setminus \text{supp } F$. Then, $p' > p$ exists such that $F(p'-) = F(p-)$. Therefore,

$$\Pi(p, F) = p(1 - F(p-)) < p'(1 - F(p-)) = p'\left(1 - F(p'-)\right) = \Pi\left(p', F\right),$$

where the inequality follows from $p' > p$ and the second equality follows from $F(p'-) = F(p-)$. This inequality chain implies S is strictly better off setting price $p'$ than price $p$, a contradiction to $p \in P(F)$.

# D Proof of Lemma 2

If S uses $H$ and B chooses $F$, the difference between B's payoff generated by $F_0$ and that of $F$ can be written as

$$\begin{aligned} U_0(H, F_0) - U_0(H, F) &= \int \left[\int_p^1 (s-p) \ \mathrm{d}F_0(s) - \int_p^1 (s-p) \ \mathrm{d}F(s)\right] dH(p) \\ &= \int \left[\int_p^1 F(s)\, ds - \int_p^1 F_0(s) \ \mathrm{d}s\right] \ \mathrm{d}H(p) = \int I_F(p) \ \mathrm{d}H(p), \end{aligned}$$

where the first equality follows from (4) and the third one from $\int_p^1 (F - F_0) \ \mathrm{d}s = \int_0^p (F_0 - F) \ \mathrm{d}s = I_F(p)$. Because $F \in \mathcal{A}$ and $I_F(\cdot)$ is continuous, we conclude $F$ generates the same payoff as perfect learning if and only if $I_F(p) = 0$ for all $p \in \text{supp } H$; that is, if and only if $\text{supp } H \subseteq S(F)$.

# E Proof of Lemma 3

Suppose $p \in S(F)$. We begin by proving (i). By the definition of $S(F)$, $I_F(p) = 0$. Recall that $I_F(x) \geq 0$ for all $x \in [0, 1]$, so

$$p \in \arg\min_{x \in [0,1]} I_F(x). \tag{6}$$

Because $I_F(x) = \int_0^x (F_0 - F) \ \mathrm{d}s$, it can be differentiated from both sides at $p$. Therefore, (6) implies

$$\begin{aligned} 0 &\geq I'_{F-}(p) = F_0(p-) - F(p-), \\ 0 &\leq I'_{F+}(p) = F_0(p) - F(p). \end{aligned}$$

From these two inequalities, it follows that $F_0(p-) \leq F(p-) \leq F(p) \leq F_0(p)$. Because $F_0$ is regular, it does not have an atom at $p$, so $F_0(p-) = F(p)$. Hence, all the inequalities in the previous inequality chain are equalities. The lemma's part (i) follows.

Now, we prove part (ii). Suppose, by contradiction, that a $p' > p$ exists such that $F(p'') > F_0(p'')$ for all $p'' \in (p, p')$. Then,

$$0 \leq I_F(p') = \int_0^{p'} (F_0 - F)(s) \ \mathrm{d}s = I_F(p) + \int_p^{p'} (F_0 - F)(s) \ \mathrm{d}s = \int_p^{p'} (F_0 - F)(s) \ \mathrm{d}s < 0,$$

where the first inequality follows from $F \in \mathcal{A}$, the third equality from $p \in S(F)$, and the last inequality from $F(p'') > F_0(p'')$ for all $p'' \in (p, p')$.

# F Proof of Lemma 4

We show $P(F) \cap S(F) \subseteq [p_0^*, 1]$ for every $F \in \mathcal{A}$. To see why this inclusion is sufficient, recall that S is optimizing against B only if all his prices are profit maximizing, whereas $F$ is a best reply for B only if all of S's prices are $F$-separating (Lemma 2). In other words, supp $H \subseteq P(F) \cap S(F)$ for every free-learning equilibrium $(H, F)$.

Suppose for a contradiction that a $p \in P(F) \cap S(F)$ exists such that $p < p_0^*$. By part (ii) of Lemma 3, a $p' \in (\underline{p}, p_0^*)$ exists such that $F(p') \leq F_0(p')$. To obtain a contradiction, observe that

$$\Pi(p, F) \geq \Pi(p', F) \geq \Pi(p', F_0) > \Pi(p, F_0) = \Pi(p, F),$$

where the first inequality follows from $p$ being profit maximizing under $F$, the second inequality from $F(p'-) \leq F_0(p')$, the third inequality from strict quasiconcavity of $\Pi(\cdot, F_0)$ and $p' \in (p, p_0^*)$, and the equality from $p$ being $F$-separating and part (i) of Lemma 3.

# G Proof of Lemma 5

Recall that $(H, F)$ is a free-learning equilibrium only if all of S's prices are profit maximizing and $F$-separating, supp $H \subseteq P(F) \cap S(F)$. Therefore, arguing that $P(F) \cap S(F) = \{\bar{p}_{\pi_F}\}$ whenever $P(F) \cap S(F)$ is non-empty is sufficient. To see why, take any $p \in P(F) \cap S(F)$. Observe first that

$$\Pi(p, F_0) = p(1 - F_0(p-)) = p(1 - F(p-)) = \Pi(p, F) = \pi_F,$$

where the second equality follows from $p \in S(F)$ and part (i) of Lemma 3, and the last equality from $p \in P(F)$. It follows that $p \in X_{\pi_F}$.

To complete the proof, arguing that $p > \underline{p}$ for any $\underline{p} \in X_{\pi_F}$ such that $\underline{p} < \bar{p}_{\pi_F}$ is sufficient. Since $\Pi(\cdot, F_0)$ is strictly quasiconcave, such a $\underline{p}$ exists only if $X_{\pi_F} = \{\underline{p}, \bar{p}_{\pi_F}\}$ where $\underline{p} < p_0^* < \bar{p}_{\pi_F}$. Lemma 4 then delivers $p \geq p_0^* > \underline{p}$, as required.

## H Proof of Theorem 1: Free-learning equilibrium payoffs

We begin by noting that if $(H, F)$ is a free-learning equilibrium and $F_0$ is regular, B's expected utility is $\int_{\bar{p}_{\pi_F}}^{1}(v - \bar{p}_{\pi_F})\ \mathrm{d}F_0(v)$, which is a consequence of two facts. First, Lemma 5 implies $H$ puts a unit mass on $\bar{p}_{\pi_F}$; that is, $H = \mathbf{1}_{[\bar{p}_{\pi_F},1]}$. Second, full information is always optimal for B when learning is costless, meaning her expected utility in equilibrium must be the same as her expected utility under full information; that is, $U_0(\mathbf{1}_{[\bar{p}_{\pi_F},1]}, F) = U_0(\mathbf{1}_{[\bar{p}_{\pi_F},1]}, F_0) = \int_{\bar{p}_{\pi_F}}^{1}(s - \bar{p}_{\pi_F})\ \mathrm{d}F_0(s)$.

Given the above, it remains to be shown that a free-learning equilibrium, $(H, F)$, exists such that $\pi = \pi_F$ if and only if $\pi \in [\underline{\pi}, \pi_{F_0}]$. To do so, we first establish that $\underline{\pi} \leq \Pi(H, F) \leq \pi_{F_0}$ whenever $(H, F)$ is a free-learning equilibrium. Because $\underline{\pi} \leq \Pi(H, F)$ by definition of $\underline{\pi}$, it remains to be shown that $\Pi(H, F) \leq \pi_{F_0}$. To do so, notice that because supp $H \subseteq S(F)$, we have by Lemma 3 that $F(p-) \geq F_0(p-)$ for every $p \in$ supp $H$. Because $H$ maximizes S's profit, S's profit must be the same from all prices in supp $H$. We therefore have that for any $p \in$ supp $H$,

$$\Pi(H, F) = \Pi(p, F) = p(1 - F(p-)) \leq p(1 - F_0(p-)) = \Pi(p, F_0) \leq \pi_{F_0},$$

as required.

We now show that, for every $\pi \in [\underline{\pi}, \pi_{F_0}]$, a free-learning equilibrium, $(H, F)$, exists such that $\Pi(H, F) = \pi$. Because the equilibrium payoff set is closed,[40] it is sufficient to show every profit $\pi \in (\underline{\pi}, \pi_{F_0})$ can be generated by some equilibrium.[41] Fix such a $\pi$, and define for $q \in [0, 1]$ and $t \in [\pi, 1]$ the following CDF:

$$\begin{aligned} G^q_{\pi,t} &: [0,1] \to [0,1] \\ &x \mapsto \max\{G_{\pi,t}(x), \min\{q, F_0(x)\}\}. \end{aligned}$$

Our proof allows for non-regular priors. As such, we let $[\underline{x}, \bar{x}] = \mathrm{co}(\text{supp } F_0)$. Below, we prove the following lemma:

**Lemma 8** *A $q^*$ exists such that $I_{G^{q^*}_{\pi,1}} \geq 0$, with equality holding for some $\hat{x} \in [\pi, \bar{x}]$ such that $G^{q^*}_{\pi,1}(\hat{x}) = G_{\pi,1}(\hat{x}) \geq q^*$.*

[40] See footnote 28.

[41] Alternatively, notice the vanishing-cost limit of Theorem 2 is a free-learning equilibrium that gives S a profit of $\underline{\pi}$, whereas having B collect full information and S best respond is an equilibrium yielding S a profit of $\pi_{F_0}$.

Before providing the lemma's proof, we show how to use the lemma to obtain an equilibrium. Take $q^*$ and $\hat{x}$ to be as in the lemma. We explain how to find a $t \geq \hat{x}$ such that $G^{q^*}_{\pi,t}$ is a signal. Let $y = \max\{x \in [\underline{x}, \bar{x}] : I_{G^{q^*}_{\pi,1}}(x) = 0\}$. Because $I_{G^{q^*}_{\pi,1}}(\hat{x}) = 0$ and $\hat{x} \in [\pi, \bar{x}] \subseteq [\underline{x}, \bar{x}]$, $y \geq \hat{x}$. As such, $x \in [y, 1]$ implies $G_{\pi,1}(x) \geq q^*$, and therefore, $G^{q^*}_{\pi,1}(x) = G_{\pi,1}(x)$. Thus,

$$I_{G^{q^*}_{\pi,y}}(1) = \int_y^{\bar{x}} (F_0(s) - 1) \, \mathrm{d}s \leq 0 \leq I_{G^{q^*}_{\pi,1}}(1).$$

Because $x \mapsto I_{G^{q^*}_{\pi,x}}(1)$ is continuous, we have that a $t \in [y, 1]$ exists such that $I_{G^{q^*}_{\pi,t}}(1) = 0$. It remains to be verified that $G^{q^*}_{\pi,t}$ is a signal. For $x \leq t$, $G^{q^*}_{\pi,t}(x-) = G^{q^*}_{\pi,1}(x-)$, and so $I_{G^{q^*}_{\pi,t}}(x) = I_{G^{q^*}_{\pi,1}}(x) \geq 0$. For $x > t$,

$$I_{G^{q^*}_{\pi,t}}(x) = I_{G^{q^*}_{\pi,t}}(t) + \int_t^x (F_0(s) - 1)\mathrm{d}s \geq I_{G^{q^*}_{\pi,t}}(t) + \int_t^1 (F_0(s) - 1)\mathrm{d}s = I_{G^{q^*}_{\pi,t}}(1) = 0.$$

Thus, $G^{q^*}_{\pi,t}$ is a signal. We now argue that $(\mathbf{1}_{[\hat{x},1]}, G^{q^*}_{\pi,t})$ is a free-learning equilibrium yielding S a profit of $\pi$. To do so, notice first that $G^{q^*}_{\pi,t}(x-) \geq G_{\pi,1}(x-)$ for all $x$, with equality holding for $x = \hat{x} \geq \pi$. Therefore, $\hat{x} \in P(G^{q^*}_{\pi,t})$, and

$$\pi_{G^{q^*}_{\pi,t}} = \Pi(\hat{x}, G^{q^*}_{\pi,t}) = \Pi(\hat{x}, G_{\pi,t}) = \pi.$$

Moreover, $I_{G^{q^*}_{\pi,t}}(\hat{x}) = I_{G^{q^*}_{\pi,1}}(\hat{x}) = 0$ by choice of $\hat{x}$ and in view of $t \geq y \geq \hat{x}$. Hence, $\hat{x} \in S(I_{G^{q^*}_{\pi,t}}(\hat{x}))$, and so $G^{q^*}_{\pi,t}$ is optimal for B given $\mathbf{1}_{[\hat{x},1]}$.

Hence, all that remains is to prove Lemma 8, which we do now.

## H.1 Proof of Lemma 8

We first show that mean-preserving spreads increase the convex hull of a CDF's support.

**Lemma 9** *Suppose $F \succeq G$. Then,* $\mathrm{co}\,(\mathrm{supp}\ F) \supseteq \mathrm{co}\,(\mathrm{supp}\ G)$.

**Proof.** Let $[x, y] = \mathrm{co}\,(\mathrm{supp}\ F)$ and $[w, z] = \mathrm{co}\,(\mathrm{supp}\ G)$, and suppose $w < x$ for a contradiction (the proof for $z > y$ is analogous). Take $\epsilon > 0$ to be such that $w + \epsilon < x$. Because $w$ must be in $G$'s support, $G(w + \epsilon) > 0$. By contrast, $F(w + \epsilon) = 0$ as $w + \epsilon$ is below $F$'s support. Because these observations are true for every $\epsilon \in (0, x - w)$, we have $\int_0^x F \, \mathrm{d}s = 0 < \int_0^x G \, \mathrm{d}s$, contradicting that $F \succeq G$. ■

Because the support of every signal is contained in $[\underline{x}, \bar{x}] = \mathrm{co}(\mathrm{supp}\ F_0)$ (by Lemma 9), and a truncated Pareto signal exists that is associated with $\underline{\pi}$ (which follows from Theorem 2), $\pi > \underline{\pi} \geq \underline{x}$. We now prove a useful lemma about $G_{\pi,1}$.

**Lemma 10** *$I_{G_{\pi,1}}(x) \geq 0$ for all $x$, with a strict inequality whenever $x > \underline{x}$.*

**Proof.** Note $\pi > \underline{\pi}$ implies $G_{\pi,1}(s) \leq G_{\underline{\pi},1}(s)$ for all $s$, with a strict inequality for $s > \underline{\pi} \geq \underline{x}$. As such, for every $x > \underline{x}$,

$$I_{G_{\pi,1}}(x) = \int_0^x (F_0 - G_{\pi,1})\ \mathrm{d}s \geq \int_0^x (F_0 - G_{\underline{\pi},1})\ \mathrm{d}s \geq \int_0^x (F_0 - G_{\underline{\pi},\bar{t}})\ \mathrm{d}s = I_{G_{\underline{\pi},\bar{t}}}(x) \geq 0,$$

where the first inequality is strict whenever $x \geq \underline{\pi}$. Because $I_{G_{\pi,1}}(\cdot)$ is continuous, we also have that $I_{G_{\pi,1}}(\underline{x}) \geq 0$. ■

Let

$$A = \{x \in [\pi, \bar{x}] : G_{\pi,1}(x) \geq F_0(x-)\}.$$

Note $A$ is closed in view of upper semicontinuity of $G_{\pi,1}(\cdot)$ and lower semicontinuity of $x \mapsto F_0(x-)$. We now show $A$ is non-empty. In particular, we show $A \supseteq P(F_0)$, which is non-empty due to upper semicontinuity of $\Pi(\cdot, F_0)$. By Lemma 1 and $\pi < \pi_{F_0}$, $P(F_0) \subseteq [\pi_{F_0}, \bar{x}] \subseteq [\pi, \bar{x}]$. Moreover, for any $x \in P(F_0)$, $\pi < \pi_{F_0}$ implies

$$F_0(x-) = G_{\pi_{F_0},1}(x-) < G_{\pi,1}(x-) \leq G_{\pi,1}(x).$$

That $P(F_0) \subseteq A$ follows.

In view of the above, $x^* := \min A$ is well defined. We now prove a $q^*$ exists such that the minimal value of $I_{G^{q^*}_{\pi,1}}$ over $A$ is zero.

**Lemma 11** *A $q^* \leq F_0(x^*-)$ exists such that $\min I_{G^{q^*}_{\pi,1}}(A) = 0$.*

**Proof.** The proof is based on the Intermediate Value Theorem. To use this theorem, we note that the mapping

$$(q, x) \mapsto I_{G^q_{\pi,t}}(x) = \int_0^x (F_0 - G^q_{\pi,t})\ \mathrm{d}s$$

is continuous, being the difference between two continuous functions of $(q, x)$. As such, $q \mapsto \min\ I_{G^q_{\pi,1}}(A)$ is continuous in view of the maximum theorem. Moreover, $\min I_{G^0_{\pi,1}}(A) = \min I_{G_{\pi,1}}(A) \geq 0$. In light of the Intermediate Value Theorem, it is sufficient to find a $q > 0$ for which $\min I_{G^0_{\pi,1}}(A) \leq 0$. To do so, note that because $G_{\pi,1}(s) < F_0(s-)$ for all $s < x^*$, we have that

$$\begin{aligned} I_{G^{F_0(x^*-)}_{\pi,1}}(x^*) &= \int_0^{x^*} (F_0 - \max\{G_{\pi,1}(s), \min\{F_0(x^*-), F_0(s)\}\})\ \mathrm{d}s \\ &= \int_0^{x^*} (F_0 - \max\{G_{\pi,1}(s), F_0(s)\})\ \mathrm{d}s = 0. \end{aligned}$$

Because $x^* \in A$, $\min\ I_{G_{\pi,1}^{F_0(x^*-)}}(A) \le I_{G_{\pi,1}^{F_0(x^*-)}}(x^*) = 0$. Thus, we have shown $\min\ I_{G_{\pi,1}^{F_0(x^*-)}}(A) \le 0 = \min I_{G_{\pi,1}^{0}}(A)$, as required. The proof is now complete. ■

The next lemma assures us that $G_{\pi,1}^{q}$ is not a signal only if it has too high a mean.

**Lemma 12** *For all* $x \in [0,1]$, $I_{G_{\pi,1}^{q^*}}(x) \ge 0$.

**Proof.** Divide $[0,1]$ into three subintervals, $[0,\pi)$, $[\pi, x^*]$, and $(x^*, 1]$, showing the desired inequality holds for each at a time. We first show $\inf I_{G_{\pi,1}^{q^*}}([0,\pi)) \ge 0$. To see why, recall that $\pi \ge \underline{x}$, meaning $x < \pi$ only if $G_{\pi,1}(x) = 0$. As such, whenever $x < \pi$,

$$G_{\pi,1}^{q^*}(x) = \max\{0, \min\{q^*, F_0(x)\}\} = \min\{q^*, F_0(x)\} \le F_0(x).$$

Thus, $I_{G_{\pi,1}^{q^*}}(x) \ge \int_0^x (F_0 - F_0)\ \mathrm{d}s = 0$ for all $x \in [0,\pi)$. We now show $\min\ I_{G_{\pi,1}^{q^*}}([\pi, x^*]) \ge 0$. Let $x \in [\pi, x^*]$, and recall that $G_{\pi,1}(s-) < F_0(s-) \le F_0(s)$ must hold for all $s < x$ by choice of $x^*$. As a consequence,

$$\begin{aligned} I_{G_{\pi,1}^{q^*}}(x) &= \int_0^x F_0(s) - \max\{G_{\pi,1}(s), \min\{q^*, F_0(s)\}\}\ \mathrm{d}s \\ &\ge \int_0^x F_0(s) - \max\{G_{\pi,1}(s), F_0(s)\}\ \mathrm{d}s \\ &= \int_0^x F_0(s) - F_0(s)\ \mathrm{d}s = 0. \end{aligned}$$

We thus have that $\min I_{G_{\pi,1}^{q^*}}([0, x^*]) \ge 0$. To complete the proof that $\min I_{G_{\pi,1}^{q^*}}([0,1]) \ge 0$, suppose for a contradiction that $x \in (x^*, 1]$ exists such that $I_{G_{\pi,1}^{q^*}}(x) < 0$. Take

$$x_0 \in \arg\min_{x \in [0,1]} I_{G_{\pi,1}^{q^*}}(x) = \arg\min_{x \in (x^*,1]} I_{G_{\pi,1}^{q^*}}(x).$$

Because $I_{G_{\pi,1}^{q^*}}(x)$ is right differentiable, we have that

$$0 \le I'_{G_{\pi,1}^{q^*}-}(x_0) = F_0(x_0-) - G_{\pi,1}(x_0-),$$

in view of $q^* \le F_0(x^*) \le G_{\pi,1}(x^*)$. Therefore, $F_0(x_0) \ge F(x_0)$; that is, $x_0 \in A$, in contradiction to $\min\ I_{G_{\pi,1}^{q^*}}(A) = 0$. Thus, $I_{G_{\pi,1}^{q^*}}(x) \ge 0$ for all $x$. ■

To conclude the proof of Lemma 8, notice that $x \in A$ only if $G_{\pi,1}(x) \ge F_0(x-) \ge F_0(x^*-) \ge q^*$. Taking $x_1 \in \arg\min_{x \in A} I_{G_{\pi,1}^{q^*}}(x)$, we therefore have

$$G_{\pi,1}^{q^*}(x_1) = \max\{G_{\pi,1}(x_1), \min\{q^*, F_0(x_1)\}\} = \max\{G_{\pi,1}(x_1), q^*\} = G_{\pi,1}(x_1).$$

Thus, $x_1$ is in $A \subseteq [\pi, \bar{x}]$, has $I_{G_{\pi,1}^{q^*}}(x) = 0$, and satisfies $G_{\pi,1}^{q^*}(x) = G_{\pi,1}(x) \ge q^*$; that is, our proof is complete.

## I Proof of Corollary 1: Pareto Payoff Ranking

We prove the corollary by showing $\bar{p}_\pi$ is strictly decreasing in $\pi$ over the interval $[\underline{\pi}, \pi_{F_0}]$. To see why this monotonicity is sufficient, recall that B's free-learning equilibrium payoff is equal to $\int_{\bar{p}_\pi}^1 (s - \bar{p}_\pi)\ \mathrm{d}F_0$, where $\pi$ is S's profit. Hence, B's utility decreases in S's price. If S's price decreases with her profit, we find that higher profits correspond to lower prices and therefore higher B utility. We now show $\bar{p}_\pi$ decreases over the range of feasible free-learning equilibrium profits. For this purpose, take any $\pi < \pi'$ in $[\underline{\pi}, \pi_{F_0}]$. We prove $\bar{p}_{\pi'} < \bar{p}_\pi$ by showing $X_\pi$ contains a price strictly larger than $\bar{p}_{\pi'}$. To find such a price, we make two observations. First, because $\pi < \pi'$, we have that

$$F_0(\bar{p}_{\pi'}) = G_{\pi',1}(\bar{p}_{\pi'}-) = 1 - \frac{\bar{p}_{\pi'}}{\pi'} < 1 - \frac{\bar{p}_{\pi'}}{\pi} = G_{\pi,1}(\bar{p}_{\pi'}-).$$

Second, because $F_0$ is regular, $G_{\pi,1}(1-) = 1 - 1/\pi < 1 = F_0(1-)$. Combining the two observations, we have that $G_{\pi,1}(\bar{p}_{\pi'}) - F_0(\bar{p}_{\pi'}) > 0 > G_{\pi,1}(1-\epsilon) - F_0(1-\epsilon)$ for any small positive $\epsilon$. Because the difference $G_{\pi,1} - F_0$ is continuous on $[0,1)$, we can apply the Intermediate Value Theorem to find some $p \in (\bar{p}_{\pi'}, 1)$ for which $G_{\pi,1}(p) - F_0(p) = 0$. Therefore, $p \in X_\pi$, meaning $\bar{p}_\pi \geq p$. We have thus concluded that $\bar{p}_\pi \geq p > \bar{p}_{\pi'}$, meaning the higher profit level corresponds to a lower price, thereby proving the corollary.

## J Proof of Proposition 1: Costly Learning Equilibria

We show supp $H =$ supp $F =$ co(supp $F$), meaning supp $F$ is a convex set over which S is indifferent; that is, $F$ is a truncated Pareto. Because supp $H \subseteq$ supp $F \subseteq$ co(supp $F$) by Lemma 1, our task is to show co(supp $F$) $\subseteq$ supp $H$.

Letting $[w,z] :=$ co(supp $F$), we wish to show $[w,z] \subseteq$ supp $H$. Suppose otherwise for a contradiction; that is, $[w,z] \cap$ supp $H \neq [w,z]$. We show $x < y$ in supp $F$ exist such that $(x,y) \cap$ supp $H = \emptyset$. To do so, we note supp $H \cap [w,z]$ is a closed set, meaning $[w,z] \setminus$ supp $H$ is open (in $\mathbb{R}$), and so must contain a non-empty open subinterval of $[w,z]$. Let $(x,y)$ be a maximal such subinterval with respect to set containment; that is, $(x,y)$ is such that $(x',y') \cap$ supp $H \neq \emptyset$ for all $(x',y') \supsetneq (x,y)$.[42] Because supp $H$ is closed, if $x \neq w$, then $x \in$ supp $H$; otherwise, $(x-\epsilon, x+\epsilon) \subseteq [w,z] \setminus$ supp $H$ for all small $\epsilon > 0$, meaning $(x,y) \subsetneq (x-\epsilon, y) \subseteq [w,z] \setminus$ supp $H$, a contradiction to maximality of

[42] One can find the subinterval $(x,y)$ by fixing some $(x',y') \subset [w,z] \setminus$ supp $H$, and taking the union of all $(x'',y'') \subseteq [w,z] \setminus$ supp $H$ that contain $(x',y')$.

$(x, y)$. An analogous argument gives $y \neq z$ only if $y \in \text{supp } H$. Hence, we have shown $x, y \in \{w, z\} \cup \text{supp } H$. Because supp $H \subseteq$ supp $F$ (Lemma 1) and $\{w, z\} \subseteq \text{supp } F$, we thus have that $x, y \in \text{supp } F$.

We now construct a family of deviations indexed by $\epsilon > 0$, $F^*_\epsilon$, and obtain a contradiction by showing these deviations must be strictly profitable for B when $\epsilon > 0$ is sufficiently small.

Fix a small $\epsilon > 0$, and note the following are all well defined due to $x, y \in \text{supp } F$:

$$\begin{aligned}
F_{1,\epsilon} &= F(\cdot | s \in [x - \epsilon, x + \epsilon]), \\
F_{2,\epsilon} &= F(\cdot | s \in [y - \epsilon, y + \epsilon]), \\
\beta_{1,\epsilon} &= F(x + \epsilon) - F((x - \epsilon)-) > 0, \\
\beta_{2,\epsilon} &= F(y + \epsilon) - F((y - \epsilon)-) > 0.
\end{aligned}$$

Moreover, take

$$\begin{aligned}
\beta_{0,\epsilon} &= 1 - \beta_{1,\epsilon} - \beta_{2,\epsilon}, \\
F_{0,\epsilon} &= \begin{cases} F(\cdot | s \notin [x - \epsilon, y + \epsilon]) & \text{if } \beta_{0,\epsilon} > 0, \\ \text{arbitrary } F' \in \mathcal{A} & \text{otherwise.} \end{cases}
\end{aligned}$$

Clearly, $F = \sum_{i=0}^{2} \beta_{i,\epsilon} F_{i,\epsilon}$. Moreover, because $x, y \in \text{supp } F$, both $\beta_{1,\epsilon}$ and $\beta_{2,\epsilon}$ are strictly positive for all $\epsilon > 0$. Define

$$\begin{aligned}
s_\epsilon &= \frac{1}{2} \int s \, \mathrm{d}\left(F_{1,\epsilon} + F_{2,\epsilon}\right), \\
\eta_\epsilon &= \min\{\beta_{1,\epsilon}, \beta_{2,\epsilon}, \epsilon\} > 0, \\
F^*_\epsilon &= \beta_{0,\epsilon} F_{0,\epsilon} + \eta_\epsilon \mathbf{1}_{[s_\epsilon, 1]} + (\beta_{1,\epsilon} - 0.5\eta_\epsilon) F_{1,\epsilon} + (\beta_{2,\epsilon} - 0.5\eta_\epsilon) F_{2,\epsilon}.
\end{aligned}$$

In words, $F^*_\epsilon$ takes $0.5\eta_\epsilon$ mass from the $\epsilon$-ball around $x$ and $0.5\eta_\epsilon$ mass from the $\epsilon$-ball around $y$ and pools them to create an $\eta_\epsilon > 0$ mass on $s_\epsilon$. Because $0.5(F_{1,\epsilon} + F_{2,\epsilon}) \succ \mathbf{1}_{[s_\epsilon, 1]}$, $F^*_\epsilon$ is less informative than $F$, which, in turn, is less informative than $F_0$. By transitivity of the information ordering, $F_0$ is more informative than $F^*_\epsilon$; that is, $F^*_\epsilon \in \mathcal{A}$.

Let $T_H(s) = \int_0^s (s - p) \, \mathrm{d}H(p)$ denote B's expected trade surplus conditional on signal realization $s$. Below, we prove

$$\lim_{\epsilon \searrow 0} \int \frac{T_H}{\eta_\epsilon} \, \mathrm{d}(F - F^*_\epsilon) = 0, \tag{7}$$

$$\lim_{\epsilon \searrow 0} \left( \frac{C(F^*_\epsilon) - C(F)}{\eta_\epsilon} \right) < 0, \tag{8}$$

and so obtain the following contradiction to $F$ maximizing $U_\kappa(H,F)$,

$$0 \le \lim_{\epsilon \searrow 0} \frac{U_\kappa(H,F) - U_\kappa(H,F^*_\epsilon)}{\eta_\epsilon} = \lim_{\epsilon \searrow 0} \left[ \int \frac{T_H}{\eta_\epsilon} \,\mathrm{d}(F - F^*_\epsilon) + \kappa \frac{C(F^*_\epsilon) - C(F)}{\eta_\epsilon} \right] < 0, \quad (9)$$

hence completing the proof.

We now explain why (7) and (8) both hold. Because $(x,y) \cap \text{supp } H = \emptyset$, B's trading surplus from receiving a signal $s \in [x,y]$ is given by

$$T_H(s) = \int_0^s (s-p) \,\mathrm{d}H(p) = \int_0^x (s-p) \,\mathrm{d}H(p) = H(x)s - \int_0^x p \,\mathrm{d}H(p). \qquad (10)$$

As such, $T_H$ is affine over $[x,y]$, and so (7) obtains as follows:

$$\begin{aligned} \int \frac{T_H}{\eta_\epsilon} \,\mathrm{d}(F - F^*_\epsilon) &= 0.5 \left( \int T_H \,\mathrm{d}F_{1,\epsilon} + \int T_H \,\mathrm{d}F_{2,\epsilon} \right) - T_H(s_\epsilon) \\ &\to 0.5 T_H(x) + 0.5 T_H(y) - T_H(0.5x + 0.5y) = 0, \end{aligned}$$

where convergence follows from continuity of $T_H(\cdot)$, $s_\epsilon \to 0.5(x+y)$, $F_{1,\epsilon} \to \mathbf{1}_{[x,1]}$, and $F_{2,\epsilon} \to \mathbf{1}_{[y,1]}$. We now use the latter three convergences to obtain (8). To do so, notice these convergences imply

$$\frac{\|F^*_\epsilon - F\|}{\eta_\epsilon} = \left\| \mathbf{1}_{[s_\epsilon,1]} - 0.5\left(F_{1,\epsilon} + F_{2,\epsilon}\right) \right\| \to \left\| \mathbf{1}_{[0.5(x+y),1]} - 0.5\left(\mathbf{1}_{[x,1]} + \mathbf{1}_{[y,1]}\right) \right\| =: M.$$

As such, Fréchet differentiability of $C$ and strict convexity of $c_F$ over co(supp $F$) $\supseteq [x,y]$ yield

$$\begin{aligned} \frac{1}{\eta_\epsilon}\left[C(F^*_\epsilon) - C(F)\right] &= \frac{1}{\eta_\epsilon} \left[ \int c_F \,\mathrm{d}(F^*_\epsilon - F) + o\left(\|F^*_\epsilon - F\|\right) \right] \\ &= \int c_F \,\mathrm{d}\left[ \mathbf{1}_{[s_\epsilon,1]} - 0.5\left(F_{1,\epsilon} + F_{2,\epsilon}\right) \right] + \frac{\|F^*_\epsilon - F\|}{\eta_\epsilon} \left[ \frac{o\left(\|F^*_\epsilon - F\|\right)}{\|F^*_\epsilon - F\|} \right] \\ &\to c_F(0.5x + 0.5y) - \left(0.5 c_F(x) + 0.5 c_F(y)\right) + M \cdot 0 < 0. \end{aligned}$$

Thus, we have (7) and (8), which together yield the contradiction (9). The proof is now complete.

## K Proof of Theorem 2: Vanishing-Cost Equilibrium

Let $\{\kappa_n\}_{n \ge 0}$ be a strictly positive sequence that converges to zero, and take $\{(H_n, F_n)\}_{n \ge 0}$ to be a corresponding sequence of equilibria. Because $\mathcal{F}$ and $\mathcal{A}$ are both compact, $\{(H_n, F_n)\}_{n \ge 0}$ can be seen as a union of convergent subsequences. Without loss, let one of these subsequences be the sequence itself, and let $(H_\infty, F_\infty) \in \mathcal{F} \times \mathcal{A}$ be its limit. To prove the theorem, it is sufficient to show $(H_\infty, F_\infty) = (\mathbf{1}_{[\bar{p}_{\underline{\pi}},1]}, G_{\underline{\pi},\bar{t}})$.

To this end, we first note that because B's objective is a continuous function of $(\kappa, H, F)$, B's best-response correspondence is upper hemicontinuous in $(\kappa, H)$. Therefore, $F_\infty \in \arg\max_{F \in \mathcal{A}} U_0(H_\infty, F)$, meaning supp $H_\infty \subseteq S(F_\infty)$ by Lemma 2. That $H_\infty$ is optimal for S against $F_\infty$ follows from upper hemicontinuity of S's mixed-best-response correspondence, $F \mapsto \arg\max_{H \in \mathcal{F}} \Pi(H, F)$.[43] Thus, the limit $(H_\infty, F_\infty)$ is a free-learning equilibrium. Because the Pareto signal set is closed and $F_\infty$ is the limit of Pareto signals (Proposition 1), we have that $F_\infty$ is itself a Pareto signal; that is, $F_\infty = G_{\pi,t}$ for some $\pi$ and $t$. Below, we argue $\pi = \underline{\pi}$, and so $t = \bar{t}$. Clearly, $\max \Pi(\cdot, G_{\underline{\pi},\bar{t}}) = \underline{\pi}$. Therefore, the free-learning equilibrium $(H_\infty, F_\infty) = (H_\infty, G_{\underline{\pi},\bar{t}})$ gives S a profit of $\underline{\pi}$. That $H_\infty = \mathbf{1}_{[\bar{p}_{\underline{\pi}},1]}$ then follows from Lemma 5.

It remains to show that $\pi = \underline{\pi}$. Suppose otherwise; that is, $\pi > \underline{\pi}$. Because $(H_\infty, G_{\pi,t})$ is a free-learning equilibrium, $H$ is profit maximizing against $G_{\pi,t}$, and so supp $H_\infty \subseteq$ supp $G_{\pi,t}$ (by Lemma 1). We show below that B is strictly better off using $G_{\underline{\pi},\bar{t}}$ against any $p \in$ supp $G_{\pi,t}$. It follows that $G_{\pi,t}$ cannot be optimal for B against $H_\infty$, a contradiction to $(H_\infty, G_{\pi,t})$ being a free-learning equilibrium.

To show $G_{\underline{\pi},\bar{t}}$ is a strictly profitable deviation for B, we first claim $\bar{t} > t$. Suppose $\bar{t} \le t$ for a contradiction. Then, $G_{\underline{\pi},\bar{t}}(s) > G_{\pi,t}(s)$ for all $s \in (\underline{\pi}, t)$, and $G_{\underline{\pi},\bar{t}}(s) = G_{\pi,t}(s)$ for all $s \in [0, \underline{\pi}] \cup [t, 1]$. We therefore get the following contradiction,

$$\begin{aligned} 0 = \int_0^1 s \ \mathrm{d}(G_{\pi,t} - G_{\underline{\pi},\bar{t}})(s) &= \int_0^1 (G_{\underline{\pi},\bar{t}} - G_{\pi,t})(s) \ \mathrm{d}s \\ &= \int_{\underline{\pi}}^t (G_{\underline{\pi},\bar{t}} - G_{\pi,t})(s) \ \mathrm{d}s > 0, \end{aligned}$$

where the first equality follows from both $G_{\pi,t}$ and $G_{\underline{\pi},\bar{t}}$ being signals, and the second equality from integration by parts.

We now conclude the proof by establishing that $U_0(p, G_{\underline{\pi},\bar{t}}) > U_0(p, G_{\pi,t})$ for all

[43] See footnote 22.

$p \in \text{supp } G_{\pi,t} = [\pi, t]$, as shown below:

$$\begin{aligned} U_0(p, G_{\underline{\pi},\bar{t}}) - U_0(p, G_{\pi,t}) &= \int_{s\in[p,1]} (s-p) \, \mathrm{d}(G_{\underline{\pi},\bar{t}} - G_{\pi,t})(s) \\ &= \int_p^1 (G_{\pi,t} - G_{\underline{\pi},\bar{t}})(s) \, \mathrm{d}s \\ &= \int_0^1 (G_{\pi,t} - G_{\underline{\pi},\bar{t}})(s) \, \mathrm{d}s - \int_0^p (G_{\pi,t} - G_{\underline{\pi},\bar{t}})(s) \, \mathrm{d}s \\ &= \int_0^p (G_{\underline{\pi},\bar{t}} - G_{\pi,t})(s) \, \mathrm{d}s = \int_{\underline{\pi}}^{\pi} 1 - \frac{\underline{\pi}}{s} \, \mathrm{d}s + \int_{\pi}^{p} \left( \frac{\pi - \underline{\pi}}{s} \right) \mathrm{d}s \\ &> \int_{\pi}^{p} \left( \frac{\pi - \underline{\pi}}{s} \right) \mathrm{d}s \geq 0, \end{aligned}$$

where the fourth equality follows from both $G_{\pi,t}$ and $G_{\underline{\pi},\bar{t}}$ being signals. The proof is now complete.

# Online Appendix

## L Non-smooth Cost Functions

### L.1 General Results

In this section, we prove our results go through for more general cost functions than specified in the main text. For this purpose, let Lip denote the set of all Lipschitz continuous functions from $[0,1]$ to $\mathbb{R}$, and take $C : \mathcal{A} \to \overline{\mathbb{R}}_+ := \mathbb{R}_+ \cup \{\infty\}$ to be a proper convex function. Given an $F \in \mathcal{A}$, define the **subdifferential** of $C$ at $F$ as

$$\partial C(F) := \left\{ \phi \in \text{Lip} : \int \phi \, \mathrm{d}(F' - F) \le C(F') - C(F) \; \forall F' \in \mathcal{A} \right\}.$$

The goal of this section is to show our results hold as long as $C$ satisfies the following properties:

(i) $C$ is lower semicontinuous.

(ii) Full information can be approximated at finite cost, $F_0 \in \text{cl } C^{-1}(\mathbb{R}_+)$.

(iii) $\phi \in \partial C(F)$ only if, for every $x, y \in \text{supp } F$, and every $\alpha \in (0,1)$,

$$\phi(\alpha x + (1-\alpha)y) < \alpha\phi(x) + (1-\alpha)\phi(y). \tag{11}$$

Properties (i) and (ii) are required for replacing the continuity assumption implied by Fréchet differentiability, whereas Property (iii) provides the appropriate relaxation of Assumption 1 for non-differentiable cost functions. In particular, Fréchet differentiability implies that the subdifferential is non-empty for all signals.[44] In general, a lower semicontinuous and convex $C$ is guaranteed to have a non-empty subdifferential only over a dense subset of its effective domain (Brøndsted and Rockafellar, 1965). By allowing the subdifferential to be empty, Property (iii) can hold for cost functions that are not Fréchet differentiable.

Below, we show Property (i) ensures an equilibrium exists, Property (ii) guarantees the vanishing-cost limit of costly learning equilibria is a free-learning equilibrium, and

[44] For convex functions with an open domain, Fréchet differentiability also implies the subdifferential is a singleton. Because $\mathcal{A}$ has an empty interior, $\partial C(F)$ contains many functions, even under Fréchet differentiability. For illustration, observe that if $\phi$ is in $\partial C(F)$, so is $s \mapsto \phi(s) + as$ for any $a \in \mathbb{R}$, because $\int s \, \mathrm{d}(F' - F'') = 0$ whenever $F', F'' \in \mathcal{A}$.

Property (iii) delivers Proposition 1. It follows that Theorem 2 holds as well. For an explanation, recall that Theorem 2's proof can be broken into three steps. First, as costs vanish, the game's equilibria converge to a free-learning equilibrium. Second, B's signal belongs to the closed class of truncated Pareto distributions, because this class contains her signal when learning is costly. And third, a unique free-learning equilibrium exists in which B uses a Pareto signal. As explained above, Property (ii) delivers the first fact, whereas the second fact is implied by Property (iii). Finally, note the last fact does not depend on B's learning cost. It follows that Theorem 2 continues to hold as long as the above-mentioned conditions are satisfied.

Before proving the above, we show Property (iii) follows from Assumption 1 whenever costs are Fréchet differentiable. To do so, we show the latter assumptions imply a condition that is sufficient for Property (iii). This condition states that one can strictly reduce the cost of any signal by approximating the slight pooling of any two distinct signal realizations. Despite being less elegant than Property (iii), this condition is often easier to verify. We now state the condition formally and show it implies Property (iii).

**Condition 1** *For any $F \in \mathcal{A}$, $\alpha \in (0,1)$, and $x, y \in \operatorname{supp} F$ such that $x \neq y$, a sequence $\{F_n, M_n\}_{n\geq 0}$ exists such that $M_n \in (0,\infty)$, $F_n \in \mathcal{A}$, $F_n \to F$,*

$$M_n\left(F - F_n\right) \xrightarrow{\mathcal{L}_1} \left(\alpha \mathbf{1}_{[x,1]} + (1-\alpha)\mathbf{1}_{[y,1]} - \mathbf{1}_{[\alpha x + (1-\alpha)y,1]}\right), \tag{12}$$

*and*

$$\liminf_n M_n\left[C(F_n) - C(F)\right] < 0. \tag{13}$$

**Lemma 13** *Condition 1 holds for $C$ only if $C$ satisfies Property (iii).*

**Proof.** Suppose $\phi \in \partial C(F)$ for some $F \in \mathcal{A}$. Since $\phi$ is Lipschitz, its derivative, $\phi'$, is well-defined up to a set of zero Lebesgue measure. Fix some distinct $x, y \in \operatorname{supp} F$ and some $\alpha \in (0,1)$. Our goal is to show $\phi$ satisfies (11) if Condition 1 holds. To do so, let

$\{F_n, M_n\}_{n\geq 0}$ be a sequence delivered by the condition. Then,

$$\begin{aligned}
0 &> \liminf M_n \left[C(F_n) - C(F)\right] \\
&\geq \liminf M_n \int_{x\in[0,1]} \phi \ \mathrm{d}(F_n - F) \\
&= \liminf \int_{x\in[0,1]} M_n[F(x) - F_n(x)] \ \mathrm{d}\phi(x) \\
&= \liminf \int_0^1 \phi'(x) M_n[F(x) - F_n(x)] \ \mathrm{d}x \\
&= \int_0^1 \phi'(x)[\alpha \mathbf{1}_{[x,1]} - (1-\alpha)\mathbf{1}_{[y,1]} - \mathbf{1}_{[\alpha x + (1-\alpha)y,1]}] \ \mathrm{d}x \\
&= \phi(\alpha x + (1-\alpha)y) - \alpha\phi(x) - (1-\alpha)\phi(y),
\end{aligned}$$

where the second inequality follows from $\phi \in \partial C(F)$, the first and last equalities from integration by parts, the second equality from $\phi$ being Lipschitz (and therefore absolutely continuous), and the penultimate equality from (12) and $\phi'$ being bounded (due to $\phi \in$ Lip). The proof is now complete. ■

We now show Property (iii) is a relaxation of Assumption 1. To do so, we use Assumption 1 to show $C$ satisfies Condition 1 by constructing a sequence similar to that constructed in the proof of Proposition 1.

**Claim 2** *Let $C : \mathcal{A} \to \mathbb{R}_+$ be a Fréchet differentiable function satisfying Assumption 1. Then, $C$ satisfies Condition 1.*

**Proof.** Take any $F \in \mathcal{A}$, $\alpha \in (0,1)$, and two distinct $x, y \in \operatorname{supp} F$. We show $C$ satisfies Condition 1 by using an argument similar to that of Proposition 1. Specifically, we obtain the sequence required for verifying Condition 1 by pooling neighborhoods of $x$ and $y$.

Fix any small $\epsilon > 0$, define $\{F_{i,\epsilon}, \beta_{i,\epsilon}\}_{i=0}^2$ and $\eta_\epsilon$ as in Proposition 1's proof. Recall from the proposition's proof that $F = \sum_{i=0}^2 \beta_{i,\epsilon} F_{i,\epsilon}$. Let

$$s_{\epsilon,\alpha} = \alpha \int s \ \mathrm{d}F_{1,\epsilon} + (1-\alpha) \int s \ \mathrm{d}F_{2,\epsilon},$$

and define

$$F_\epsilon^\alpha = \beta_{0,\epsilon} F_{0,\epsilon} + \eta_\epsilon \mathbf{1}_{[s_{\epsilon,\alpha},1]} + (\beta_{1,\epsilon} - \alpha\eta_\epsilon)F_{1,\epsilon} + (\beta_{2,\epsilon} - (1-\alpha)\eta_\epsilon)F_{2,\epsilon}. \tag{14}$$

Thus, $F_\epsilon^\alpha$ takes $\alpha\eta_\epsilon$ mass from the $\epsilon$-ball around $x$ and $(1-\alpha)\eta_\epsilon$ mass from the $\epsilon$-ball around y and pools them to create an $\eta_\epsilon > 0$ mass on $s_{\alpha,\epsilon}$. Let us see that $F_\epsilon^\alpha$ is a signal.

By construction, $\alpha F_{1,\epsilon} + (1-\alpha)F_{2,\epsilon}) \succ \mathbf{1}_{[s_{\epsilon,\alpha},1]}$. Noting that $\succeq$ is a transitive relation that respects convex combinations delivers $F_0 \succeq F \succeq F^{\alpha}_{\epsilon}$; that is, $F^{\alpha}_{\epsilon} \in \mathcal{A}$.

Letting $n_0$ be such that $F^{\alpha}_{1/n_0}$ is well defined, let $F_n = F^{\alpha}_{1/(n_0+n)}$ and $M_n = \left(\eta_{1/(n_0+n)}\right)^{-1}$. By construction,

$$\begin{aligned} M_n(F_n - F) &= \mathbf{1}_{[s_{1/(n_0+n),\alpha},1]} - \alpha F_{1,1/(n_0+n)} - (1-\alpha)F_{2,1/(n_0+n)} \\ &\to \mathbf{1}_{[\alpha x+(1-\alpha)y,1]} - \alpha \mathbf{1}_{[x,1]} - (1-\alpha)\mathbf{1}_{[y,1]}. \end{aligned}$$

Thus, (12) holds for $\{F_n, M_n\}_{n\geq 0}$. It remains to show this sequence satisfies (13). Following the argument that establishes (8) in Proposition 1's proof with $0.5x + 0.5y$ replaced by $\alpha x + (1-\alpha)y$ and $F^{\alpha}_{1/(n_0+n)}$ replacing $F^*_{\epsilon}$ delivers the equality

$$M_n(C(F_n - C(F)) \to c_F(\alpha x + (1-\alpha)y) - \alpha c_F(x) - (1-\alpha)c_F(y) < 0,$$

where the inequality follows from Assumption 1. Thus, $\{F_n, M_n\}$ satisfies (13), meaning $C$ satisfies Condition 1, as required. ∎

We now proceed to showing Property (iii) delivers the conclusions of Proposition 1. To show this result, we first prove the following trivial lemma.

**Lemma 14** *Let $\phi \in \mathcal{C}$. Then, $\phi \in \partial C(F)$ if and only if*

$$F \in \arg\max_{F' \in \mathcal{A}} \int \phi \, \mathrm{d}F' - C(F'). \tag{15}$$

**Proof.** By definition, $\phi \in \partial C(F)$ if and only if the following holds for all $F' \in \mathcal{A}$,

$$C(F') - C(F) \geq \int \phi \, \mathrm{d}(F' - F) \iff \int \phi \, \mathrm{d}F - C(F) \geq \int \phi \, \mathrm{d}F' - C(F').$$

∎

We now use Lemma 14 to show Property (iii) implies the conclusions of Proposition 1.

**Proposition 2** *Suppose $C$ satisfies Property (iii) and let $(H, F)$ be an equilibrium in the $\kappa > 0$ game. Then,*

*(i)* supp $H$ = supp $F$ = co(supp $F$)*, and*

*(ii) $F$ is a Pareto signal.*

**Proof.** Suppose $(H, F)$ is an equilibrium of the $\kappa > 0$ game. As explained in the proof of Proposition 1, it is sufficient to prove every $x < y$ in supp $F$ admits some price in between; that is, supp $H \cap (x, y) \neq \emptyset$ whenever $x, y \in$ supp $F$. To prove this claim, observe B's objective can be written as

$$U_\kappa(H, F') = \int T_H \, \mathrm{d}F' - \kappa C(F'),$$

where $T_H(s) = \int \max\{s - p, 0\} \, \mathrm{d}H$ is B's expected trade surplus conditional on her signal realization being $s$. Because $F$ maximizes $U_\kappa(H, \cdot)$, it also maximizes

$$\kappa^{-1} U_\kappa(H, F') = \int \kappa^{-1} T_H \, \mathrm{d}F' - C(F'),$$

meaning, by Lemma 14, that $\kappa^{-1} T_H \in \partial C(F)$. It follows by Property (iii) that $T_H$ satisfies (11) for $x$, $y$ and every $\alpha \in (0, 1)$. To conclude the proof, observe that if supp $H \cap (x, y) = \emptyset$, then $T_H$ is affine over $[x, y]$ by (10), and so could not possibly satisfy (11) for every $\alpha \in (0, 1)$. ∎

Knowing Property (iii) implies Proposition 1, we need to establish two facts to ensure our results continue to hold when $C$ satisfies Properties (i) through (iii). The first fact is that a costly learning equilibrium always exists. The second is that every convergent sequence of vanishing-cost equilibria attains a free-learning equilibrium as its limit. Theorem 3 below proves the former.

**Theorem 3** *Suppose $C$ satisfies Properties (i) and (ii) (and so is lower semicontinuous convex function differing from $\infty$). Then, a $\kappa$-equilibrium exists for all $\kappa \geq 0$.*

**Proof.** Because the $\kappa = 0$ case is identical to the main paper, it remains to show existence for the case in which $\kappa > 0$. By Property (ii), without loss of generality, we can assume $\inf C(\mathcal{A}) = 0$.

Let $U_\kappa^M(H, F) := \max\{U_\kappa(H, F), M\}$ for some $M < 0$. Note $U_\kappa^M$ is upper semicontinuous and quasiconcave because it is a composition of a continuous and increasing function on a concave and upper-semicontinuous function. Consider the game in which S's action set equals $\mathcal{F}$, B's action set is $\mathcal{A}$, and the player's payoffs are given by $\Pi$ and $U_\kappa^M$. We prove the modified game has an equilibrium, after which we show that this equilibrium corresponds to an equilibrium of the original game.

We prove the modified game has an equilibrium by using Corollary 3.3 of Reny (1999). Both $U_\kappa^M$ and $\Pi$ are quasiconcave, lie in $[M, 1]$, and are upper semicontinuous over $\mathcal{F} \times \mathcal{A}$, which is a compact subset of a topological vector space. Thus, this game is compact, quasiconcave, and reciprocally upper semicontinuous. To show existence, it is therefore sufficient to show the game is payoff secure.

Fix any $(H, F)$ and $\epsilon > 0$. Since $U_\kappa$ is continuous in $H$, $F$ secures a payoff of $U_\kappa(H, F) - \epsilon$ from $U_\kappa(H, F)$ (otherwise, one can find a sequence $H_n \to H$ such that $U_\kappa(H_n, F)$ does not converge to $U_\kappa(H, F)$), and so B can secure $\max\{U_\kappa(H, F) - \epsilon, M\} \geq U_\kappa^M(H, F) - \epsilon$ from $U_\kappa^M(H, F)$ using $F$. To show S can secure $\Pi(H, F)$, take $H_\epsilon$ to be the distribution of $\max\{p - \epsilon, 0\}$, where $p$ is drawn according to $H$. Note that for any sequence $\{F_n\}_{n\geq 0}$ that converges to $F$,

$$\begin{aligned}
\liminf_{n\to\infty} \Pi(H_\epsilon, F_n) &\geq \liminf_{n\to\infty} \int (p-\epsilon)[1 - F_n(p-\epsilon)]\mathrm{d}H(p) \\
&\geq \liminf_{n\to\infty} \int p[1 - F_n(p-\epsilon)]\mathrm{d}H(p) - \epsilon \\
&\geq \int \liminf_{n\to\infty} p[1 - F_n(p-\epsilon)]\mathrm{d}H(p) - \epsilon \\
&\geq \int p[1 - F(p-\epsilon)]\mathrm{d}H(p) - \epsilon \\
&\geq \int p[1 - F(p)]\mathrm{d}H(p) - \epsilon = \Pi(H, F) - \epsilon,
\end{aligned}$$

where the first inequality follows from $\max\{p - \epsilon, 0\} \geq p - \epsilon$, the second inequality from $\{s > p - \epsilon\} \subseteq \{s \geq p - \epsilon\}$ and probabilities being less than 1, the third inequality from Fatou's lemma, and the fourth inequality from the Portmanteau theorem. Thus, the modified game is payoff secure and therefore has an equilibrium.

We now show any equilibrium of the modified game must be an equilibrium of the original game. To do so, let $(H, F)$ be an equilibrium of the modified game. Clearly S is best responding, because his objective is the same in both games. To see that B best responds, let $\underline{F}$ be such that $C(\underline{F}) = 0 = \inf C(\mathcal{A})$ (observe $\underline{F}$ exists because $C$ is lower semincontinuous), and note $U_\kappa^M(H, \underline{F}) = U_\kappa(H, \underline{F}) \geq 0$, and so $M < 0 \leq U_\kappa(H, \underline{F}) \leq U_\kappa^M(H, F) = U_\kappa(H, F)$ due to $F$ being optimal in the modified game. Combined with $U_\kappa^M \geq U_\kappa$, $F$ being optimal in the modified game also implies $U_\kappa(H, F) = U_\kappa^M(H, F) \geq U_\kappa^M(H, F') \geq U_\kappa(H, F')$ for all $F' \in \mathcal{A}$. In other words, $F$ maximizes $U_\kappa(H, \cdot)$ over $\mathcal{A}$. Thus, $(H, F)$ is an equilibrium of the original game. ■

To conclude this section, we show the limit of every convergent sequence of $\kappa$-equilibria with $\kappa \to 0$ is a free-learning equilibrium. In doing so, we complete the task of showing our results hold for all cost functions satisfying properties (i) through (iii).

**Lemma 15** *Suppose $C$ satisfies Property (ii). Let $\{\kappa_n\}_{n\geq 0}$ be such that $\kappa_n \searrow 0$, and take $\{(H_n, F_n)\}_{n\geq 0}$ to be a convergent sequence of $\kappa_n$-equilibria attaining $(H_\infty, F_\infty)$ as its limit. Then, $(H_\infty, F_\infty)$ is a free-learning equilibrium.*

**Proof.** Because S's best-response correspondence is upper hemicontinuous (Lemma 7), $H_\infty$ maximizes S's profits when B's signal is $F_\infty$. To prove $F_\infty$ maximizes $U_0(H_\infty, \cdot)$, we show below that every $\epsilon > 0$ admits an $F_\epsilon$ such that $C(F_\epsilon) < \infty$ and $U_0(p, F_\epsilon) \geq U_0(p, F_0) - \epsilon$ for all $p \in [0, 1]$. Therefore,

$$\begin{aligned} U_{\kappa_n}(H_n, F_n) &\geq U_{\kappa_n}(H_n, F_\epsilon) \\ &= U_0(H_n, F_\epsilon) - \kappa_n C(F_\epsilon) \\ &= \int U_0(p, F_\epsilon)\, \mathrm{d}H_n - \kappa_n C(F_\epsilon) \\ &\geq \int U_0(p, F_0) - \epsilon\, \mathrm{d}H_n - \kappa_n C(F_\epsilon) \\ &= U_0(H_n, F_0) - \epsilon - \kappa_n C(F_\epsilon) \to U_0(H_\infty, F_0) - \epsilon, \end{aligned}$$

where the second equality follows from $U_0(H, F) = \int\int \max\{s - p, 0\}\, \mathrm{d}H(p)\, \mathrm{d}F(s)$ and Fubini's (or Tonelli's) Theorem. Recalling that $U_0$ is continuous then delivers

$$U_0(H_\infty, F_\infty) = \lim_{n\to\infty} U_0(H_n, F_n) \geq U_0(H_\infty, F_0) - \epsilon$$

for all $\epsilon > 0$. Therefore, $U_0(H_\infty, F_\infty) = U_0(H_\infty, F_0)$, meaning $F_\infty$ is optimal for B against $H_\infty$ when learning is free.

Hence, all that remains to show is existence of $F_\epsilon$ as above. Fix any $\epsilon > 0$. Because $U_0$ (viewed as a function over $[0, 1] \times \mathcal{A}$) is a continuous function with a compact domain, it is also uniformly continuous. Therefore, a $\delta > 0$ exists such that $|U_0(p, F) - U_0(p, F_0)| < \epsilon$ holds for all $p \in [0, 1]$ whenever $\|F - F_0\| < \delta$. Existence of $F_\epsilon$ then follows from Property (ii), which asserts that a sequence $\{F_m\}_{m\geq 0}$ converging to $F_0$ exists such that $C(F_m) < \infty$ for all $m$. The proof is now complete. ■

## L.2 Posterior Separable Costs

This section shows our results hold when B's learning costs are posterior separable (Caplin et al., 2017). Whereas in our model we formalized B's signal choice using the distribution of her posterior *mean*, to describe posterior separable costs one must model information acquisition using the distribution of B's posterior *belief*. After introducing the posterior-belief framework, we present the definition of posterior separable costs, and show these costs satisfy Properties (i), (ii), and (iii) from the previous section. Because these properties are sufficient for our results, our conclusions hold for posterior separable costs as well.

A few preliminary definitions are in order. Given a compact metrizable set $X$, let $\Delta X$ be the set of all Borel probability measures over it equipped with the weak* topology. When $X$ is also subset of a locally convex space, one can define the **barycenter** of any $\xi \in \Delta X$ as the unique element $x \in \overline{\text{co}}(X)$, such that $f(x) = \int f \, \mathrm{d}\xi$ for every continuous linear functional on $X$. We denote this barycenter by $E\xi$. As this notation suggests, $E\xi = \int x \, \mathrm{d}\xi$ whenever $X \subset \mathbb{R}$.

We now proceed with describing the posterior-based approach for modeling B's information acquisition. Our presentation is terse. For a more complete and extensive presentation, see a standard reference (e.g., Kamenica and Gentzkow, 2011). Taking $V = [0, 1]$ to be the set of possible valuations B may have, a belief for B is member $\mu$ of $\Delta V$. It is well known that $\Delta V$ is isomorphic to $\mathcal{F}$, with each $F$ having a unique $\mu_F \in \Delta V$ such that $F(x) = \mu[0, x]$ for all $x$. We denote B's prior by $\mu_0 = \mu_{F_0}$. B updates this prior upon observing a signal realization to obtain her posterior. Thus, every signal structure induces a distribution for B's posterior belief, $p \in \Delta\Delta V$. As noted by the literature, (e.g., Aumann and Maschler, 1966; Benoît and Dubra, 2011; Kamenica and Gentzkow, 2011) $p$ describes the distribution of B's posterior following some information structure if and only if it attains B's prior as its mean; that is, $Ep = \mu_0$. We let $\mathcal{R} = \{p \in \Delta\Delta V : Ep = \mu_0\}$ be the set of all such distributions over posteriors, and refer to elements of $\mathcal{R}$ as **random posteriors**. Thus, letting B choose a random posterior is equivalent to letting B use any signal structure to learn about her valuation.

To understand the connection between the random posterior framework and our modeling approach, recall only the mean of B's belief matters for her trade outcomes. In other words, B's purchasing decision and resulting surplus is the same for all beliefs that share the same expectations. As such, without loss, we can replace every belief $\mu$ with its expectation, $E\mu$, and replace $p$ with its induced distribution over means,

$$F_p(x) = p\{\mu : E\mu \leq x\}.$$

As explained in the main text, it is well-known (e.g., Gentzkow and Kamenica, 2016) that $F = F_p$ for some $p \in \mathcal{R}$ if and only if $F \in \mathcal{A}$.

Next, we define posterior separable costs, and explain how such costs can be mapped to our way of modeling costly learning. A cost function $\mathcal{C} : \mathcal{R} \to \mathbb{R}$ is **posterior separable** if a continuous, strictly convex function $\psi : \Delta V \to \mathbb{R}_+$ exists such that $\mathcal{C}(p) = \int \psi(\mu)\ \mathrm{d}p$. Given a $\psi$, one can obtain the cost of a particular posterior mean distribution $F$ by finding the cheapest random posterior that induces it. Specifically, let

$$\mathcal{R}(F) = \{p \in \mathcal{R} : F_p = F\}$$

be the set of random posteriors inducing $F \in \mathcal{A}$ as its mean distribution. Then the cost of each $F \in \mathcal{A}$ is given by the cheapest random posterior generating it,

$$C_\psi(F) = \min_{p \in \mathcal{R}(F)} \int \psi\ \mathrm{d}p.$$

Note the above minimization problem is well-defined because $R(F)$ is compact valued. In fact, one can show $R(\cdot)$ is a Kakutani correspondence over $\mathcal{A}$, that is, an upper hemicontinuous correspondence whose values are non-empty, compact, and convex.

In what follows, we argue $C_\psi$ satisfies properties (i) through (iii). Two of the three properties are immediate: Property (i) follows from Berge's theorem (e.g., Lemma 17.3 in Aliprantis and Border, 2006), whereas Property (ii) follows from $\psi$ being continuous and therefore bounded on $\Delta V$. Thus, showing Property (iii) is all that remains.

To show Property (iii), we prove $C_\psi$ satisfies Condition 1. Thus, fix some $F \in \mathcal{A}$, $\alpha \in (0, 1)$, and distinct $x, y \in \operatorname{supp} F$. Consider the sequence $\{F_n, M_n\}$ constructed in the proof of Claim 2. Given the claim's proof, we only need to show (13). To do so, we fix some $p \in \arg\min_{p \in \mathcal{R}(F)}$ and construct a sequence $p_n \in \mathcal{R}(F_n)$ such that

$$\liminf_n M_n \int \psi\ \mathrm{d}(p_n - p) < 0.$$

It follows that

$$\liminf_n M_n \left[C_\psi(F_n) - C_\psi(F)\right] \leq \liminf_n M_n \int \psi\ \mathrm{d}(p_n - p) < 0,$$

meaning $\{F_n, M_n\}_n$ satisfies the desired inequality.

We now construct the sequence $\{p_n\}_{n \geq 0}$. Recall that $F_n = F^\alpha_{\epsilon_n}$, where $\epsilon_n = 1/(n_0 + n)$ for some sufficiently large $n_0$ so that $\{\beta_{i,\epsilon_n}\}_{n \geq 0, i=0,1,2}$ and $\eta_{\epsilon_n}$ are strictly positive, and $\{F_{i,\epsilon_n}\}_{n \geq 0, i=0,1,2}$ are all well defined. For every $n \geq 0$, partition $\Delta V$ into three sets:

$$\begin{aligned}
D_{0,n} &= \{\mu \in \Delta\Theta : E\mu \notin [x - \epsilon_n, x + \epsilon_n] \cup [y - \epsilon_n, y + \epsilon_n]\},\\
D_{1,n} &= \{\mu \in \Delta\Theta : E\mu \in [x - \epsilon_n, x + \epsilon_n]\}, \text{ and}\\
D_{2,n} &= \{\mu \in \Delta\Theta : E\mu \in [y - \epsilon_n, y + \epsilon_n]\}.
\end{aligned}$$

By choice of $n_0$, $p(D_{i,n}) = \beta_{i,\epsilon_n} > 0$ for all $i$ and $n$. One can therefore decompose $p$ into three probability measures, $p = \sum_{i=1}^{3} p(D_{i,n})p_{i,n}$, where $p_{i,n}$ are defined via

$$\mathrm{d}p_{i,n} = \frac{\mathbf{1}_{D_{i,n}}}{p(D_{i,n})}\ \mathrm{d}p.$$

Notice $F_{p_{i,n}} = F_{i,\epsilon_n}$. Moreover, passing to a subsequence if necessary, we may assume each $p_{i,n}$ converges to some $p_i$, where $p_i \neq p_j$ for any $i \neq j$. Letting $\eta_n := \eta_{\epsilon_n}$, we can use the above decomposition to define $p_n$ as

$$p_n = p - \eta_n(\alpha p_{1,n} + (1-\alpha)p_{2,n}) + \eta_n \delta_{E(\alpha p_{1,n}+(1-\alpha)p_{2,n})},$$

where $\delta_\mu \in \Delta\Delta V$ denotes the probability measure putting mass 1 on $\mu$. It is straightforward to show $F_{p_n} = F_n$; that is, $p_n \in \mathcal{R}(F_n)$. Finally, recall $\eta_n = \eta_{\epsilon_n} = 1/M_n$, meaning

$$\begin{aligned}\liminf_n M_n \int \psi\ \mathrm{d}(p_n - p) &= \int \psi\ \mathrm{d}(\delta_{E(\alpha p_{1,n}+(1-\alpha)p_{2,n})} - \alpha p_{1,n} - (1-\alpha)p_{2,n}) \\ &= \liminf_n \left[\psi(E(\alpha p_{1,n} + (1-\alpha)p_{2,n})) - \alpha \int \psi\ \mathrm{d}p_{1,n} - (1-\alpha)\int \psi\ \mathrm{d}p_{2,n}\right] \\ &= \psi(E(\alpha p_1 + (1-\alpha)p_2)) - \alpha \int \psi\ \mathrm{d}p_1 - (1-\alpha)\int \psi\ \mathrm{d}p_2 < 0,\end{aligned}$$

where the inequality follows from strict convexity of $\psi$. The proof is now complete.

## M Random Production Costs

This section extends our model to the case in which S's production cost is random. Thus, suppose S's production cost is given by a random variable, $\mathbf{m}$, distributed according to a full-support CDF, $H^{\mathbf{m}} \in \mathcal{F}$. Given production cost $m$, a price $p$, and a B signal $F$, S's profit is now given by

$$\begin{aligned}\Pi : [0,1] \times [0,1] \times \mathcal{F} &\to \mathbb{R} \\ (p, m, F) &\mapsto (p-m)(1-F(p-)).\end{aligned}$$

For analytical convenience, we describe S's strategy space via distributional strategies (Milgrom and Weber, 1985). Recall that a behavioral strategy for S maps every cost realization to a distribution over prices. Each such strategy generates a joint distribution for S's price, $\mathbf{p}$, and S's cost, $\mathbf{m}$, whose marginal CDF over the cost coordinate is $H^{\mathbf{m}}$. The set of all such joint distributions, which we denote by $\mathcal{H}$, is the set of S's distributional strategies. As noted by Milgrom and Weber (1985), using distributional strategies rather than mixed strategies neither increases nor reduces the model's generality. Indeed, we can see every $H \in \mathcal{H}$ is generated by some behavioral S-strategy, meaning that the mapping from behavioral strategies to $\mathcal{H}$ is surjective. This mapping is also injective, at least up to an $H^{\mathbf{m}}$-almost sure equality[45]: two behavioral strategies generate the same $H \in \mathcal{H}$ if and only if they are $H^{\mathbf{m}}$-almost surely equal. We therefore use $\mathcal{H}$ as S's strategy space.

Given an S strategy $H \in \mathcal{H}$, we use $H^{\mathbf{p}}$ to denote its marginal distribution over prices, $H^{\mathbf{p}|\mathbf{m}} : \mathbb{R} \to \mathcal{F}$ to denote a version of $H$'s conditional price distribution given S's cost realization, and write $H^{\mathbf{p}|\mathbf{m}}(p|m) := H^{\mathbf{p}|\mathbf{m}}(p)(m)$. Given a $\bar{m} \in (0,1]$, we also write

$$H^{\mathbf{p}|\mathbf{m}}(\cdot|\mathbf{m} < \bar{m}) := [H^{\mathbf{m}}(\bar{m}-)]^{-1} \int_{m<\bar{m}} H^{\mathbf{p}|\mathbf{m}}(\cdot|m) \, \mathrm{d}H^{\mathbf{m}}(m)$$

to denote the price distribution conditional on $\mathbf{m}$ being strictly below $\bar{m}$.

We abuse notation and let

$$\Pi(H,F) = \int \Pi(p,m,F) \, \mathrm{d}H(p,m)$$

denote S's expected profits given a strategy profile, $(H,F)$. B's utility given $\kappa \geq 0$ is as before,

$$U_\kappa(H,F) = \int \max\{s-p,0\} \, \mathrm{d}H^p(p) \, \mathrm{d}F(s) - C(F).$$

A $\kappa$-equilibrium is a profile, $(H,F) \in \mathcal{H} \times \mathcal{A}$ such that

1. $H$ maximizes $\Pi(\cdot,F)$ over $\mathcal{H}$;

2. $F$ maximizes $U_\kappa(H,\cdot)$ over $\mathcal{A}$.

As before, we let $\mathcal{E}_\kappa$ denote the set of all $\kappa$-equilibrium profiles. In addition, let $\mathcal{U}_\kappa$ be the set of $\kappa$-equilibrium payoffs.

When B uses a signal whose support lies below S's cost, the best S can do is obtain zero profits. In such cases, prices below S's cost may be optimal, because they lead to no trade. The possibility of S using such weakly dominated prices creates technical difficulties. In particular, such prices can result in S's best-response correspondence violating upper hemicontinuity. To circumvent this difficulty, we often concentrate on the set of equilibria in which S's strategy is not dominated. Specifically, we focus on the set $\mathcal{E}^*_\kappa$, which is all

[45]We occasionally abuse terminology and identify CDFs over $[0,1]^n$ for any $n$ with their induced Borel probability measures.

$\kappa$-equilibria such that S's price is above his production cost almost surely; that is, $\mathcal{E}_\kappa^*$ consists of all $(H, F) \in \mathcal{E}_\kappa$ such that supp $H \subseteq \{(m,p) : p \geq m\}$. Similarly, let $\mathcal{U}_\kappa^*$ be the set of all payoffs attainable via an equilibrium in which S's price is always above his cost, that is, all equilibria in $\mathcal{E}_\kappa^*$.

The next section proves $\mathcal{E}_\kappa^*$ is non-empty, and so an equilibrium exists. Moreover, we establish that $\mathcal{E}_\kappa^*$ is sufficient for discussing the player's possible equilibrium payoffs. In other words, we show $\mathcal{U}_\kappa^* = \mathcal{U}_\kappa$.

## M.1 Equilibrium Existence

In this section, we prove an equilibrium always exists, and that the equilibrium correspondence taking $\kappa$ to equilibrium strategy profiles and expected values is closed.

Observe that B's objective is still continuous in $(H, \kappa, F)$ and concave in $F$, and so B's value function remains continuous, and her best-response correspondence is Kakutani.

We now turn to proving a Berge-like maximum theorem for S. To do so, we first observe a certain semicontinuity of S's expected profit.

**Lemma 16** $\Pi$ *is upper semincontinuous for all* $(p, m, F)$ *such that* $p \geq m$.

**Proof.** Take $\{(p_n, m_n, F_n)\}_{n \geq 0}$ to be a convergent sequence whose limit $(p_\infty, m_\infty, F_\infty)$ satisfies $p_\infty \geq m_\infty$. We show

$$\limsup \Pi(p_n, m_n, F_n) \leq \Pi(p_\infty, m_\infty, F_\infty)$$

by separately considering two cases. Suppose first $p_\infty = m_\infty$. Then,

$$\Pi(p_n, m_n, F_n) = (p_n - m_n)(1 - F_n(p_n-)) \leq |(p_n - m_n)| \to 0 = \Pi(p_\infty, m_\infty, F_\infty).$$

Suppose now that $p_\infty > m_\infty$. In this case, $p_n > m_n$ for all sufficiently large $n$. Moreover, fixing any continuity point $p'_k \in (p_\infty - 1/k, p_\infty)$ of $F_\infty$ (which exists because $F_\infty$ can have at most countably many discontinuities), one has $p_n > p'_k$ for all $n$ large. Hence, for all sufficiently large $n$,

$$\begin{aligned}\Pi(p_n, m_n, F_n) = (p_n - m_n)(1 - F_n(p_n-)) &\leq (p_n - m_n)(1 - F_n(p'_k-)) \\ &\to (p_\infty - m_\infty)(1 - F_\infty(p'_k)) \\ &\leq (p_\infty - m_\infty)(1 - F_\infty(p_\infty - 1/k)),\end{aligned}$$

where the convergence follows from $p'_k$ being a continuity point of $F_\infty$. Observing that $F_\infty(p_\infty - 1/k) \to F_\infty(p_\infty-)$ as $k \to \infty$ completes the proof. ■

Using the previous lemma, we now show S's maximal profit,

$$\pi^* : [0,1] \times \mathcal{F} \to \mathbb{R}$$
$$(m, F) \mapsto \sup_{p \in [0,1]} \Pi(m, p, F),$$

can be attained by some price $p \geq m$. Observe $\Pi \leq 1$, and so $\pi^*$ is finite.

**Lemma 17** *Every $(m, F)$ admits a $p \in [m, 1]$ such that $\Pi(p, m, F) = \pi^*(m, F)$.*

**Proof.** Observe first that $\pi^*(m, F) \geq 0$, because $\Pi(m, m, F) = 0$ for all $(m, F)$. Moreover, $p = m$ attains $\pi^*(m, F)$ whenever $\pi^*(m, F) = 0$. Thus, only the case in which $\pi^*(m, F) > 0$ remains. Because $[0, 1]$ is compact, a sequence $\{p_n\}_{n \geq 0}$ exists attaining some $p_\infty \in [0, 1]$ as its limit such that $\Pi(p_n, m, F) \to \pi^*(m, F)$. That $\pi^*(m, F) > 0$ implies $\Pi(p_n, m, F) = (p_n - m)(1 - F(p_n-)) > 0$ for all large $n$, meaning $p_\infty > m$. To complete the proof, apply Lemma 16 to obtain

$$\pi^*(m, F) = \lim_n \Pi(p_n, m, F) \leq \Pi(p_\infty, m, F) \leq \pi^*(m, F),$$

where the last equality comes from the definition of $\pi^*$. ■

The above lemma implies $\pi^*(m, F) = \max_{p \in [0,1]} \Pi(p, m, F) = \max_{p \in [m,1]} \Pi(p, m, F)$. Moreover, the correspondence

$$P^* : [0,1] \times \mathcal{F} \rightrightarrows \mathbb{R}$$
$$(m, F) \mapsto \arg\max_{p \in [m,1]} \Pi(p, m, F),$$

taking $(m, F)$ pairs to the corresponding profit-maximizing prices above $m$ is non-empty. By Lemma 16, $P^*(m, F)$ is closed valued, and so is compact valued by virtue of being bounded. Next, we show some continuity properties of $\pi^*$ and $P^*$.

**Lemma 18** *$\pi^*$ is continuous and $P^*$ is upper hemicontinuous.*

**Proof.** Let $\{(m_n, F_n)\}_{n \geq 0}$ be some sequence in $[0, 1] \times \mathcal{F}$ attaining $(m_\infty, F_\infty)$ as its limit.

We begin by showing $\pi^*$ is continuous. Recall $\pi^*(m_n, F_n) = \max_{p \in [m_n,1]} \Pi(p, m_n, F_n)$, and that $\Pi$ is upper semicontinuous over the graph of the correspondence $m \mapsto [m, 1]$ (by Lemma 16). Because this correspondence is continuous (and so upper hemicontinuous), $\pi^*$ is upper semicontinuous (Aliprantis and Border, 2006, Lemma 17.30). As such, it suffices to show $\pi^*$ is lower semincontinuous; that is, $\liminf_{n \to \infty} \pi^*(m_n, F_n) \geq$

$\pi^*(m_\infty, F_\infty)$. Observing $\pi^*(m_n, F_n) \geq \Pi(m_n, m_n, F_n) = 0$ delivers this inequality whenever $\pi^*(m_\infty, F_\infty) = 0$. Therefore, suppose $\pi^*(m_\infty, F_\infty) > 0$, and take any $p \in P(m_\infty, F_\infty)$. Then, for all small $\epsilon > 0$,

$$\pi^*(m_n, F_n) \geq \Pi(p - \epsilon, m_n, F_n) \geq (p - \epsilon - m_n)(1 - F_n(p - \epsilon)).$$

In particular, the above holds for $\epsilon \in (0, p - m_\infty)$. Therefore,

$$\begin{aligned}
\liminf_n \pi^*(c_n, F_n) &\geq \liminf_n (p - \epsilon - m_n)(1 - F_n(p - \epsilon)) \\
&= (p - \epsilon - m_\infty) \liminf_n (1 - F_n(p - \epsilon)) \\
&\geq (p - \epsilon - m_\infty)(1 - F_\infty(p - \epsilon)) \geq (p - m_\infty)(1 - F_\infty(p-)) - \epsilon,
\end{aligned}$$

where the second inequality follows from the Portmanteau theorem. Taking $\epsilon > 0$ above to be arbitrarily small delivers the desired inequality.

We now show $P^*$ is upper hemicontinuous. To do so, we take any convergent sequence $p_n \in P^*(m_n, F_n)$ attaining $p_\infty$ as its limit. Observe $p_n \geq m_n$, and so $p_\infty \geq m_\infty$. Moreover, because $\Pi$ is upper semicontinuous over the graph of $m \mapsto [m, 1]$ and $\pi^*$ is continuous,

$$\pi^*(m_\infty, F_\infty) = \lim_n \pi^*(m_n, F_n) = \limsup_n \Pi(p_n, m_n, F_n) \leq \Pi(p_\infty, m_\infty, F_\infty) \leq \pi^*(m_\infty, F_\infty)$$

Thus, $\Pi(p_\infty, m_\infty, F_\infty) = \pi^*(m_\infty, F_\infty)$ and $p_\infty \geq m_\infty$, that is, $p_\infty \in P^*(m_\infty, F_\infty)$. ■

Next, we use upper hemicontinuity of $P^*$ to show S's undominated best-response correspondence,

$$\begin{aligned}
&\mathcal{H}^* : \mathcal{F} \rightrightarrows \mathcal{H} \\
&\qquad F \mapsto \{H \in \mathcal{H} : \text{supp } H \subseteq \text{graph } P^*(\cdot, F)\},
\end{aligned}$$

is upper hemicontinuous, too. Since $\mathcal{H}$ is compact, it is enough to take any sequence $\{F_n\}_n$ attaining some $F_\infty$ as its limit, and any $H_n \in \mathcal{H}^*(F_n)$ such that $H_n \to H_\infty$, and show supp $H_\infty \subseteq$ graph $P^*(\cdot, F_\infty)$. Now, on the one hand, supp$(\cdot)$ is lower hemicontinuous, and so $(p_\infty, m_\infty) \in$ supp $H_\infty$ only if we can pass to a subsequence that admits a $(p_n, m_n) \in$ supp $H_n$ for every $n$ such that $(p_n, m_n) \to (p_\infty, m_\infty)$. On the other hand, $P^*$ is upper hemicontinuous, and so $p_n \in P(m_n, F_n)$ implies $p_\infty \in P(m_\infty, F_\infty)$. Therefore, $(p_\infty, m_\infty) \in$ supp $H_\infty$, only if $p_\infty \in P(m_\infty, F_\infty)$; that is, $H_\infty \in \mathcal{H}^*(F_\infty)$, as required.

We are now ready to show an equilibrium exists.

**Theorem 4** *$\mathcal{E}^*_\kappa$ is non-empty for all $\kappa \geq 0$.*

**Proof.** The result follows from Kakutani's fixed-point theorem. To see how to apply the theorem, observe that, because B's object is continuous and concave, her best-response correspondence, $\mathcal{A}^*(H) := \arg\max_{F \in \mathcal{A}} U_\kappa(H, F)$, is non-empty, compact valued, convex valued, and upper hemicontinuous. As for S, we have already shown $\mathcal{H}^*(\cdot)$ is non-empty and upper hemicontinuous. It is also convex valued (because S's objective is affine in $H$) and compact valued (because $\Pi(\cdot, \cdot, F)$ is upper semicontinuous over the relevant range). Thus, Kakutani's fixed-point theorem applies to the correspondence $(H, F) \rightrightarrows \mathcal{H}^*(F) \times \mathcal{A}^*(H)$, delivering an equilibrium. ■

The above also delivers that the graphs of both $\kappa \mapsto \mathcal{E}^*_\kappa$ and $\kappa \mapsto \mathcal{U}^*_\kappa$ are closed. Next, we show $\mathcal{E}^*_\kappa$ is sufficient for analyzing the player's equilibrium payoffs.

**Lemma 19** *Let $(H, F) \in \mathcal{E}_\kappa$. Then, a $H^* \in \mathcal{H}$ exists such that $(H^*, F) \in \mathcal{E}^*_\kappa$. Moreover, $(H, F)$ and $(H^*, F)$ yield the same payoffs to both players, and so $\mathcal{U}^*_\kappa = \mathcal{U}_\kappa$.*

**Proof.** Take any $(H, F) \in \mathcal{E}_\kappa \setminus \mathcal{E}^*_\kappa$. We find a $H^* \in \mathcal{H}$ such that $(H^*, F) \in \mathcal{E}^*_\kappa$, $U_\kappa(H^*, F) = U_\kappa(H, F)$ and $\Pi(H^*, F) = \Pi(H, F)$. To construct $H^*$, we replace every losing price with a price of 1. Clearly, this replacement can only increase S's profits. Because $H$ was optimal against $F$, $H^*$ must be optimal against $F$ as well. As for B, note this replacement results in a first-order dominance increase in S's price distribution, and so weakly reduces B's utility across all signals. B's utility from $F$, however, remains unchanged: because $H$ was optimal against $F$, B must never accept any losing prices under $F$, and so replacing these prices with 1 does not change $B$'s trade surplus from $F$. In other words, we have $U_\kappa(H^*, F) = U_\kappa(H, F) \geq U_\kappa(H, F') \geq U_\kappa(H^*, F')$ for all $F' \in \mathcal{A}$. It follows that $(H^*, F)$ is a $\kappa$-equilibrium attaining the same value to both players as $(H, F)$, and such that S's prices are always above his costs, namely $(H^*, F) \in \mathcal{E}^*_\kappa$. The proof is now complete. ■

An implication of the above results is that the correspondence $\kappa \mapsto \mathcal{U}_\kappa$ is closed, a fact that we use when analyzing the case of vanishing learning costs.

## M.2 S's Best Response

This section develops several lemmas regarding the set of profit maximizing prices. Whereas many of these results are known, we record them here for completeness. We begin by noting S's objective exhibits increasing differences in costs.

**Lemma 20** *If $m > m'$ and $p \geq p'$,*

$$\Pi(p, m, F) - \Pi(p', m, F) \geq \Pi(p, m', F) - \Pi(p', m', F).$$

*Moreover, if $F(p-) > F(p'-)$, the inequality is strict.*

**Proof.** Consider $m > m'$ and $p \geq p'$ as above. Then,

$$\begin{aligned}\Pi(p, m, F) - \Pi(p', m, F) &= (p - m)(1 - F(p-)) - (p' - m)(1 - F(p'-)) \\ &= p(1 - F(p-)) - p'(1 - F(p')) + m(F(p-) - F(p'-)) \\ &\geq p(1 - F(p-)) - p'(1 - F(p')) + m'(F(p-) - F(p'-)) \\ &= \Pi(p, m', F) - \Pi(p', m', F),\end{aligned}$$

where the inequality is strict whenever $F(p'-) > F(p-)$, as required. ■

We now use the above lemma to show S's best response is increasing in its costs.

**Lemma 21** *Suppose $p \in P(m, F)$ and $p' \in P(m', F)$. Then, $p \vee p' \in P(m \vee m', F)$ and $p \wedge p' \in P(m \wedge m', F)$. Moreover, if $F(p-) \neq F(p'-)$ and $m \neq m'$, then either $p \vee p' \notin P(m \wedge m', F)$ or $p \wedge p' \notin P(m \vee m', F)$.*

**Proof.** Without loss, assume $m \geq m'$. Clearly, the lemma holds if $m = m'$. Suppose $m > m'$. By Lemma 20,

$$\Pi(p \vee p', m, F) - \Pi(p \wedge p', m, F) \geq \Pi(p \vee p', m', F) - \Pi(p \wedge p', m', F), \tag{16}$$

where the inequality is strict whenever $F(p-) \neq F(p'-)$. Suppose now that $p \vee p' \notin P(m, F)$. Then $p \vee p' \in P(m', F)$ and $p \wedge p' \in P(m, F)$. In this case, (16) delivers $\Pi(p \vee p', m, F) \geq \Pi(p \wedge p', m, F)$, and so $p \vee p' \in P(m, F)$, a contradiction. An analogous argument establishes $p \wedge p' \in P(m', F)$.

Finally, note that if $F(p-) \neq F(p'-)$, Lemma 20 says that the inequality in (16) is strict. Hence, if $p \vee p' \in P(m', F)$, then $\Pi(p \vee p', m, F) > \Pi(p \wedge p', m, F)$; that is, $p \vee p' \in P(m', F)$ only if $p \wedge p' \notin P(m, F)$. Taking the contra-positive then delivers that $p \wedge p' \in P(m, F)$ only if $p \vee p' \notin P(m', F)$. ■

We now point out an unrelated fact: dominated prices are never optimal for S when the maximum of $F$'s support lies above his costs $m$.

**Lemma 22** *If $\max \operatorname{supp} F > m$, then $P(m, F) = P^*(m, F)$ and $\pi^*(m, F) > 0$.*

**Proof.** Let $\bar{x} = \max \text{supp } F$. If $\bar{x} > m$, then a $p \in (m, \bar{x})$ exists such that $F(p-) < 1$, and so $\Pi(p, m, F) > 0 \geq \Pi(p', m, F)$ for all $p' < m$. It follows that $P(m, F) \subseteq (m, 1]$, and so equals $P^*(m, F)$, as required. ■

Next, we generalize the result from Lemma 1: for production costs that attain a strictly positive profit, a price $p$ is optimal only if it is in the support of B's signal.

**Lemma 23** *If $p \in P^*(m, F)$ and $\pi^*(m, F) > 0$, then $p \in \text{supp } F$.*

**Proof.** Together, $p \in P^*(m, F)$ and $\pi^*(m, F) > 0$ imply $p > m$. Moreover, $p \in P^*(m, F)$ implies that, for all $p'$,

$$(p - m)(1 - F(p-)) = \Pi(p, m, F) \geq \Pi(p', m, F) = (p' - m)(1 - F(p'-)).$$

Therefore, since $p > m$, it follows that if $p' > p$, then $1 - F(p'-) < 1 - F(p-)$. In other words, $F(p'-) > F(p'-)$. Because $p'$ was an arbitrary price above $p$, $p \in \text{supp } F$ follows. ■

We conclude this section by showing that, for $F \in \mathcal{A}$, the bottom of $H^{\mathbf{p}}$'s support is always a best response for S when he faces zero costs.

**Lemma 24** *Suppose $F \in \mathcal{A}$ and that $H \in \mathcal{H}$ maximizes $\Pi(\cdot, F)$ over $\mathcal{H}$. Then,*

$$\min \text{supp } H^{\mathbf{p}} \in P^*(0, F).$$

**Proof.** Let $p_* = \min \text{supp } H^{\mathbf{p}}$. Then $p_* \in P(m, F)$ for some $m \in [0, 1]$. If $m = 0$, we are done. Suppose then that $m > 0$. Take any $m' \in (0, m)$ strictly smaller than $\bar{v}$. We argue below that $p_* \in P^*(m', F)$. Because $m'$ is arbitrary, $p_* \in P^*(0, F)$ follows from upper hemicontinuity.

To show $p_* \in P^*(m', F)$, we show $p_* \in P(m', F)$. This inclusion is sufficient because $m' < \bar{v} \leq \max \text{supp } F$ (where the latter inequality follows from $F \in \mathcal{A}$), and so $\pi^*(m', F) > 0$ and $P^*(m', F) = P(m', F)$ by Lemma 22.

For a contradiction, suppose $p_* \notin P(m', F)$. We claim that $P(m'', F) \subset [0, p_*)$ for all $m'' < m'$, and so $H^{\mathbf{p}}(p_*-) \geq H^{\mathbf{m}}(m'-) > 0$, delivering the contradiction $\min \text{supp } H^{\mathbf{p}} < p_* = \min \text{supp } H^{\mathbf{p}}$. Towards this goal, note Lemma 21 implies $P(m', F) \subseteq [0, p_*)$: if a $p \in P(m', F) \cap [p_*, 1]$ exists, then Lemma 21 would imply $p_* \wedge p = p_* \in P(m', F)$, because $m' < m$. Recalling that $P(m', F)$ is non-empty, one can then find a $p \in P(m', F) \cap [0, p_*)$. Using this $p$, one can apply Lemma 21 again to obtain that $P(m'', F) \cap [p_*, 1] \subseteq P(m', F) \cap [p_*, 1] = \varnothing$ for any $m'' < m'$, concluding the proof. ■

## M.3 Free Learning

The goal of the current section is to characterize the set of free-learning equilibrium payoffs when S faces random production costs. The section has three main results. Proposition 3 describes the players' strategies in a free-learning equilibrium, showing they can be parameterized by the price that S charges under free-production. S charges this price whenever $m$ is such that he would have charged a lower price under perfect learning, and the CDF of B's signal above this price agrees with the value distribution. For higher values of $m$, S sets the same price as he would under full information. Our second main result is Theorem 5, which notes S's free-production price in this equilibrium determines S's payoff conditional on production being free, which in turn is sufficient for determining both players' equilibrium payoffs. Finally, Corollary 2 points out that both players' ex-ante payoffs are increasing functions of S's profit under costless production, implying all free-learning equilibria are strongly Pareto ranked.

We begin by noting that, because of regularity, S has a unique profit-maximizing price for every cost when B learns perfectly. Moreover, this price is an increasing and continuous function of S's costs.

**Lemma 25** *For every $m \in [0,1]$, a $p_0(m) \geq m$ exists such that $P(m, F_0) = \{p_0(m)\}$. Moreover, viewed as a function, $p_0$ is increasing and continuous.*

**Proof.** Define
$$p_0(m) := \sup\{p \in [0,1] : p - \frac{1 - F_0(p)}{f_0(p)} \leq m\},$$
which is obviously increasing. Because $F_0$ is regular, $p < (>) p_0(m)$ if and only if
$$m > (<) p - \frac{1 - F_0(p)}{f_0(p)}.$$
Note, however, that
$$\frac{\partial \Pi}{\partial p} = 1 - F_0(p) - (p - m) f_0(p).$$
Therefore, $\Pi(\cdot, m, F_0)$ is strictly increasing on $[0, p_0(m))$ and strictly decreasing on $(p_0(m), 1]$. It follows that $P(m, F) \cap ([0, p_0(m)) \cup (p_0(m), 1]) = \varnothing$, meaning $P(m, F_0) = \{p_0(m)\}$ because $P(m, F_0)$ is non-empty.

Finally, note $\max \text{supp } F_0 = 1 > m$, and so $P(m, F_0) = P^*(m, F_0)$ for all $m < 1$. Because $P^*$ is upper hemicontinuous and upper hemicontinuity reduces to continuity for singleton-valued correspondences, $p_0(\cdot)$ is continuous on $[0, 1)$. Continuity at 1 follows from noting $p_0(m) \in [m, 1]$ for all $m$, and so $m_n \to 1$ implies $p_0(m_n) \to 1 = p_0(1)$. ■

The following lemma generalizes a conclusion mentioned in the main text to the random-production-cost case. Specifically, the lemma points out that in a free-learning equilibrium, S's profits are below his profits under full information for $H^{\mathbf{m}}$-almost all costs.

**Lemma 26** *If* $(H, F) \in \mathcal{E}$, *then* $\pi^*(m, F) \leq \pi^*(m, F_0)$.

**Proof.** By Lemma 2, supp $H^{\mathbf{p}} \subseteq S(F)$, and so $F(p) = F(p-) = F_0(p)$ for all $p \in$ supp $H^{\mathbf{p}}$ by Lemma 3. Because supp $H \subseteq$ graph $P(\cdot, F)$ by optimality, it follows that for $H^{\mathbf{m}}$-every $m$, the following holds for every $p \in \text{supp} H^{\mathbf{p}|\mathbf{m}}(\cdot|m)$:

$$\pi^*(m, F) = \Pi(p, m, F) = (p-m)(1-F(p-)) = (p-m)(1-F_0(p)) = \Pi(p, m, F_0) \leq \pi^*(m, F_0),$$

as required. ■

Next we note that if a CDF lies above (below) B's CDF at some point, it must do the same for a small neighborhood above it. The result follows rather directly from continuity of $F_0$ and CDFs being right continuous.

**Lemma 27** *If* $F \in \mathcal{F}$ *is such that* $F(x) > F_0(x)(F(x) < F_0(x))$ *at some* $x \in [0, 1]$, *then a* $y > x$ *exists such that the same holds for every* $x' \in [x, y]$.

**Proof.** We prove the lemma for $F(x) > F_0(x)$. The proof for the other inequality is analogous. For a contradiction, suppose no such $y$ exists. Then, one can find a decreasing sequence $x_n \searrow x$ such that $F(x_n) - F_0(x_n) \leq 0$ for all $x_n$. By continuity of $F_0$ and right continuity of $F$, it follows that $0 \geq \lim_n F(x_n) - F_0(x_n) = F(x) - F_0(x)$, a contradiction. ■

Using the above lemma and preceding results, we now show B's signal in a free-learning equilibrium equals her value distribution for all values above S's minimal price. Before we state the lemma, recall that given a $\pi \in [0, \pi^*(0, F_0)]$, $\bar{p}_\pi$ denotes the highest price yielding a zero-cost S a profit of $\pi$ when B learns perfectly, $\bar{p}_\pi = \max\{p : \Pi(p, 0, F_0) = \pi\}$.

**Lemma 28** *Let* $(H, F) \in \mathcal{E}_0$. *Let* $p_* = \min \text{supp} H^{\mathbf{p}}$. *Then* $p_* = \bar{p}_{\pi^*(0,F)}$ *and* $F(x-) = F(x) = F_0(x)$ *for all* $x \geq p_*$.

**Proof.** Suppose $(H, F) \in \mathcal{E}_0$. By Lemma 24, $p_* \in P^*(0, F)$. Because $p_* \in \text{supp } H^{\mathbf{p}}$, it follows from Lemma 2 that $p_*$ is $F$-separating. Lemma 5 then implies $p_* = \bar{p}_{\pi^*(0,F)}$.[46] Appealing to Lemma 3 then delivers that $F(p_*-) = F(p_*) = F_0(p_*)$.

We now establish that $F(x) \geq F_0(x)$ for all $x \geq p_*$. Suppose, for a contradiction, that an $x \geq p_*$ exists such that $F(x) < F_0(x)$. By Lemma 27, a $y > x$ exists such that $F(p') < F_0(p')$ for all $p' \in [x, y]$. Now, observe that $x \geq p_* = \bar{p}_{\pi^*(0,F)} \geq p_0(0)$, and so $p_0^{-1}[x, y]$ is a non-trivial interval, $[\min\ p_0^{-1}(x), \max\ p_0^{-1}(y)]$, by Lemma 25. In particular, $p_0^{-1}[x, y]$ is a set of positive $H^{\mathbf{m}}$-measure. Observe, however, that every $m \in p_0^{-1}[x, y]$ satisfies

$$\begin{aligned}
\Pi(p_0(m), m, F) &= (p_0(m) - m)(1 - F(p_0(m)-)) \\
&> (p_0(m) - m)(1 - F_0(p_0(m))) \\
&= \Pi(p_0(m), 0, F_0) = \pi^*(m, F_0),
\end{aligned}$$

a contradiction to Lemma 26. It follows that $F(x) \geq F_0(x)$ for all $x \geq p_*$.

To conclude the proof, suppose $F(x) > F_0(x)$ for some $x \geq p_*$. Then Lemma 27 delivers a $y > x$ such that $F(s) > F_0(s)$ for all $s \in [x, y]$. Evaluating the mean-preserving-spread constraint at $y$ then gives

$$\begin{aligned}
I_F(y) &= I_F(p_*) + \int_{p_*}^{y} (F_0 - F)(s)\ \mathrm{d}s \\
&= \int_{p_*}^{y} (F_0 - F)(s)\ \mathrm{d}s \\
&\leq \int_{x}^{y} (F_0 - F)(s)\ \mathrm{d}s < 0,
\end{aligned}$$

where the second equality follows from $p_*$ being $F$-separating, the first inequality from $F_0(s) \leq F(s)$ holding for all $s \geq p_*$, and the last inequality from choice of $y$. Thus, $I_F(y) < 0$, a contradiction to $F \in \mathcal{A}$. Thus, $F(x) = F_0(x)$ for all $x > p_*$. All that remains is to show $F(x-) = F_0(x)$ for all such $x$, which follows from continuity of $F_0$. ■

Next, we characterize S's optimal prices when B's strategy satisfies the structure given by the previous lemma. This characterization is useful both for characterizing and for constructing free-learning equilibria.

**Lemma 29** *Suppose $F \in \mathcal{A}$ is such that a $p_* \in P(0, F) \cap S(F)$ exists for which $F(x) = F(x-) = F_0(x)$ for all $x \geq p_*$. Then, $P(m, F) = \{p_* \vee p_0(m)\}$ for all $m > 0$.*

[46] Whereas Lemma 5 takes as given a free-learning equilibrium, its proof actually establishes that $P^*(0, F) \cap S(F) \subseteq \{\bar{p}_{\pi^*(0,F)}\}$.

**Proof.** We begin by showing $P(m,F) \cap [0,p_*) = \varnothing$ for all $m > 0$. To do so, suppose a $p \in P(m,F) \cap [0,p_*)$ exists for a contradiction. By Lemma 21, $p \in P(0,F)$. To get a contradiction, we argue next $F(p-) < F(p_*-)$, and so $p \notin P(m,F)$ by Lemma 21 (because $p_* \in P(0,F)$). To obtain the desired inequality, observe $F(p_*-) < 1$, because $\pi^*(0,F) > 0$ for all $F \in \mathcal{A}$. Therefore, if $F(p_*-) = F(p-)$, S's profits under $p_*$ would be strictly larger than under $p$ for all production costs, meaning $p$ cannot be in $P(m,F)$, as required.

Next, we establish that $P(m,F) = \{p_0(m)\}$ whenever $p_0(m) \geq p_*$. Because $P(m,F) \cap [0,p_*) = \varnothing$, it is sufficient to show $p_0(m)$ strictly maximizes $\Pi(\cdot,m,F)$ over $[p_*,1]$. But this maximization is obvious: $p_0(m)$ strictly maximizes $\Pi(\cdot,m,F_0)$ (Lemma 25), which equals $\Pi(\cdot,m,F)$ over $[p_*,1]$ (because $F(x-) = F_0(x)$ for all $x \geq p_*$).

To complete the proof, we prove $P(m,F) = \{p_*\}$ whenever $p_0(m) < p_*$. Again, because no price strictly below $p_*$ is optimal, and because $\Pi(\cdot,m,F)$ coincides with $\Pi(\cdot,m,F_0)$, showing $p_*$ maximizes $\Pi(\cdot,m,F_0)$ over $[p_*,1]$ is enough. Recalling that $\Pi(\cdot,m,F_0)$ is a strictly quasiconcave function (because of regularity) that attains its maximum at $p_0(m) < p_*$ completes the proof. ■

Using the above lemmas, we now provide a characterization of a free-learning equilibrium when S's production cost is random. Given some price $p$, let $\hat{H}_p \in \mathcal{H}$ be the S strategy satisfying

$$\hat{H}_p(p',m) = \int_{m' \leq m} \mathbf{1}_{[p_0(m') \vee p, 1]}(p') \, \mathrm{d}H^{\mathbf{m}}(m').$$

In words, $\hat{H}_p$ is the strategy obtained by having S charge $p_0(m') \vee p$ whenever his production cost is $m'$. That is, S either charges her optimal perfect-learning price, or $p$, whichever is higher. The following proposition shows free-learning equilibria are characterized by three properties. First, S uses $\hat{H}_p$ for some $p$. Second, $p$ is a separating price that maximizes S's profits when costs are zero. And third, B's signal coincides with the value distribution above $p$.

**Proposition 3** *Fix some strategy profile $(H,F) \in \mathcal{H} \times \mathcal{A}$, and let $p_* = \min \operatorname{supp} H^{\mathbf{p}}$. Then, $(H,F)$ is a free-learning equilibrium if and only if*

*(i)* $p_* \in S(F) \cap P(0,F)$.

*(ii)* $F(x) = F(x-) = F_0(x)$ *for all* $x \geq p_*$.

*(iii)* $H = \hat{H}_{p_*}$.

**Proof.** To see the above three conditions are necessary for a free-learning equilibrium, take any free-learning equilibrium, $(H, F)$. Condition (i) then follows from Lemma 24 and Lemma 2. Lemma 28 then delivers (ii), whereas (iii) is implied from Lemma 29.

For sufficiency, suppose $(H, F)$ satisfy the above conditions. That $H$ is profit maximizing follows from Lemma 29. For $F$ to be optimal for B, we need to show S only charges separating prices, that is, that $I_F(p) = 0$ for all $p \in \text{supp } H^{\mathbf{p}}$. To see this equality holds, observe that for any such $p$,

$$I_F(p) = I_F(p_*) + \int_{p_*}^{p} F_0 - F \, \mathrm{d}s = 0,$$

where the last equality comes from $p_*$ being $F$-separating and condition (ii). ∎

Using the above characterization, we now describe the set of feasible free-learning equilibrium payoffs. For this purpose, recall that $\underline{\pi} = \min_{F \in \mathcal{A}} \max_{p \in [0,1]} \Pi(p, 0, F)$ is S's minmax profit, and that given a $\pi \in [0, \pi^*(0, F_0)]$, we use

$$\bar{p}_\pi = \max\{p : \Pi(p, 0, F_0) = \pi\}$$

to denote the highest price yielding S a profit of $\pi$ when his production cost is zero and B learns her value. Using these definitions, we now provide the random-production-cost generalization of Theorem 1. The generalization shows that, just like the free-production case studied in the main text, the set of free-learning equilibrium payoffs can be parameterized by S's profit when his costs are zero. Specifically, define the following mappings:

$$\Pi^0 : [\underline{\pi}, \pi^*(0, F_0)] \to \mathbb{R},$$
$$\pi^0 \mapsto \int \Pi(\bar{p}_{\pi^0} \vee p_0(m), m, F_0) \, \mathrm{d}H^{\mathbf{m}}(m),$$
$$U^0 : [\underline{\pi}, \pi^*(0, F_0)] \to \mathbb{R},$$
$$\pi^0 \mapsto \int U_0(\bar{p}_{\pi^0} \vee p_0(m), F_0) \, \mathrm{d}H^{\mathbf{m}}(m).$$

The theorem below proves the image of $(\Pi^0, U^0)$ gives the set of free-learning equilibrium payoffs. Moreover, given a free-learning equilibrium $(H, F)$, players' payoffs are given by $(\Pi^0, U^0)$ evaluated at $\pi^*(0, F)$.

**Theorem 5** *The following hold:*

*(i) A profile $(H, F)$ is a free-learning equilibrium only if $\Pi(H, F) = \Pi^0(\pi^*(0, F))$ and $U_0(H, F) = U^0(\pi^*(0, F))$, and $\pi^* \in [\underline{\pi}, \pi^*(0, F_0)]$.*

*(ii) If* $(\pi, u) = (\Pi^0, U^0)(\pi^0)$ *for some* $\pi^0 \in [\underline{\pi}, \pi^*(0, F_0)]$*, then* $(\pi, u) \in \mathcal{U}_0$*.*

**Proof.** Suppose first $(H, F) \in \mathcal{E}_0$. By Proposition 3, $p_* = \min \text{supp } H^{\mathbf{p}}$ is in $S(F) \cap P(0, F)$, meaning $p_* = \bar{p}_{\pi^*(0,F)}$ by Lemma 5. Let $\pi_0 := \pi^*(0, F)$. Clearly, $\pi_0 \geq \underline{\pi}$ by definition of $\underline{\pi}$. Moreover, we have by Lemma 26 that $\pi_0 \leq \pi^*(0, F_0)$. Proposition 3 and direct computation then delivers that $(\pi, u) = (\Pi^0, U^0)(\pi_0)$, as required.

Suppose now $(\pi, u) = (\Pi^0, U^0)(\pi^0)$ for some $\pi^0 \in [\underline{\pi}, \pi^*(0, F_0)]$. We construct an equilibrium that delivers $(\pi, u)$ as its payoff profile. If $\pi^0 = \pi^*(0, F_0)$, then $(\pi, u)$ is the player's payoffs under perfect learning, and so can be attained by having B use $F_0$ and S best respond. Suppose then that $\pi^0 < \pi^*(0, F_0)$. Let $\bar{p} := \bar{p}_{\pi^0}$ and $\underline{p} = \min\{p : \Pi(p, 0, F_0) = \pi^0\}$. Observe that $\underline{p} < \bar{p}$ because $\pi^0 < \pi^*(0, F_0)$, meaning $X_{\pi^0} = \{p : \Pi(p, 0, F_0) = \pi^0\}$ has two distinct prices. For any $y \in [0, \underline{p}]$, define the function

$$\tilde{G}_y : [0, 1] \to \mathbb{R}_+$$
$$x \mapsto \begin{cases} F_0(x) & \text{if } x \in [y, \underline{p}] \cup [\bar{p}, 1] \\ G_{\pi^0,1}(x) & \text{if } x \in [0, y) \cup [\underline{p}, \bar{p}]. \end{cases}$$

Observe that $\tilde{G}_y$ is well-defined because $F_0$ and $G_{\pi,1}$ coincide on $\{\underline{p}, \bar{p}\}$. Moreover, because $F_0 \geq G_{\pi,1}$ on $[0, \underline{p}] \cup [\bar{p}, 1]$, we have that $\tilde{G}_y$ is right continuous and increasing. Since $\tilde{G}_y$ is positive and $\tilde{G}_y(1) = 1$, it constitutes a CDF over $[0, 1]$, that is, $\tilde{G}_y \in \mathcal{F}$. In addition, these properties deliver $\bar{p} \in P(0, \tilde{G}_y)$ for all $y \in [0, \underline{p}]$. Figure 8 below illustrates $\hat{G}_y$ for the case in which $F_0$ is uniform.

Below we show a $y^* \in [0, \underline{p}]$ exists for which $\tilde{G}_{y^*} \in \mathcal{A}$ and $\bar{p}$ is $\tilde{G}_{y^*}$-separating. Proposition 3 then establishes that $(\hat{H}_{\bar{p}}, \tilde{G}_{y^*})$ is a free-learning equilibrium. Direct computation reveals the players' payoffs in this equilibrium equal $(\pi, u)$.

We now find a $y^*$ such that $I_{\tilde{G}_{y^*}}(\bar{p}) = 0$. To do so, observe that, because $\Pi(\cdot, 0, F_0)$ is strictly quasiconcave (due to regularity of $F_0$), $\Pi(p, 0, F_0) > \pi^0$ for all $p \in (\underline{p}, \bar{p})$, and so $F_0(p) < G_{\pi^0,1}(p)$ for all $p \in (\underline{p}, \bar{p})$. Hence,

$$I_{\tilde{G}_0}(\bar{p}) = \int_{\underline{p}}^{\bar{p}} F_0 - G_{\pi^0,1} \, \mathrm{d}s < 0.$$

Moreover, notice that $\tilde{G}_{\underline{p}}$ coincides with $G_{\pi^0,1}$ over $[0, \bar{p}]$. Therefore,

$$I_{\hat{G}_{\bar{p}}}(\bar{p}) = I_{G_{\pi^0,1}}(\bar{p}) \geq 0,$$

where the inequality follows from Lemma 10 for $\pi^0 > \underline{\pi}$, and $I_{G_{\pi^0,1}}(\bar{p}) = I_{G_{\underline{\pi},\bar{t}}}(\bar{p}_{\underline{\pi}}) = 0$ when $\pi^0 = \underline{\pi}$.[47] Because $y \mapsto I_{\tilde{G}_y}(\bar{p})$ is continuous, a $y^* \in [0, \bar{p}]$ exists for which $I_{\tilde{G}_{y^*}}(\bar{p}) = 0$.

All that remains is to show $\tilde{G}_{y^*} \in \mathcal{A}$, that is, $I_{\tilde{G}_{y^*}}(x) \geq 0$ for all $x$, with equality at $x = 1$. To prove this claim, suppose first that $x \geq \bar{p}$. In this case,

$$I_{\tilde{G}_{y^*}}(x) = I_{\tilde{G}_{y^*}}(\bar{p}) + \int_{\bar{p}}^{x} F_0 - \tilde{G}_{y^*} \,\mathrm{d}s = 0,$$

where the last equality follows from $\tilde{G}_{y^*}(s) = F_0(s)$ for all $s \geq \bar{p}$. Suppose now that $x \leq y^*$. Then,

$$I_{\tilde{G}_{y^*}}(x) = \int_0^x F_0 - \tilde{G}_{y^*} \,\mathrm{d}s = \int_0^x F_0 - G_{\pi^0,1} \,\mathrm{d}s = I_{G_{\pi^0,1}}(x) \geq 0,$$

where $\tilde{G}_{y^*}(s) = G_{\pi,1}(s)$ for all $s \leq y^*$ implies the first equality, and the inequality follows from Lemma 10 when $\pi^0 > \underline{\pi}$ and from $I_{G_{\underline{\pi},1}}(s) \geq I_{G_{\underline{\pi},1}}(\bar{p}) \geq I_{G_{\underline{\pi},\bar{t}}}(\bar{p}) \geq 0$ for $\pi^0 = \underline{\pi}$. Finally, consider the remaining case in which $x \in (y^*, \bar{p})$. Then,

$$I_{\tilde{G}_{y^*}}(x) = I_{\tilde{G}_{y^*}}(\bar{p}) - \int_x^{\bar{p}} F_0 - \tilde{G}_{y^*} \,\mathrm{d}s = \int_x^{\bar{p}} \tilde{G}_{y^*} - F_0 \,\mathrm{d}s \geq 0,$$

where the last equality follows from $\tilde{G}_{y^*}(s) \geq F_0(s)$ for all $s \in (y^*, \bar{p})$.[48] Thus, we have established $\tilde{G}_{y^*} \in \mathcal{A}$, completing the proof. ■

Next, we make two observations. First, the functions $\Pi^0$ and $U^0$ are strictly increasing. And, second, equilibrium payoffs are strongly Pareto ranked. Combined with Theorem 5, the corollary below implies players' free-learning utilities are strictly increasing functions of S's profits under costless production.

**Corollary 2** *The following hold:*

*(i) Both $\Pi^0$ and $U^0$ are strictly increasing.*

*(ii) For any $(\pi, u), (\pi', u') \in \mathcal{U}_0$, $\pi \geq \pi'$ if and only if $u \geq u'$.*

*(iii) If $(H, F) \in \mathcal{E}_0$, then $\Pi(H, F) = \Pi^0(\pi^*(0, F))$ and $U_0(H, F) = U^0(\pi^*(0, F))$.*

[47] Recall $(\mathbf{1}_{\bar{p}_{\underline{\pi}}}, G_{\underline{\pi},\bar{t}})$ is a free-learning equilibrium in the original game. It follows that $\bar{p}_{\underline{\pi}} \in \text{supp } G_{\underline{\pi},\bar{t}} \subseteq [0, \bar{t}]$, and that $\bar{p}_{\underline{\pi}} \in S(G_{\underline{\pi},\bar{t}})$.

[48] To see this last inequality, observe $\tilde{G}_{y^*}(s) = F_0(s)$ for all $s \in (y^*, \underline{p}]$, and $\tilde{G}_{y^*}(s) = G_{\pi^0,1}(s) > F_0(s)$ for all $s \in (\underline{p}, \bar{p})$ (because $\Pi(s, 0, F_0) > \pi^0$ for all such $s$ due to strict quasiconcavity).

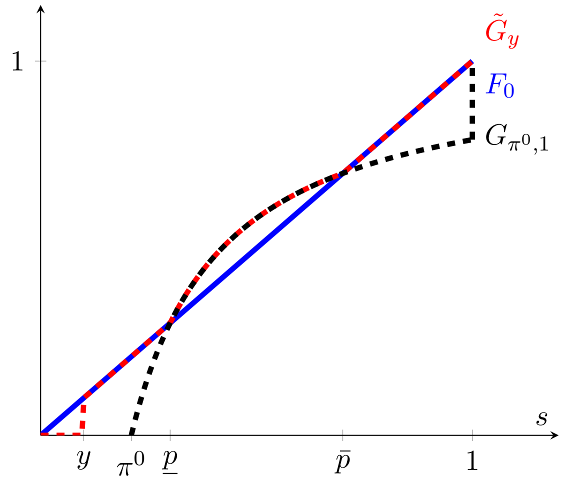


Figure 8: The CDF $\tilde{G}_y$ constructed in the proof of Theorem 5.

**Proof.** Note (ii) follows immediately from (i) and Theorem 5. Therefore, showing (i) is sufficient. Towards this goal, fix some $\pi < \pi'$. Observe $\bar{p}_\pi > \bar{p}_{\pi'}$ (see proof of Corollary 1). Because $\bar{p}_\pi > \bar{p}_{\pi'} \geq p_0(0)$ and $p_0$ is continuous, strictly increasing, and satisfies $p_0(1) = 1$, it follows that the set $p_0^{-1}(\bar{p}_{\pi'}, \bar{p}_\pi)$ has a strictly positive $H^{\mathbf{m}}$-probability. Direct computation then shows $U^0(\pi) < U^0(\pi')$. To see $\Pi^0(\pi) < \Pi^0(\pi')$, recall $\Pi(\cdot, m, F_0)$ is a strictly quasiconcave function attaining its maximum at $p_0(m)$. Hence, $\Pi(p_0(m) \vee \bar{p}_\pi, m, F_0) \leq \Pi(p_0(m) \vee \bar{p}_{\pi'}, m, F_0)$, where the inequality is strict for any $m \leq \min p_0^{-1}(\bar{p}_\pi)$. The desired inequality then follows from noting $[0, \min p_0^{-1}(\bar{p}_\pi)]$ has positive $H^{\mathbf{m}}$-probability. ∎

Thus, we have established that free-learning equilibria are strongly Pareto ranked. Moreover, both players' payoffs in such an equilibrium are a strictly increasing function of S's profits conditional on production being costless.

## M.4 Costly Learning

In this section, we provide a partial characterization of equilibria under costly learning. In particular, we show an analogue of Proposition 1 holds when one conditions on S's costs being below the top of the support of B's signal.

**Proposition 4** *Suppose $(H, F)$ is an equilibrium of the $\kappa$-learning game. Letting $\bar{x} = \max \operatorname{supp} F$, we have that*

$$\operatorname{supp} F = \operatorname{co}(\operatorname{supp} F) = \operatorname{supp} H^{\mathbf{p}|\mathbf{m}}(\cdot|\mathbf{m} < \bar{x}).$$

**Proof.** In the next paragraph we that if $x < y$ are both in supp $F$, then $H^{\mathbf{p}|\mathbf{m}}(\cdot|\mathbf{m} < \bar{x})$ contains some prices strictly in between; that is, $(x, y) \cap \text{supp } H^{\mathbf{p}|\mathbf{m}}(\cdot|\mathbf{m} < \bar{x}) \neq \varnothing$. Lemmas 22 and 23 then deliver supp $H^{\mathbf{p}|\mathbf{m}}(\cdot|\mathbf{m} < \bar{x}) \subseteq \text{supp } F$. It follows that supp $F$ is convex, because supports are closed. Once we know supp $F$ is convex, the claim proven below implies supp $H^{\mathbf{p}|\mathbf{m}}(\cdot|\mathbf{m} < \bar{x})$ contains a dense subset of supp $F$, and so contains supp $F$, again due to closedness of supports. The proposition then follows.

We now show that if $x, y \in \text{supp } F$ are such that $x < y$, then $(x, y) \cap \text{supp } H^{\mathbf{p}|\mathbf{m}}(\cdot|\mathbf{m} < \bar{x})$ is non-empty. To do so, observe first that the argument in Proposition 1 establishes that $(x, y) \cap \text{supp } H^{\mathbf{p}} \neq \varnothing$. Because $\text{supp } H^{\mathbf{p}} = \text{supp } H^{\mathbf{p}|\mathbf{m}}(\cdot|m \geq \bar{x}) \cup \text{supp } H^{\mathbf{p}|\mathbf{m}}(\cdot|m < \bar{x})$, $\text{supp } H^{\mathbf{p}|\mathbf{m}}(\cdot|m \geq \bar{x}) \subseteq [\bar{x}, 1]$ (because $H$ maximizes S's profits), and $y \leq \bar{x}$, it follows that $(x, y) \cap \text{supp } H^{\mathbf{p}} \subseteq \text{supp } H^{\mathbf{p}|\mathbf{m}}(\cdot|m < \bar{x})$, as required. The proof is now complete. ■

## M.5 Vanishing Costs

In this section, we study what happens when learning costs vanish when S's production cost is random. We show these equilibria converge to the worst free-learning equilibrium. This result shows our conclusions are robust to random perturbations in S's production cost.

We begin by showing that the vanishing-cost limit of every sequence of costly-learning equilibria is a free-learning equilibrium. A proof is needed for this result because S's best-response correspondence is not upper hemicontinuous when S's production costs are non-zero. To establish this convergence, we therefore rely on Lemma 19, which shows replacing dominated prices with non-dominated ones preserves equilibrium and payoffs.

**Lemma 30** *Let $\{(H_n, F_n)\}_{n \geq 0}$ be a sequence of $\kappa_n$-equilibria with $\kappa_n > 0$ that attains $(H_\infty, F_\infty)$ as a limit, where $\kappa_n \to 0$. Then, $(H_\infty, F_\infty)$ is a free-learning equilibrium.*

**Proof.** To prove the lemma, we begin with a sequence as in the lemma's premise. We then replace S's strategy in each of the sequence's elements with a strategy in non-dominated prices that preserves equilibrium and payoffs. Whereas the new sequence need not converge, it attains a subsequence that does. In the subsequence's limit, B's strategy is given by $F_\infty$. Because a limit of equilibria in non-dominated S strategies is an equilibrium itself, one obtains that $F_\infty$ satisfies the conditions of Proposition 3. Using these conditions, we establish that, eventually, every S-cost type other than 1 can earn strictly positive payoffs in all but finitely many elements of the original sequence. Using upper hemicontinuity of S's non-dominated best-response correspondence then delivers that the original sequence's limit is indeed a free-learning equilibrium.

We now proceed with the detailed arguments. Let $\{(H_n, F_n)\}_{n\geq 0}$ be as required by the Lemma. By Lemma 19, every $n$ admits a $H_n^* \in \mathcal{H}$ such that $(H_n^*, F_n) \in \mathcal{E}_{\kappa_n}^*$ and delivers the same payoffs to both players (under $\kappa_n$) as $(H_n, F_n)$. Because $\{(H_n^*, F_n)\}_{n\geq 0}$ lives in a compact set, it has a convergent subsequence in which $\{H_{n_k}^*\}_{k\geq 0}$ converges to some limit $H_\infty^*$, and so $(H_{n_k}^*, F_{n_k}) \to (H_\infty^*, F_\infty)$. Because the correspondence $\kappa \mapsto \mathcal{E}_\kappa^*$ has a closed graph, we have $(H_\infty^*, F_\infty) \in \mathcal{E}_0^*$, and so satisfies the conditions of Proposition 3.

Letting $p_* = \min \operatorname{supp} H^{*\mathbf{p}}$, Proposition 3 implies $p_* \in S(F_\infty) \cap P(0, F_\infty)$, and so $p_* \leq \bar{p}_{\underline{\pi}} < 1$ (see proof of Corollary 1). Because Proposition 3 also implies $F_\infty(x) = F_0(x) < 1$ for all $x \in [p_*, 1)$, it follows that $\max\ \operatorname{supp}\ F_\infty = 1$.

Fix any $(p, m) \in \operatorname{supp} H_\infty$. We show $p \in P(m, F_\infty)$, meaning $H_\infty$ is profit maximizing against $F_\infty$. To show this inclusion, recall that supports are lower hemicontinuous, and so, passing to a subsequence if necessary, we can find a $(p_n, m_n) \in \operatorname{supp} H_n$ that attains $(p, m)$ as its limit. Suppose first $m < 1$. Then, given any $m' \in (m, 1)$, an $N$ exists such that $m_n < m'$ for all $n > N$. Appealing again to lower hemicontinuity of supports and $\max\ \operatorname{supp}\ F_\infty = 1$, we have (passing to a further subsequence if necessary) that $\max \operatorname{supp}\ F_n > m'$ for all $n$ larger than some finite $N'$. Therefore, an $N''$ exists for which $m_n < \max \operatorname{supp}\ F_n$ for all $n > N''$. Lemma 22 implies that $P(m_n, F_n) = P^*(m_n, F_n)$ holds for such $n$. Upper hemicontinuity of $P^*$ (see Lemma 18) then gives that $p \in P^*(m, F_\infty)$. Suppose now $m = 1$. Appealing to Lemma 22, we then have that either $m_n > \max \operatorname{supp}\ F_n$ for all large $n$, or a subsequence $\{(p_{n_k}, m_{n_k})\}_{k\geq 0}$ exists such that $p_{n_k} \in P^*(m_{n_k}, F_{n_k})$ for all $k$. Either way, $p_n \to 1$: in the first case, $p_n \geq \max \operatorname{supp}\ F_n \to 1$ due to optimality of S's strategy, whereas in the second case, $(p_{n_k}, m_{n_k}) \to (p, 1)$, and so $p \in P^*(1, F_\infty)$ due to upper hemicontinuity of $P^*$. Hence, we have shown $\operatorname{supp} H_\infty \subseteq \operatorname{graph}\ P$, meaning $H_\infty$ is optimal for S against $F_\infty$.

To conclude the proof, observe that B's best response is upper hemicontinuous due to Berge's theorem, and so $F_\infty$ is optimal against $H_\infty$ at $\kappa = 0$. It follows that $(H_\infty, F_\infty)$ is a free-learning equilibrium. ■

Our next task is to prove the main result of this appendix section. To describe this result, define the following function,

$$F_{\underline{\pi}} : [0,1] \to \mathbb{R},$$
$$s \mapsto \begin{cases} G_{\underline{\pi},\bar{t}}(s) & \text{if } s \le \bar{p}_{\underline{\pi}} \\ F_0(s) & \text{if } s \ge \bar{p}_{\underline{\pi}}, \end{cases}$$

which is well defined because $\bar{p}_{\underline{\pi}}$ is $G_{\underline{\pi}}$-separating, and so $G_{\underline{\pi},\bar{t}}(\bar{p}_{\underline{\pi}}) = F_0(\bar{p}_{\underline{\pi}})$. The latter fact also implies $F_{\underline{\pi}}$ is a CDF. Moreover, $F_{\underline{\pi}} \in \mathcal{A}$. To see this inclusion, observe that for every $x \in [0,1]$,

$$I_{F_{\underline{\pi}}}(x) = \int_0^{x \wedge \bar{p}_{\underline{\pi}}} (F_0 - G_{\underline{\pi},\bar{t}})(s) \, \mathrm{d}s = I_{G_{\underline{\pi},\bar{t}}}(x \wedge \bar{p}_{\underline{\pi}}) \ge 0,$$

where the inequality follows from $G_{\underline{\pi},\bar{t}} \in \mathcal{A}$. Furthermore, at $x = 1$ the inequality holds with equality because $\bar{p}_{\underline{\pi}} \in S(G_{\underline{\pi},\bar{t}})$.

We now prove Theorem 6, which establishes that every sequence of costly learning equilibria converges to a Pareto worst free-learning equilibrium as learning becomes cheap. In this equilibrium, S's strategy is $\hat{H}_{\bar{p}_{\underline{\pi}}}$, and B's signal is $F_{\underline{\pi}}$.

**Theorem 6** *For $\kappa > 0$, let $(H_\kappa, F_\kappa)$ be any equilibrium in $\mathcal{E}_\kappa$. Then,*

$$\lim_{\kappa \to 0} (H_\kappa, F_\kappa) = (\hat{H}_{\bar{p}_{\underline{\pi},\bar{t}}}, F_{\underline{\pi}}) \in \mathcal{E}_0.$$

*Moreover, $\pi^*(0, F_{\underline{\pi}}) = \underline{\pi}$, and so $(\hat{H}_{\bar{p}_{\underline{\pi},\bar{t}}}, F_{\underline{\pi}})$ is a Pareto-worst free-learning equilibrium.*

**Proof.** Take any sequence $\{(H_n, F_n)\}_{n \ge 0}$ of $\kappa_n$-equilibria with $\kappa_n > 0$ where $\kappa_n \to 0$. Below, we show the equilibria sequence attains $(\hat{H}_{\bar{p}_{\underline{\pi},\bar{t}}}, F_{\underline{\pi}})$ as its limit. It follows that this limit is a free-learning equilibrium, by Lemma 30. Moreover, because $\bar{p}_{\underline{\pi}} = \min \operatorname{supp} \hat{H}^{\mathbf{p}}_{\bar{p}_{\underline{\pi}}}$, Lemma 18 delivers $\bar{p}_{\underline{\pi}} \in P^*(0, F_{\underline{\pi}})$. Since $F_{\underline{\pi}}$ is B-optimal under free learning, it also follows that $\bar{p}_{\underline{\pi}} \in S(F_{\underline{\pi}})$. Hence, $F_{\underline{\pi}}(\bar{p}-) = F_0(\bar{p}-) = F_0(\bar{p})$, and so

$$\pi^*(0, F_{\underline{\pi}}) = \Pi(\bar{p}_{\underline{p}}, 0, F_{\underline{\pi}}) = \Pi(\bar{p}_{\underline{\pi}}, 0, F_0) = \underline{\pi}.$$

Thus, all that remains is to establish that $\{(H_n, F_n)\}_{n \ge 0}$ converges to the desired limit. Because this sequence lives in a compact set, it is the union of convergent subsequences. Therefore, to show the desired convergence, it suffices to show that if the sequence converges to some $(H_\infty, F_\infty)$, then $(H_\infty, F_\infty) = (\hat{H}_{\bar{p}_{\underline{\pi},\bar{t}}}, F_{\underline{\pi}})$.

Towards showing the desired equality, observe $(H_\infty, F_\infty) \in \mathcal{E}_0$ by Lemma 30. Hence, $H_\infty = \hat{H}_{p_*}$ for $p_* = \min \operatorname{supp} H^{\mathbf{p}}_\infty$ by Proposition 3. The same proposition also implies

$p_* \in \text{supp } F$ (because $F_0$ is strictly increasing). Let $\underline{x} = \min \text{supp } F$. We now establish $[\underline{x}, p_*] \subseteq P(0, F_\infty)$. Note $p_* \in P(0, F_\infty)$ follows from Lemma 24. To show $[\underline{x}, p_*) \subseteq P(0, F_\infty)$, recall that supports are lower hemicontinuous, and so (passing to a subsequence if necessary), every $n$ admits some $x_n$ and $p_n$, both in supp $F_n$, such that $p_n \to p_*$ and $x_n \to \underline{x}$. Proposition 4 then delivers that, for every $\alpha \in (0, 1]$, $y_n = \alpha \underline{x}_n + (1 - \alpha) p_n$ lies in supp $H_n^{\mathbf{p}|\mathbf{m}}(\cdot|\mathbf{m} < \bar{x}_n)$, where $\bar{x}_n = \max \text{supp } F_n$. Therefore, an $m_n$ exists such that $y_n \in P^*(m_n, F_n)$. Passing to a subsequence if necessary, one can also attain that $m_n$ converges to some limit $m_\infty$. Because $P^*$ is upper hemicontinuous (Lemma 18), it follows that $\lim y_n = y_\alpha := \alpha \underline{x} + (1 - \alpha) p_* \in P^*(m_\infty, F_\infty)$. Because $y_\alpha < p_*$, Lemma 21 gives that $y_\alpha \in P(0, F_\infty)$. Thus, we have shown $[\underline{x}, p_*] \subseteq P(0, F_\infty)$. As such, it follows that $F_\infty(x-) = G_{\pi^*(0,F_\infty),1}(x-)$ for all $x \in [\underline{x}, p_*]$ (and $\underline{x} \geq \pi(0, F_\infty)$). Because $p_* < 1$, $G_{\pi^*(0,F_\infty),1}$ is continuous on $[\underline{x}, p_*]$, and so $F_\infty(x) = G_{\pi^*(0,F_\infty),1}(x)$ for all $x \in [\underline{x}, p_*)$. Recalling that $F_\infty$ is a best response against $\hat{H}_{p_*}$ gives that $p_*$ is $F_\infty$-separating, and so $F_\infty$ is continuous at $p_*$. It follows that $F_\infty = G_{\pi^*(0,F_\infty),1}$ on $[\underline{x}, p_*]$, and $\underline{x} = \pi(0, F_\infty)$. Combined with $p_*$ being $F_\infty$-separating, this delivers

$$I_{G_{\pi^*(0,F_\infty),1}}(p_*) = I_{F_\infty}(p_*) = 0.$$

Appealing to Lemma 10 then gives that $\pi^*(0, F_\infty) = \underline{\pi}$. It follows that $p_* = \bar{p}_{\underline{\pi}}$ and $G_{\underline{\pi},1}(x) = G_{\underline{\pi},\bar{t}}(x)$ for all $x \leq \bar{p}_{\underline{\pi}} < \bar{t}$. To summarize, we have shown that $H_\infty = \hat{H}_{\bar{p}_{\underline{\pi}}}$, $F_\infty(x) = G_{\underline{\pi},1}(x) = G_{\underline{\pi},\bar{t}}$ for all $x \leq \bar{p}_{\underline{\pi}}$, and $F_\infty(x) = F_0(x)$ for all $x \geq \bar{p}_{\underline{\pi}}$. In other words, we have shown $(H_\infty, F_\infty) = (\hat{H}_{\bar{p}_{\underline{\pi}},\bar{t}}, F_{\underline{\pi}})$, which completes the proof. ■