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

July 1, 2019

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, online resources enable buyers to learn about a mechanic's reputation, a contractor's ability, or an over-the-counter (OTC) asset's value. This information acquisition often takes

*We would like to thank Ben Brooks, Piotr Dworczak, Navin Kartik, Elliot Lipnowski, George Mailath, Stephen Morris, Philip Reny, Andrzej Skrzypacz, and Rakesh Vohra for extremely useful comments and suggestions.

†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

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. 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 contains less information than the valuation itself. Because 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 CDFs.

[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).

In particular, we postulate that the cost function is Fréchet differentiable.

Let us 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 will be argued, a cost function is monotonic whenever its Fréchet derivative, which is a function itself, is convex.[2] 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 CDF's support. Imposing this assumption in addition to monotonicity resembles stipulating that a strictly increasing function has a strictly positive derivative everywhere.

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 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

[2]See Machina (1982) for a similar result.

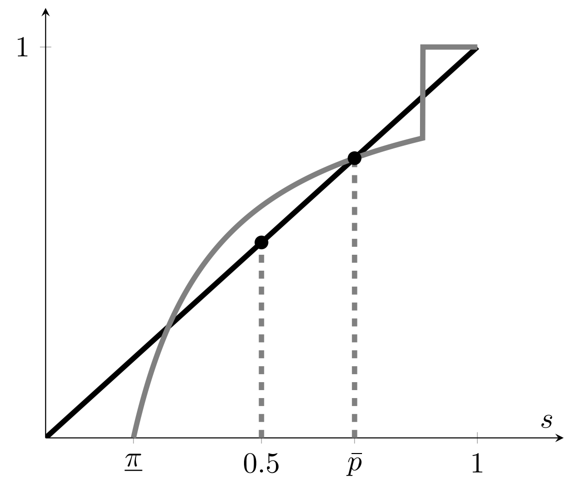

Figure 1: An Illustration of the best and worst equilibria in the uniform case.

equilibria, the buyer learns her valuation perfectly. 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 so 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 on 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 CDF. 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 compared to those in a full-information environment, leading to considerable welfare losses. Mitigating these losses may justify certain market features such as the existence of professional intermediaries. By making sure traders are informed, intermediaries can substantially increase social surplus. More broadly, our results highlight the importance of regulating the provision of product information even when data are cheap, because cheap data do not necessarily approximate full information. Special care should be taken when designing the information channels through which market participants learn. For example, mandatory information sessions appear to be more desirable than supplying brochures, because being able to know something is not the same as actually knowing it.

Our paper serves as a cautionary tale on interpreting recent papers characterizing consumer and producer surplus pairs that can arise as an equilibrium outcome for *some* information structure (e.g., Bergemann et al. (2015), Roesler and Szentes (2017)). Of particular relevance is Roesler and Szentes (2017), who consider a setting similar to ours in which the buyer's signal is observed by the seller before he sets a price. Their key result identifies the signal-equilibrium pair that maximize the buyer's payoff. The buyer-optimal signal turns out to be the same Pareto signal as in our worst free-learning equilibrium.[3]

[3]Pareto distributions also arise in robust auction design (e.g., Bergemann and Schlag (2008), Carrasco et al. (2018)). Of particular relevance is Du (2018), who studies an auctioneer's revenue guarantee from an exponential mechanism. In the single buyer case, he shows that his mechanism and the signal in our worst free-learning equilibrium form a saddle point in a zero-sum game between seller and nature.

At first glance, their result might seem surprising given that the worst free-information equilibrium minimizes the buyer's payoff. However, because the seller sets a price only after observing the buyer's signal in Roesler and Szentes's (2017) model, 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 him 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.

Our paper also adds to the recent literature on the relationship between free-learning equilibria and the vanishing-cost limits of equilibria. For example, Yang (2015) studies a 2-by-2 coordination game in which players can learn about their stochastic benefits from coordination. When the learning cost is proportional to entropy reduction, infinitely many equilibria can be attained in the limit. Morris and Yang (2016) consider a related regime-switching game and show a unique vanishing-cost limit exists if the learning cost admits a "continuous choice" property, that is, if only signals whose distribution varies continuously with the state can be optimal.[4] 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. Therefore, the vanishing-cost limit can be viewed as an equilibrium-selection device from a symmetric-information game. By contrast, learning is asymmetric in our model because the seller cannot acquire information about the buyer's valuation. Consequently, as we explained above, 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.

Costly consumer learning is extensively studied in the literature on rational inattention initiated by Sims (1998, 2003, 2006). In these models, information cost is proportional to the resulting expected reduction in entropy. For example, Matějka (2015) studies a dynamic pricing model with a consumer who is rationally inattentive to prices. The author finds that rational inattention leads to rigid pricing, because such pricing structures are easier for the consumer to assess. Ravid (2018) studies a dynamic, repeated-offer

[4] Denti (2018) shows that allowing players to learn about others' information yields a unique vanishing cost limit in Yang's (2015) model, whereas Hoshino (2018) argues the limits selected by Denti's (2018) model depend on the fine details of the cost function.

bargaining game in which the buyer is rationally inattentive and can learn about both her valuation and the seller's offers. He finds the buyer benefits from her inattention, and that such benefits remain large even when offers are frequent and costs vanish. 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.

Several papers examine buyers' incentives to acquire costly information about their valuations before participating in auctions. The buyers' learning strategies depend on the selling mechanism announced by the seller. Persico (2000) shows that if the buyers' signals are affiliated, they acquire more information in a first-price auction than in a second-price one. Compte and Jehiel (2007) show dynamic auctions tend to generate higher revenue than simultaneous ones. Shi (2012) also analyzes models in which it is costly for the buyers to learn about their valuations, and identifies the revenue-maximizing auction in private-value environments. In all of these setups, the seller is able to commit to a selling 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.[5]

Condorelli and Szentes (2018) also consider a bilateral trade model. In contrast to our setup, the distribution of the buyer's valuation is not given exogenously. Instead, the buyer chooses her value distribution and perfectly observes its realization. The seller observes the buyer's distribution but not her valuation and sets a price. The authors show that, as in our model, the equilibrium distribution generates a unit-elastic demand.

2 The Model

A seller, S, has an object to sell to a single buyer, B. 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 observe $\mathbf{v}$ 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.

[5] 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, Ganuza (2004), Bergemann and Pesendorfer (2007), and Ganuza and Penalva (2010).

Then S, without observing B's information-acquisition strategy and signal, makes a take-it-or-leave-it price offer, $p \in [0,1]$, which B accepts if and only if her expected valuation conditional on her signal weakly exceeds p.[6] Both players are risk-neutral expected-payoff maximizers.

Signal structures and B's action space. Note that both B's trading decision and her expected payoff from trading depend only on her posterior expectation, $\mathbb{E}[\mathbf{v}|\mathbf{s}]$. Assuming that acquiring more information[7] is more costly, restricting attention to signal structures for which B's posterior expectation is the signal itself – that is, $\mathbb{E}[\mathbf{v}|\mathbf{s}] = \mathbf{s}$, is without loss of generality. As a consequence, both B and S only care about the signal's marginal distribution. Thus, we identify each signal with the CDF of its marginal distribution. We let $\mathcal{F}$ denote the space of all CDFs over $[0,1]$, which we endow with the $\mathcal{L}_1$-norm, denoted by $\|\cdot\|$. [8] For any subset $A \subseteq [0,1]$, we take $\mathbf{1}_A$ to be the indicator function that is equal to 1 on A, and zero otherwise. Therefore, $\mathbf{1}_{[x,1]} \in \mathcal{F}$ is the CDF corresponding to a unit atom at $x \in [0,1]$. Given a CDF, $F \in \mathcal{F}$, we let $F(x-)$ be its left limit at x.[9]

Comparing the informativeness of different signals turns out to be useful. We say that $\mathbf{s}$ is more informative than $\mathbf{s}'$ if observing $\mathbf{s}$ is equivalent to observing $\mathbf{s}'$ and an additional signal $\mathbf{t}$, that is, $\mathbf{s} = \mathbb{E}[\mathbf{v}|\mathbf{s}', \mathbf{t}]$. Of course, one can assume $\mathbf{t}$ is just the difference between the two signals, so $\mathbf{s} = \mathbf{s}' + \mathbf{t}$. Furthermore, by the Law of Iterated Expectation, $\mathbb{E}[\mathbf{t}|\mathbf{s}'] = 0$.[10] In other words, $\mathbf{s}$ is a mean-preserving spread of $\mathbf{s}'$. Conversely, for any signal $\mathbf{s}$ whose distribution is a mean-preserving spread of F', a less informative signal $\mathbf{s}'$ exists whose distribution is F'.[11] Hence, if $F, F' \in \mathcal{F}$, we say that F is **more informative** than F' (denoted by $F \succeq F'$) if and only if F is a mean-preserving spread of F'; that is,[12]

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

The CDF F is said to be **strictly more informative** than F' (denoted by $F \succ F'$) if both $F \succeq F'$ and $F' \neq F$.[13]

[6] Assuming B trades if indifferent has no effect on our results but makes the analysis simpler.

[7] Using Blackwell's (1953) information ranking.

[8] 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).

[9] That is, $F(x-) = \sup F((-\infty, x))$.

[10] More precisely, $\mathbb{E}\left[\mathbf{t}|\mathbf{s}'\right] = \mathbb{E}\left[\mathbf{s} - \mathbf{s}'|\mathbf{s}'\right] = \mathbb{E}\left[\mathbb{E}\left[\mathbf{v}|\mathbf{s}', \mathbf{t}\right]|\mathbf{s}'\right] - \mathbf{s}' = \mathbb{E}\left[\mathbf{v}|\mathbf{s}'\right] - \mathbf{s}' = 0$.

[11] See Gentzkow and Kamenica (2016) and references therein.

[12] See Rothschild and Stiglitz (1970) for the statement and Leshno et al. (1997) for the corrected proof. Blackwell and Girshick (1979) proves the result for discrete distributions.

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

We allow B to choose *any* signal to learn about $\mathbf{v}$, and we identify B's action space with the set of those CDFs that correspond to a signal about her valuation. Of course, observing the valuation perfectly is more informative than any signal. Thus, B can choose any CDF F that is less informative than the prior F_0, that is, any $F \in \mathcal{F}$ such that $F_0 \succeq F$. We denote this set by $\mathcal{A}$, and refer to CDFs in $\mathcal{A}$ as signals. Letting $I_F(x)$ denote $\int_0^x (F_0 - F)\ \mathrm{d}s$, (1) implies $F \in \mathcal{A}$ if and only if $I_F(x) \geq 0$ for all x and $I_F(1) = 0$.

The cost of information acquisition. Information acquisition is costly. In general, different information structures generating the same distribution of posterior expectation 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 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 also require the function C to be sufficiently smooth. More precisely, we assume C is **Fréchet differentiable**; that is, it is continuous, and for each $F \in \mathcal{A}$, a Lipschitz function, $c_F : [0,1] \to \mathbb{R}$, exists such that for every $F' \in \mathcal{A}$,

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

where o is a function that equals zero at zero and $\lim_{x \searrow 0} [o(x)/x] = 0$. We refer to c_F as C's **derivative at** F.[14]

The assumption that acquiring more information is more costly is natural. We say that C is **monotone** if $C(F) \geq C(F')$ whenever F is strictly 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 convex and Fréchet differentiable. Then, C is monotone if and only if c_F is convex for each $F \in \mathcal{A}$.*

Proof. See appendix. ■

For the intuition behind the claim and for better understanding the concept of Fréchet differentiablility, let us restrict attention to signals whose support lies in a finite set, say,

[14] Formally, $c_F(x) = \int_0^x \phi_F\ \mathrm{d}s$ for some $\phi_F \in \mathcal{L}_\infty[0,1]$, and so c_F is unique Lebesgue-a.e.

[15] See Machina (1982) for a related result.

$\{s_1, \ldots, s_N\}$. Then, each $F \in \mathcal{F}$ can be represented by the vector in the n-dimenional simplex $(\alpha_1, \ldots, \alpha_N) \in \Delta^n$ for which $F = \sum_{n=1}^{N} \alpha_n \mathbf{1}_{[s_n,1]}$. In this case, the function C is a mapping from Δ^n to $\mathbb{R}$, and the Fréchet derivative at F 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.

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* $\mathrm{co}(\mathrm{supp}\ F)$.

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, $F \in \mathcal{A}$.[16] If B's signal is F, S's expected payoff from the random price H is given by

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

We denote S's maximal profit by $\pi_F := \max_{p \in [0,1]} \Pi(p, F)$ and the set of profit-maximising prices by $P(F) = \arg\max_{p \in [0,1]} \Pi(p, F)$[17] In Appendix B, we establish continuity of S's maximal profit and upper hemicontinuity of the profit-maximizing prices, $P(\cdot)$.[18]

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

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

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 $\mathcal{A}$.

[16] We can assume B uses a pure strategy because C is convex, S's objective is linear, and $\mathcal{A}$ is convex.

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

[18] Notice S's profit is only upper semicontinuous, and so said properties do not follow from Berge's Maximum Theorem.

Because B's best response and S's (mixed) best response are upper hemicontinuous, nonempty, convex and compact valued, an equilibrium exists by Kakutani's Fixed-Point Theorem.[19]

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

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

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 characterizes their Fréchet derivatives.

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

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

Then, C's Fréchet derivative equals c for all F.

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

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

[19] 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_∞. Because $\mathcal{F}$ is compact, it is enough to show H_∞ is an S best response to F_∞, 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_∞ only if a sequence $p_n \in$ supp H_n exists that attains p_∞ 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_∞, only if $p_\infty \in P(F_\infty)$, that is, supp $H_\infty \subseteq P(F_\infty)$.

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 convex for all F, and strictly convex for any $F \neq \mathbf{1}_{[\bar{v},1]}$.

Example 3. (Quadratic Costs) Let $c : [0,1] \times [0,1] \to \mathbb{R}_+$ be some strictly convex, 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[20]

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

is convex and Fréchet differentiable, with the derivative being given by the strictly convex 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 B and S, respectively. We then 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, $F = F_0$. Later, we also show that the worst equilibrium outcome is attainable with a Pareto signal.

3.1 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

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

[20]Example 3 is essentially the functional form for quadratic preferences as introduced by Machina (1982).

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 an 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 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 shows the argument of this paragraph can be extended to the case where S randomizes over prices. We show that by choosing F, B achieves the same payoff as with perfect learning if and only if S only charges F-separating prices.

Lemma 1 *The signal F is a best response against H if and only if* $\operatorname{supp} H \subseteq S(F)$.

Proof. 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 $\int_p^1 (F_0 - F)\ \mathrm{d}s = I_F(p)$. Because $F \in \mathcal{A}$ and $I_F(\cdot)$ is continuous, we conclude that F generates the same payoff as perfect learning if and only if $I_F(p) = 0$ for all $p \in \operatorname{supp} H$, that is, $\operatorname{supp} H \subseteq S(F)$. ■

Next, we show the graphs of F and F_0 must intersect at any F-separating price. Intuitively, 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.

Lemma 2 *If $F \in \mathcal{A}$ and $p \in S(F)$, then F is continuous at p and*

$$F(p) = F_0(p). \tag{5}$$

Proof. Suppose $p \in S(F)$. Then, 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 follows. ■

3.2 The Seller's Best Responses

We now characterize 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, then 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 π. Otherwise, for each $p \in [\pi, 1]$, the cumulative probability, y, which guarantees profit π is $1 - \pi/p$. Observe that $1 - \pi/p$ is the CDF corresponding to the Pareto distribution parameterized by π. Becaue $p \leq 1$, we conclude that the iso-profit curve of the seller corresponding to profit π is essentially identical to the truncated Pareto distribution, $G_{\pi,1}$.

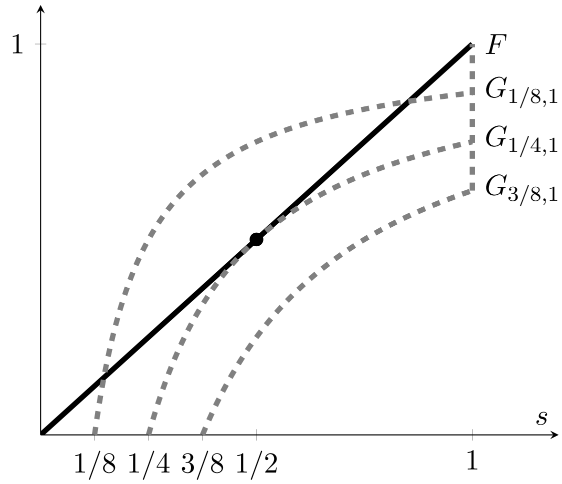

Figure 2: The seller's best response against the uniform distribution.

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 are $\{(p, F(p-)) : p \in [0,1]\}$. Therefore, S's profit is defined by the largest π, 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 largest 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. Because iso-profit curves are strictly increasing, the CDF F must also be strictly increasing at any point of tangency, and hence any such points must lie in the support of F. In Figure 2, only a single point of tangency at $p = 0.5$ exists. The following lemma summarizes these observations.

Lemma 3 *Fix any $F \in \mathcal{A}$. Then,*

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

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

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

Proof of Lemma 3. To prove part (i), note S's profit from setting a certain price cannot exceed π_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 that $s \in P\,(F)$ 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 that $P\,(F) \subseteq \text{supp } F$. Suppose, by contradiction, that a p exists such that $p \in P\,(F) \backslash \text{supp } F$. 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 with setting price p' than price p, a contradiction to $p \in P\,(F)$. ■

3.3 Free-Learning Equilibrium Characterization

We now turn to characterizing the set of free-learning equilibrium payoffs. We begin by showing 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 π_F even if B learns perfectly instead of getting signal F. To state this result precisely, for each π, let X_π be the set of prices that yield profit π under F_0, that is, $X_\pi := \{p : \Pi\,(p,F_0) = \pi\}$. We first explain that, in any free-learning equilibrium, (H,F),

$$\text{supp } H \subseteq X_{\pi_F}.$$

To see this inclusion, note that any equilibrium price must be profit maximizing as well as F separating (see Lemma 1), and hence supp $H \subseteq P(F) \cap S(F)$. Therefore, it is enough to show that

$$P(F) \cap S(F) \subseteq X_{\pi_F}. \tag{7}$$

To explain this inclusion, we consider a price p that is both profit maximizing under F and F separating, and explain that p generates profit π_F under perfect learning, that is, $p \in X_{\pi_F}$. Because $p \in S\,(F)$, B's demand is the same under F and under perfect learning, see Lemma 2. Consequently, p generates the same profit irrespective of whether B's signal is F or F_0. Furthermore, this profit is π_F because $p \in P\,(F)$, and hence $p \in X_{\pi_F}$.

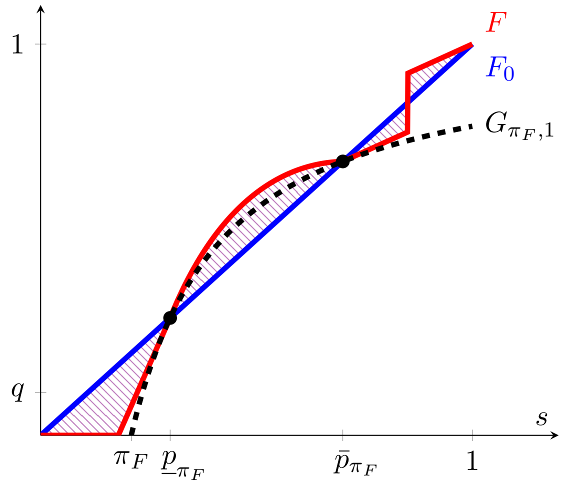

Figure 3: An illustration of Lemma 4. The blue line corresponds to the prior, F_0, the red curve is the the signal, F, and the dashed curve is the π_F iso-profit curve, G_{π_F}. Although the signal is such that both prices in $X_{\pi_F} = \{\underline{p}_{\pi_F}, \bar{p}_{\pi_F}\}$ are profit maximizing, only $\bar{p}_F$ can be separating.

The next lemma states that S's equilibrium price must the largest element of X_{π_F}. Before we state this result, note that because F_0 is regular, the function $\Pi(\cdot, F_0)$ is strictly concave, so X_π contains at most two such prices for every π. Because $\Pi(\cdot, F_0)$ is continuous, it attains any value between 0 and π_{F_0}.[21] Therefore, for each $\pi \in [0, \pi_{F_0}]$, X_π 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 $\bar{p}_{\pi_F}$ is the unique price that can be both profit maximizing and F separating.

Lemma 4 *Let (H, F) be a free-learning equilibrium. Then,* supp $H = \{\bar{p}_{\pi_F}\}$.

Proof. See the Appendix. ■

For an explanation, recall that X_π has at most two elements. If X_{π_F} is a singleton, the statement of the lemma immediately follows from the observation that supp $H \subseteq P(F) \cap S(F)$ and (7). Suppose now that X_{π_F} is binary, that is, $X_{\pi_F} = \{\underline{p}_{\pi_F}, \bar{p}_{\pi_F}\}$ and $\underline{p}_{\pi_F} < \bar{p}_{\pi_F}$. Figure 3 illustrates this case and depicts the prior, F_0, the signal, F, and the π_F-iso-profit curve, $G_{\pi_F,1}$. By Lemma 2 and the definition of X_{π_F},

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

these three curves intersect at $\underline{p}_{\pi_F}$ and $\bar{p}_{\pi_F}$. We now argue that

$$\int_{\underline{p}_{\pi_F}}^{\bar{p}_{\pi_F}} F_0\left(s\right)\ \mathrm{d}s < \int_{\underline{p}_{\pi_F}}^{\bar{p}_{\pi_F}} G_{\pi_F,1}\left(s\right)\ \mathrm{d}s \leq \int_{\underline{p}_{\pi_F}}^{\bar{p}_{\pi_F}} F\left(s\right)\ \mathrm{d}s.$$

The first inequality follows from the observation that the strict concavity of $\Pi(\cdot, F_0)$ implies $\Pi(\cdot, F_0)$ is strictly larger than π_F on $\left(\underline{p}_{\pi_F}, \bar{p}_{\pi_F}\right)$, so $F_0 < G_{\pi_F,1}$ on this interval. The second inequality follows from the fact that S's maximal profit is π_F if B's signal is F, therefore, by part (i) of Lemma 3, $F \geq G_{\pi_F,1}$. An immediate consequence of this inequality chain is that $I_F(\underline{p}_{\pi_F}) - I_F(\bar{p}_{\pi_F}) = \int_{\underline{p}_{\pi_F}}^{\bar{p}_{\pi_F}} \left(F - F_0\right)(s)\ \mathrm{d}s > 0$. Because $F \in \mathcal{A}$, $I_F(\bar{p}_{\pi_F}) \geq 0$, so it must be that $I_F(\underline{p}_{\pi_F}) > 0$, meaning $\underline{p}_{\pi_F}$ is not F separating.

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,[22]

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

Theorem 1 shows S's minimal and maximal equilibrium profits are $\underline{\pi}$ and π_{F_0}, respectively, and that S can attain any profit in between. If S's equilibrium profit is π, 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 $u = \int_{\bar{p}_\pi}^{1} \left(s - \bar{p}_\pi\right) \mathrm{d}F_0\left(s\right)$.*

Proof. See the Appendix. ∎

The "only if" part of this theorem implies that in a free-learning equilibrium, S can never attain a profit above his full-information profit. This result is a straightforward consequence of Lemma 4. Recall that this lemma states that if B's signal is F, the equilibrium price is the largest price that generates profit π_F under perfect learning, $\overline{p}_{\pi_F}$. But if learning is perfect, S can achieve π_{F_0} by setting the optimal price instead of $\overline{p}_{\pi_F}$, showing that $\pi_F \leq \pi_{F_0}$. 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 $\bar{p}_{\pi_F}$. This conclusion follows from the facts that S sets price $\overline{p}_{\pi_F}$ in every equilibrium where his profit is π_F

[22]The minmax is well defined because $\mathcal{A}$ 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).

(see Lemma 4) 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 π. Existence of an equilibrium with profit $\underline{\pi}$ follows from the equilibrium payoff set being closed.[23] Figure 4 illustrates our construction, which obtains an equilibrium by applying two modifications to the π-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 π. 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. This CDF, however, fails to be a signal, due to having too large of a mean. To make the CDF into a signal, we reduce the mean using the second modification: truncating the distribution at the top at some value t. The result is a signal corresponding to the red curve, $G^q_{\pi,t}$, in Figure 4. Noting the truncation point t is larger than p means p still yields S a profit of π, 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')*,*

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

Proof. We prove the corollary by showing that $\bar{p}_\pi$ is strictly decreasing in π 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 π 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

[23] 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.

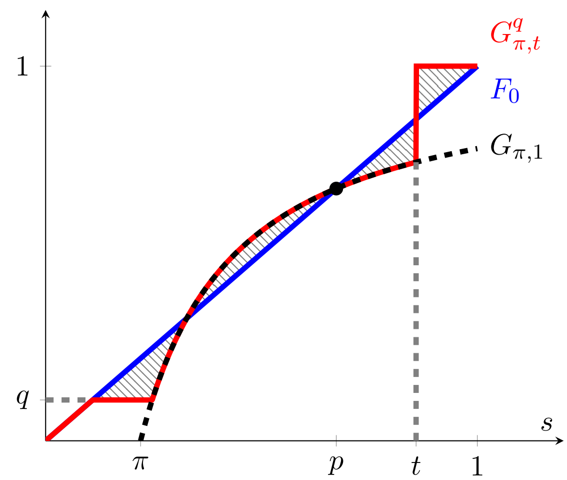

Figure 4: A constructed free-learning equilibrium, $(\mathbf{1}_{[p,1]}, F)$.

correspond to lower prices and therefore higher B utility. We now show that $\bar{p}_\pi$ 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_π 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 ϵ. 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. ∎

We have thus shown that when learning is free, our model admits a continuum of equilibria, all of which can be Pareto ranked. In the next section, we discuss the shape of equilibria when learning is costly, and show that as costs vanish, the equilibrium must converge to a Pareto-worst free-learning equilibrium.

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.

Proof. See the Appendix. ■

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.

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,

$$\text{co}(\text{supp } F) \subseteq \text{supp } H. \tag{8}$$

If this inclusion 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 fact, we show that if (x, y) is maximal among such intervals, x, y must both lie in supp F. So, to prove (8), it is enough to show supp $H \cap (x, y) \neq \emptyset$ if $x, y \in$ supp F. Suppose first that F places atoms at both x and y. Then, B can profitably deviate by bunching together all the signals x and y; that is, instead of observing these signals, she only learns the 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 (8), the second follows from the observation that S never sets a price that is not a possible signal realization (see part (ii) of Lemma 3). This chain of inclusion implies part (i) of the theorem.

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, $G_{\underline{\pi},\bar{t}}$.[24]

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

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

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 the 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.

The proof of this theorem is based on connecting our analysis of costly learning with our observations regarding free-learning equilibria. When costs are positive, B uses a Pareto signal (see Proposition 1). Because the set of Pareto signals is closed, she must also be using a Pareto signal in the limit, say, $G_{\pi,t}$. In turn, when learning is free, S must set a $G_{\pi,t}$-separating price in the support of $G_{\pi,t}$ (see Lemma 1 and part (ii) of Lemma 3). The key step in the proof, which we explain in detail in the next paragraph, is to show a Pareto signal that has a non-empty set of separating prices in its support is associated

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

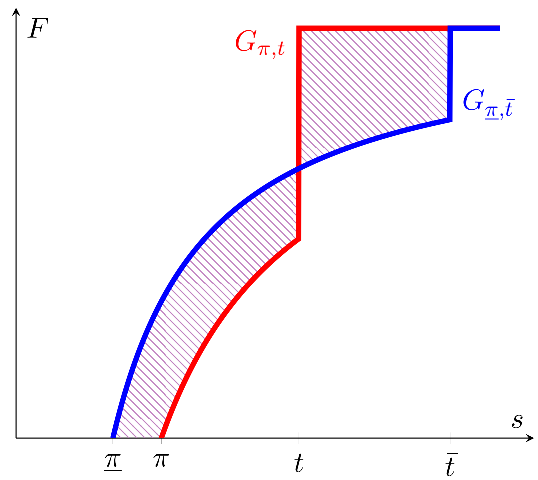

Figure 5: The minmax Pareto signal, $G_{\underline{\pi},\bar{t}}$, and another Pareto signal, $G_{\pi,t}$, with $\pi > \underline{\pi}$.

with the minmax profit, $\underline{\pi}$. To conclude the theorem, we note that if S's profit is $\underline{\pi}$, he must charge $\bar{p}_{\underline{\pi}}$ by Lemma 4.

Let us return to the key step of the proof and explain that the only Pareto signal that has a separating price in its support is $G_{\underline{\pi},\bar{t}}$. We first observe that the support of each Pareto signal, $G_{\pi,t} \in \mathcal{A}$, with $\underline{\pi} < \pi$, is contained in the support of $G_{\underline{\pi},\bar{t}}$, that is, $[\pi, t] \subset [\underline{\pi}, \bar{t}]$. The reason is that the mean of a Pareto signal, $G_{\pi,t}$, is strictly increasing in both π and t. Because the mean of each Pareto signal is $\overline{v}$ $\left(= \int v \, \mathrm{d}F_0(v)\right)$ and $\underline{\pi} < \pi$, it follows that $t < \bar{t}$. We now argue that the information constraint of $G_{\underline{\pi},\bar{t}}$ is point-wise tighter than that of $G_{\pi,t}$ over $G_{\pi,t}$'s support. In other words, we demonstrate that for all $x \in \operatorname{supp} G_{\pi,t}$,

$$0 < I_{G_{\pi,t}}(x) - I_{G_{\underline{\pi},\bar{t}}}(x) = \int_0^x \left[G_{\underline{\pi},\bar{t}}(s) - G_{\pi,t}(s)\right] \, \mathrm{d}s. \tag{9}$$

The right-hand side is just the area between the CDFs $G_{\pi,t}$ and $G_{\underline{\pi},\bar{t}}$ on $[0, x]$. Figure 5 illustrates these CDFs and the area between them. Note that this area is zero for all $x \in [0, \underline{\pi}]$ and strictly increasing over $[\underline{\pi}, t]$. Therefore, the area must be strictly positive for all $x \in [\pi, t] = \operatorname{supp} G_{\pi,t}$; that is, (9) holds. Because $G_{\underline{\pi},\bar{t}} \in \mathcal{A}$, $I_{G_{\underline{\pi},\bar{t}}}(x) \geq 0$, and so $I_{G_{\pi,t}}(x) > 0$ must hold for for all $x \in \operatorname{supp} G_{\pi,t}$. Hence, $G_{\pi,t}$ has no separating prices in its support.

We now turn to proving Theorem 2.

Proof of Theorem 2. Let $\{\kappa_n\}_{n\geq 0}$ be a strictly positive sequence that converges to zero, and take $\{(H_n, F_n)\}_{n\geq 0}$ to be a corresponding sequence of equilibria. Because $\mathcal{F}$ and $\mathcal{A}$ are both compact, $\{(H_n, F_n)\}_{n\geq 0}$ 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 that $(H_\infty, F_\infty) = (\mathbf{1}_{[\bar{p}_{\underline{\pi}},1]}, G_{\underline{\pi},\bar{t}})$.

To this end, we first note that because B's objective is a continuous function of (κ, H, F), B's best-response correspondence is upper hemicontinuous in (κ, 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 1. That H_∞ is optimal for S against F_∞ follows from upper hemicontinuity of S's mixed-best-response correspondence, $F \mapsto \arg\max_{H\mathcal{F}} \Pi(H, F)$.[25] Thus, the limit (H_∞, F_∞) is a free-learning equilibrium. Because the Pareto signal set is closed and F_∞ is the limit of Pareto signals (Proposition 1), we have that F_∞ is itself a Pareto signal; that is, $F_\infty = G_{\pi,t}$ for some π and t. Below, we argue that $\pi = \underline{\pi}$, 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 4.

All that remains is to show that $\pi = \underline{\pi}$. Suppose for a contradiction that $\pi > \underline{\pi}$. Because $(H_\infty, G_{\pi,t})$ is a free-learning equilibrium, the support of H_∞ contains only prices that are both profit maximizing and separating, meaning $P(G_{\pi,t}) \cap S(G_{\pi,t}) \neq \emptyset$. Take any $p \in P(G_{\pi,t}) \cap S(G_{\pi,t})$. On the one hand, p is separating, and so $I_{G_{\pi,t}}(p) = 0$. On the other hand, $p \in P(G_{\pi,t}) \subseteq \text{supp } G_{\pi,t} = [\pi, t]$ (by Lemma 3 Part (ii)) and so

$$
\begin{aligned}
I_{G_{\pi,t}}(p) = \int_0^p (F_0 - G_{\pi,t})\ \mathrm{d}s &= \int_0^\pi F_0\ \mathrm{d}s + \int_\pi^p (F_0 - (1 - \pi/s))\ \mathrm{d}s \\
&> \int_0^\pi (F_0 - G_{\underline{\pi},\bar{t}})\ \mathrm{d}s + \int_\pi^p (F_0 - (1 - \underline{\pi}/s))\ \mathrm{d}s = I_{G_{\underline{\pi},\bar{t}}}(p) \geq 0,
\end{aligned}
$$

where the inequality follows from $G_{\underline{\pi},\bar{t}}(s) > 0$ for all $s \in (\underline{\pi}, \pi)$, and the last equality follows from $G_{\underline{\pi},\bar{t}}(s) = 1 - \underline{\pi}/s$ for all $s \in [\underline{\pi}, \bar{t}] \supset [\pi, t]$.[26] Thus, we have shown that $I_{G_{\pi,t}}(p) = 0 < I_{G_{\pi,t}}(p)$ – a contradiction. It follows that $\pi = \underline{\pi}$, completing the proof. ■

[25] See footnote 19.

[26] Recall that a Pareto signal's truncation point is decreasing with its associated profit.

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 $p - c$, where $c \in (0, 1)$. For $c \in (0, \bar{v})$, our analysis goes through with the c-**shifted truncated Pareto** signal,

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

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

For $c \geq \bar{v}$, 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 $c > 0$, 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 a c-shifted truncated Pareto. But when $c \geq \bar{v}$, no informative signal can have a c-shifted truncated Pareto distribution.[27] 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 describe what happens if S has a random production cost with full support in $[0, 1]$ that is independent of B's valuation. S privately observes the cost realization, c, before setting a price. Then, his

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

utility from trade at price p is $p - c$, where c 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 $c = 0$. This price is offered for all values of c 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 c, 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 $c = 0$. 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. 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 that 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 that any mechanism is equivalent to setting a random price; see, 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.[28] When learning is free, equilibrium requires S's price to be separating, and the full-information outcome remains profit maximizing regardless of the prior. Similarly, the regularity of the prior plays no role in showing that B uses a Pareto signal when learning is costly, and the same holds in the costless limit. Because the

[28]Our results hold without change when the buyer's prior is 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 that there is a unique free-learning equilibrium. Such uniqueness arises if and only if $\pi_{F_0} = \underline{x}$.

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 profi-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.

References

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

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

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.

Blackwell, David (1953), "Equivalent Comparisons of Experiments." *The Annals of Mathematical Statistics*, 24, 265–272.

Blackwell, David A and Meyer A Girshick (1979), *Theory of Games and Statistical Decisions*. Courier Corporation.

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

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.

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

Condorelli, Daniele and Balazs Szentes (2018), "Information Design in the Hold-Up Problem." *Working Paper*.

Denti, Tommaso (2018), "Unrestricted Information Acquisition." *Working Paper*.

Du, Songzi (2018), "Robust mechanisms under common valuation." *Econometrica*, 86, 1569–1588.

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

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

Leshno, Moshe, Haim Levy, and Yishay Spector (1997), "A Comment on Rothschild and Stiglitz's "Increasing Risk: I. A Definition"." *Journal of Economic Theory*, 77, 223–228.

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.

Morris, Stephen and Ming Yang (2016), "Coordination and Continuous Choice." *Working Paper*.

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

Ravid, Doron (2018), "Bargaining with Rational Inattention." *Working Paper*.

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

Rothschild, Michael and Joseph E Stiglitz (1970), "Increasing Risk: I. A Definition." *Journal of Economic Theory*, 2, 225–243.

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

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 for every F, $w, z \in \operatorname{int}(\operatorname{co}(\operatorname{supp} F))$, 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 5 *Fix some $F \in \mathcal{F} \setminus \{\mathbf{1}_{[x,1]} : x \in [0,1]\}$, let $[x, x'] = \operatorname{co}(\operatorname{supp}\ F)$, and take $\bar{w} = \int s \ \mathrm{d}F$. 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 that $F_{\lambda,0} \succeq \mathbf{1}_{[\bar{w},1]}$ for all $\lambda > 0$ because $\lambda y + (1-\lambda)x_\lambda = \bar{w}$. We now show that $F_{\lambda,\beta} \succeq F_{\lambda,0}$ 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 $\bar{s} \in [0,1]$,

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

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_λ being continuous in λ 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

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

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

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

and so φ is continuous in the product topology. Therefore,

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

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 $\int_0^{\bar{s}} (F - \mathbf{1}_{[\bar{w},1]})(s) \, \mathrm{d}s = \int_x^{\bar{s}} F(s) \, \mathrm{d}s > 0$. Similarly, for all $s \in [\bar{w}, x')$, $F(s) < 1 = \mathbf{1}_{[\bar{w},1]}(s)$. As such, if $\bar{s} \in [\bar{w}, x^*]$, $\int_{\bar{s}}^1 (1 - F(s)) \, \mathrm{d}s > 0 = \int_{\bar{s}}^1 (1 - \mathbf{1}_{[\bar{w},1]}(s)) \, \mathrm{d}s$, and so $\int_{\bar{s}}^1 (F - \mathbf{1}_{[\bar{w},1]})(s) \, \mathrm{d}s < 0$. Since $\int_0^1 (F - \mathbf{1}_{[\bar{w},1]})(s) \, \mathrm{d}s = 0$, we obtain $\int_0^{\bar{s}} (F - \mathbf{1}_{[\bar{w},1]})(s) \, \mathrm{d}s > 0$ 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 φ^*, one must then have $\varphi^*(\lambda, \beta) > 0$ for all $\lambda, \beta > 0$ small enough. Fixing any such λ and β, we now show that $\int_0^{\bar{s}} (F - F_{\lambda,\beta}) \, \mathrm{d}s \geq 0$ for all $\bar{s}$ by considering three cases. First, if $\bar{s} \in [x_*, x^*]$,

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

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

$$\begin{aligned}\int_0^{\bar{s}} (F - F_{\lambda,\beta}) \, \mathrm{d}s &= \int_0^{x^*} (F - F_{\lambda,\beta}) \, \mathrm{d}s + \int_{x^*}^{\bar{s}} (F - 1) \, \mathrm{d}s \\ &\geq \int_0^{x^*} (F - F_{\lambda,\beta}) \, \mathrm{d}s + \int_{x^*}^1 (F - 1) \, \mathrm{d}s = \int_0^1 (F - F_{\lambda,\beta}) \, \mathrm{d}s = 0,\end{aligned}$$

in view of $\operatorname{supp} F_{\lambda,\beta} \subseteq [x_*, x^*]$ and $F_{\lambda,\beta} \succeq \mathbf{1}_{[\bar{w},1]}$. We have therefore shown that for all sufficiently small λ and β, $\int_0^{\bar{s}} (F - F_{\lambda,\beta}) \, \mathrm{d}s \geq 0$ 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 $F' \succeq F$. Because C is convex, we have that

$$C(F') - C(F) \geq \int c_F \ \mathrm{d}(F' - F) \geq 0,$$

where the last inequality follows from c_F being convex. Hence, C is monotone.

Suppose now that C is monotone. Fix any $w, y, z \in \mathrm{co}(\mathrm{supp}\ F_0)$ such that $y = \alpha w + (1-\alpha)z$ for some $\alpha \in (0,1)$. Because c_F is only unique Lebesgue almost everywhere (see footnote 14), we may as well assume $w, y, z \in \mathrm{int}\,(\mathrm{co}(\mathrm{supp}\ F_0))$. Our task is to show $c_F(y) \leq \alpha c_F(w) + (1-\alpha)c_F(z)$.

By Lemma 5, 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 6 *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_∞ as its limit. We show $\lim_{n\to\infty} \pi_{F_n} = \pi_{F_\infty}$. Because Π is upper semicontinuous, $F \mapsto \pi_F$ is also upper semicontinuous.[29] 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 ϵ 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_∞ as its limit. Because Π 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 4: Free-learning equilibrium prices

If (H, F) is an equilibrium, F is a best response to H, and hence, by Lemma 1, supp $H \subseteq S(F)$. Furthermore, because H is a best response to F, each price in the support of H must be profit maximizing, that is, supp $H \subseteq P(F)$. Therefore, it is enough to prove $P(F) \cap S(F) = \{\bar{p}_{\pi_F}\}$. We have already shown that $P(F) \cap S(F) \subseteq X_{\pi_F}$; see equation (7). Thus, it remains to be shown that if $p \in X_{\pi_F}$ but $p < \bar{p}_{\pi_F}$ then $p \notin S(F)$.

To this end, note that for all $s \in (p, \bar{p}_{\pi_F})$, it must be that

$$G_{\pi_F,1}(s) > F_0(s). \tag{10}$$

The reason is that because F_0 is regular, the profit function $\Pi(\cdot, F_0)$ is strictly concave, and hence any price between p and $\bar{p}_{\pi_F}$ generates a profit strictly above $\pi_F (= \Pi(p, F) = \Pi(p, F_0))$. Thus, F_0 is strictly below the π_F-iso-profit curve at these prices; that is, (10) holds. Now, observe that

$$I_F(p) = I_F(\bar{p}_{\pi_F}) - \int_p^{\bar{p}_{\pi_F}} (F_0 - F)\, ds \geq I_F(\bar{p}_{\pi_F}) - \int_p^{\bar{p}_{\pi_F}} (F_0 - G_{\pi_F,1})\, ds > I_F(\bar{p}_{\pi_F}) \geq 0,$$

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

where the first inequality follows from part (i) of Lemma 3, the strict inequality follows from (10), and the last inequality is implied by $F \in \mathcal{A}$. Thus, we have shown that $I_F(p) > 0$, and hence $p \notin S(F)$.

D 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 4 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 $\underline{\pi} \le \Pi(H, F) \le \pi_{F_0}$ whenever (H, F) is a free-learning equilibrium. Because $\underline{\pi} \le \Pi(H, F)$ by definition of $\underline{\pi}$, it remains to be shown that $\Pi(H, F) \le \pi_{F_0}$. To do so, notice that because $\operatorname{supp} H \subseteq S(F)$, we have by Lemma 2 that $F(p-) \ge F_0(p-)$ for every $p \in \operatorname{supp} H$. Because H maximizes S's profit, S's profit must be the same from all prices in $\operatorname{supp} H$. We therefore have that for any $p \in \operatorname{supp} H$,

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

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,[30] it is sufficient to show that every profit $\pi \in (\underline{\pi}, \pi_{F_0})$ can be generated by some equilibrium.[31] Fix such a π, and define for $q \in [0, 1]$ and $t \in [\pi, 1]$ the following CDF:

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

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

[30] See footnote 23.

[31] 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 π_{F_0}.

Lemma 7 *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 $G^{q^*}_{\pi,1}(\hat{x}) = G_{\pi,1}(\hat{x}) \geq q^*$.*

Before providing the lemma's proof, let us 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 $x > t$,

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

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 π. 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 7, which we do now.

D.1 Proof of Lemma 7

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

Lemma 8 *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 8), and a truncated Pareto signal 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 9 *$I_{G_{\pi,1}}(x) \geq 0$ for all x, with a strict inequality whenever $x > \underline{x}$.*

Proof. Note that $\pi > \underline{\pi}$ 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 that 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 3 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 10 *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 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 $x^* \in A$, $\min\ I_{G^{F_0(x^*-)}_{\pi,1}}(A) \leq I_{G^{F_0(x^*-)}_{\pi,1}}(x^*) = 0$. Thus, we have shown that $\min\ I_{G^{F_0(x^*-)}_{\pi,1}}(A) \leq 0 = \min I_{G^0_{\pi,1}}(A)$, as required. The proof is now complete. ■

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

Lemma 11 *For all $x \in [0,1]$, $I_{G^{q^*}_{\pi,1}}(x) \geq 0$.*

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 that $\inf I_{G^{q^*}_{\pi,1}}([0,\pi)) \geq 0$. To see this, recall that $\pi \geq \underline{x}$, meaning $x < \pi$ only if $G_{\pi,1}(x) = 0$. As such, whenever $x < \pi$,

$$
G^{q^*}_{\pi,1}(x) = \max\{0, \min\{q^*, F_0(x)\}\} = \min\{q^*, F_0(x)\} \leq F_0(x).
$$

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

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

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

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

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

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

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

To conclude the proof of Lemma 7, notice that $x \in A$ only if $G_{\pi,1}(x) \geq F_0(x-) \geq F_0(x^*-) \geq q^*$. Taking $x_1 \in \arg\min_{x \in A} I_{G^{q^*}_{\pi,1}}(x)$, we therefore have

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

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

E 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 3, our task is to show co(supp F) $\subseteq$ supp H.

Letting $[w, z] :=$ co(supp F), we wish to show that $[w, z] \subseteq$ 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 that 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 $(x', y') \supsetneq (x, y)$.[32] 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 (x, y). An analogous argument gives $y \neq z$ only if $y \in$ supp H. Hence, we have shown $x, y \in \{w, z\} \cup$ supp H. Because supp $H \subseteq$ supp F (Lemma 3) and $\{w, z\} \subseteq$ supp F, we thus have that $x, y \in$ supp F.

We now construct a family of deviations indexed by $\epsilon > 0$, F^*_ϵ, 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 $x, y \in$ supp F:

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

[32] 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').

Moreover, take

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

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

$$s_\epsilon = \frac{1}{2} \int s \ \mathrm{d}\left(F_{1,\epsilon} + F_{2,\epsilon}\right),$$
$$\eta_\epsilon = \min\{\beta_{1,\epsilon}, \beta_{2,\epsilon}\} > 0,$$
$$F_\epsilon^* = \beta_{0,\epsilon} F_{0,\epsilon} + \eta_\epsilon \delta_{s_\epsilon} + (\beta_{1,\epsilon} - 0.5\eta_\epsilon) F_{1,\epsilon} + (\beta_{2,\epsilon} - 0.5\eta_\epsilon) F_{2,\epsilon}.$$

In words, F_ϵ^* takes $0.5\eta_\epsilon$ mass from the ϵ-ball around x and $0.5\eta_\epsilon$ mass from the ϵ-ball around y and pools them to create an $\eta_\epsilon > 0$ mass on s_ϵ. Because $0.5(F_{1,\epsilon} + F_{2,\epsilon}) \succ \delta_{s_\epsilon}$, F_ϵ^* 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_ϵ^*; 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

$$\lim_{\epsilon \searrow 0} \int \frac{T_H}{\eta_\epsilon} \ \mathrm{d}(F - F_\epsilon^*) = 0, \tag{11}$$
$$\lim_{\epsilon \searrow 0} \left(\frac{C(F_\epsilon^*) - C(F)}{\eta_\epsilon} \right) < 0, \tag{12}$$

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

$$0 \leq \lim_{\epsilon \searrow 0} \frac{U_\kappa(H,F) - U_\kappa(H, F_\epsilon^*)}{\eta_\epsilon} = \lim_{\epsilon \searrow 0} \left[\int \frac{T_H}{\eta_\epsilon} \ \mathrm{d}(F - F_\epsilon^*) + \kappa \frac{C(F_\epsilon^*) - C(F)}{\eta_\epsilon} \right] < 0, \tag{13}$$

hence completing the proof.

We now explain why (11) and (12) 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

$$T_H(s) = \int_0^s (s-p) \ \mathrm{d}H(p) = \int_0^x (s-p) \ \mathrm{d}H(p) = H(x)s - \int_0^x p \ \mathrm{d}H(p). \tag{14}$$

As such, T_H is affine over $[x, y]$, and so (11) obtains as follows:

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

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 (12). 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

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

Thus, we have (11) and (12), which together yield the contradiction (13), which completes the proof.