Consumer Profiling via Information Design*

Itay P. Fainmesser† Andrea Galeotti‡ Ruslan Momot§

Ludovic Renou¶

This Draft: October 26, 2026.

Abstract

A platform profiles consumers using their purchase histories and selectively discloses information to successive short-lived sellers. By shaping sellers' prices, the platform generates new purchase data and improves its profiles. With short-lived consumers, the platform can learn valuations below a threshold but cannot distinguish valuations above it. The threshold is the lowest price that sellers would not want to raise for any consumer, even if they knew that all consumers would accept it. With long-lived consumers, the platform can learn every valuation. We also characterize attainable combinations of consumer and producer surplus as the platform becomes arbitrarily patient.

*For helpful comments and discussions, we thank Nageeb Ali, Heski Bar-Isaac, Dirk Bergemann, Darrell Duffie, Alexander Frug, Ben Golub, Sanjeev Goyal, Ali Khan, Andrew Koh, David Myatt, Steve Tadelis, Omer Tamuz, Quitzé Valenzuela-Stookey, and seminar and conference participants at Queen's University, Pennsylvania State University at Harrisburg, Tel Aviv University, University of Michigan, University of Toronto, University of Georgia, Stanford University, Boston University, the U.S. Naval Academy, Caltech, UC Berkeley, Duke University, the Second Federal Trade Commission Conference on Marketing and Public Policy, the 2024 ESIF Economics and AI+ML Meeting, and the Data, Markets, and Privacy Conference at The Johns Hopkins University. This project was supported by ERC Grant 101142008 (Andrea Galeotti).

†The Johns Hopkins Carey Business School and The Department of Economics, The Johns Hopkins University, Baltimore, MD 21202 (e-mail: itay_fainmesser@jhu.edu)

‡Department of Economics, London Business School, London, NW1 4SA, United Kingdom (e-mail: agaleotti@london.edu)

§Technology & Operations, Ross School of Business, University of Michigan, Ann Arbor, MI 48109 (e-mail: momot@umich.edu)

¶Department of Economics, Arizona State University, USA (e-mail: lrenou.econ@gmail.com)

1. Introduction

Data collection in digital marketplaces has raised concerns among policymakers and antitrust authorities about platforms' ability to profile consumers and use this information to shape market outcomes in ways that may reduce social welfare. Platforms collect data on consumer behavior and decide what information to share with sellers. Amazon's Brand Analytics, for example, provides participating sellers with aggregate information about customer searches, purchase patterns, and demographic characteristics.[1] Such information can help sellers choose which consumers to target, which products to present to them—a practice often referred to as product steering—and what prices to charge.[2] A platform's choices about what sellers learn can therefore influence the offers consumers encounter, the trades that occur, and the distribution of gains from trade.

The information available to a platform, however, is itself shaped by these offers. Purchase records do not reveal consumers' valuations directly: what can be inferred from a purchase or rejection depends on the offer the consumer faced. Moreover, offers that maximize a seller's profits need not generate the information most useful for learning about demand. By controlling the information available to sellers, a platform can influence both current offers and what subsequent purchase decisions reveal. Its disclosure choices thus help determine the information the platform learns over time.

To study these interactions, we develop a model in which a platform initially lacks information about consumers' valuations and learns from their purchase decisions. Sellers choose prices, while the platform controls the information they receive about consumer demand. We characterize what the platform can learn through selective disclosure and how

[1] Amazon describes these tools as helping sellers select and position products, tailor advertising, and identify opportunities for bundling. See Amazon, *Brand Analytics*, https://sell.amazon.com/tools/amazon-brand-analytics.

[2] Empirical evidence documents how platform-provided information affects sellers' decisions and outcomes. Expedia's Rev+ provides hotels with market and revenue insights to guide pricing decisions. Bar-Gill et al. (2024) find that access to eBay's seller analytics increases transactions, revenues, and service quality. Chen et al. (2026) show that access to market data on Taobao leads retailers to expand their product assortments.

the resulting knowledge affects trade and the distribution of surplus.

In our model, each consumer has a characteristic that determines his valuation of the product. We think of a characteristic as a vector of demographic and socioeconomic attributes, such as age, sex, race, education, or income. Before trade begins, a valuation is drawn for each characteristic. All consumers with that characteristic share this valuation, and valuations may be affiliated across characteristics.[3] Such dependence captures common demand factors—such as product quality, seasonality, or fashion—alongside horizontal differences. At the beginning of each period, a short-lived seller enters the platform with no information about previous offers, purchasing decisions, or consumers' characteristics. The platform, which observes all these, sends a message to the seller. The seller then makes an offer to the consumer, who accepts or rejects it. Consumers are short-lived in Section 3 and long-lived in Section 4.

We provide two general insights. The first concerns learning. Because sellers are short-lived, they have no incentive to experiment with prices. To learn about consumers' valuations, the platform must therefore use information design to change the demand sellers expect to face. Information design can do this in two ways. First, by selectively concealing past offers, purchase decisions, and consumers' characteristics, the platform can change sellers' beliefs about the distribution of valuations. With short-lived consumers, this is the only available channel, and we show that the platform can learn valuations below a threshold but cannot distinguish valuations above it (Theorem 1). This threshold is the lowest price that sellers find optimal to offer to any consumer when they know that all consumers' valuations are at least that price. With one characteristic, the threshold equals the static monopoly price. Second, with long-lived consumers, information design can also influence consumers' willingness to pay today by changing their opportunity cost of waiting. It does so by shaping their expectations of future offers. We show that this additional channel allows the platform to learn every valuation to any desired precision (Theorem 2).

The second insight concerns welfare. In the case of a single characteristic, consumer

[3]High valuations for some characteristics shift the conditional distributions of the remaining valuations upward.

profiling allows the platform to recover the entire deadweight loss of static monopoly and achieve efficient trade. How the gains from trade can be distributed, however, depends on what the platform can learn. With short-lived consumers, the platform can learn only valuations below the static monopoly price. This implies that consumers and sellers must each retain at least their static monopoly surplus, so only the recovered deadweight loss can be divided arbitrarily between them. With long-lived consumers, the platform can learn every valuation. In this case, we show that all gains from trade can be divided arbitrarily between consumers and sellers (Theorem 3).

In the remainder of the introduction, we describe the main features of the information designs that achieve these learning outcomes and then discuss how we construct the attainable combinations of consumer and producer surplus.

With short-lived consumers, the information design supports learning in three ways. First, the platform shapes sellers' beliefs about the valuation associated with a characteristic by selectively concealing the purchase decisions of past consumers with that characteristic. This allows it to learn valuations below the corresponding monopoly price to arbitrary precision. Second, the platform can extend learning by concealing consumers' characteristics. Suppose characteristic c' has a higher monopoly price than characteristic c. By leaving sellers uncertain whether the current consumer has characteristic c or c', the platform can induce offers above the monopoly price for c. Observing how consumers with characteristic c respond to these offers allows the platform to learn their valuation up to the monopoly price for c'. Finally, affiliation between the valuations associated with different characteristics allows learning to extend further. Evidence of strong demand from consumers with both characteristics can make sellers more optimistic about demand from consumers with either characteristic. By selectively disclosing past acceptances of high offers, the platform can induce prices above the higher of the two initial monopoly prices and learn additional valuations.

We next consider long-lived consumers: each consumer remains active until he purchases and is then replaced by a new consumer. We focus on the case of one characteristic. Following

the classical literature on bargaining with one-sided incomplete information (Fudenberg, Levine, and Tirole, 1985, Ausubel, Cramton, and Deneckere, 2002), we restrict attention to equilibria in which consumers' strategies are stationary and monotone.[4] These restrictions provide a natural basis for comparison with short-lived consumers, whose strategies are also stationary and monotone.

The classical literature on bargaining with one-sided incomplete information shows that stationary-equilibrium prices converge to the lowest possible valuation as bargaining frictions vanish. This might suggest that purchase decisions provide little scope for learning about consumers' valuations. In contrast, we show that, with sufficiently patient consumers, the platform can learn every valuation to arbitrary precision. Along the equilibrium path, prices rise in small steps, so the minimum valuation needed to accept each offer also rises gradually. Each acceptance reveals a lower bound on the consumer's valuation, while the first rejection reveals a nearby upper bound. Hence, after the first rejection the platform identifies the valuation with high precision. The information design makes the prescribed prices optimal for sellers by shaping the prices consumers expect after a rejection. Suppose a seller deviates by charging more than the prescribed price. If the consumer rejects, the platform conceals earlier purchases from the next seller and tells her that the first consumer rejected the initial offer and that all subsequent offers were made at a very low price. Believing the consumer's valuation to be low, this seller offers a low price to the same consumer. Anticipating this lower future offer, the consumer is less willing to accept a higher current price, making a price increase unprofitable for the seller.

Finally, we characterize attainable consumer and producer surplus in the one-characteristic environment as the platform becomes patient; with long-lived consumers, the platform and consumers share the same discount factor. In the benchmark where the platform knows consumers' valuations, Bergemann, Brooks, and Morris (2015) show that every feasible surplus pair is attainable provided sellers receive at least their static monopoly profit. In our setting

[4] A consumer's acceptance rule depends only on his valuation, the public belief when he enters the market, and the current price; higher prices are accepted with lower probability.

with short-lived consumers, since consumers' valuations above the monopoly price cannot be learned, consumer surplus must also remain at least at its static monopoly level. However, learning valuations below that price allows the platform to recover the entire monopoly deadweight loss. By disclosing the learned valuations, the platform can allocate this recovered surplus entirely to sellers. Alternatively, adapting the consumer-optimal information design of Bergemann, Brooks, and Morris (2015) allows it to allocate the recovered surplus entirely to consumers.

With long-lived consumers, the platform can learn every valuation and disclose this information to sellers, allowing them to price discriminate and capture all gains from trade. Although learning takes longer as consumers become more patient, the discounted loss of surplus during the learning phase vanishes. The platform can also allocate all gains from trade to consumers by withholding information about past offers and purchases. Under this design, there is a stationary equilibrium with a constant price that converges to zero as consumers become patient. Trade therefore becomes efficient, and consumers receive all gains from trade in the limit.

Section 2 introduces the model. Sections 3 and 4 characterize learning outcomes when consumers are short-lived and long-lived, respectively. Section 5 characterizes attainable surplus outcomes. Section 6 concludes by relating our findings to the literature on dynamic persuasion and on information and price discrimination.

2. Model

We consider a platform-mediated marketplace where consumers and sellers meet and trade. The platform observes all transactional data and commits to a disclosure policy that governs the information available to sellers. The platform uses disclosure to induce prices that generate information about consumers' valuations. We characterize the information the platform can learn about these valuations. Our baseline model assumes that consumers are also short-lived;

we relax this assumption in Section 4.

Time is discrete, $t = 1, 2, \ldots$. In each period t, a new seller with one unit of a good enters the marketplace and meets a new consumer with unit demand. The seller offers the price $p_t \in \mathbb{R}_+$, and the consumer either accepts ($a_t = 1$) or rejects the offer ($a_t = 0$).

Consumers. The newly entered consumer has one of finitely many characteristics $c \in \mathcal{C} := \{1, \ldots, C\}$. We think of characteristics as vectors of socio-economic attributes, which determine the consumer's valuation $v_c^* \in V := [0, 1]$ for the good, a realization of the random variable v_c. All consumers with characteristic c value the item at v_c^*. Consumers know their characteristics and their valuations. If offered price p, a consumer with characteristic c obtains $v_c^* - p$ from purchasing and zero from rejecting.

Assumption 1 (Short-lived consumers)**.** *A consumer with characteristic c accepts price p if and only if $v_c^* \geq p$.*

Since consumers are short-lived, the only substantive requirement that Assumption 1 imposes is that consumers purchase when indifferent. The profile of valuations $\mathbf{v}^* = (v_c^*)_{c \in \mathcal{C}}$ is realized before period 1. We denote the associated random vector by $\mathbf{v}$ and its distribution by F, with density f. The consumer's characteristic c_t at period t is drawn independently across periods, with $\lambda_c > 0$ the probability of characteristic $c_t = c$. (It is also independent of $\mathbf{v}$.)

Assumption 2 (Affiliation and log-concavity)**.** *The density f is continuous and strictly positive on V^C, log-concave, and affiliated, i.e., for every $(\mathbf{v}, \mathbf{v}') \in V^C \times V^C$,*

$$f(\mathbf{v} \wedge \mathbf{v}')f(\mathbf{v} \vee \mathbf{v}') \geq f(\mathbf{v})f(\mathbf{v}'), \tag{1}$$

where $\wedge$ and $\vee$ denote the coordinate-wise minimum and maximum.

When valuations associated with different characteristics are independent, (1) holds with equality.

Sellers. At period t, the newly entered seller observes only the calendar time and a

message m_t from the platform. If the seller offers the price p_t and the consumer accepts, their payoff is p_t. Otherwise, it is zero.

Assumption 3 (Short-lived sellers)**.** *After receiving the message m_t, the period-t seller posts a price that maximizes their expected period-t profit.*

Platform. The platform observes the current characteristic c_t and the history of characteristics, messages, prices, and purchase decisions. Formally, let $H^0 := \{\emptyset\}$ and define recursively

$$H^t := H^{t-1} \times \mathcal{C} \times M_t \times \mathbb{R}_+ \times \{0, 1\}.$$

Before play begins, the platform commits to an information design $\mathcal{M} = (\mathcal{M}_t)_{t \geq 1}$, where

$$\mathcal{M}_t : H^{t-1} \times \mathcal{C} \longrightarrow \Delta(M_t).$$

The policy is common knowledge. Conditional on (h^{t-1}, c_t), the platform draws the seller's message from $\mathcal{M}_t(\cdot \mid h^{t-1}, c_t)$. This formulation permits full disclosure, no disclosure, or partial disclosure of the current characteristic and past transactions.

A belief system for the period-t seller maps each message into a posterior over the current characteristic, the past history, and the valuation profile:

$$\mu_t : M_t \longrightarrow \Delta(\mathcal{C} \times H^{t-1} \times V^C).$$

Assumption 4 (Consistent seller beliefs)**.** *At every period t, the seller's posterior is derived by Bayes' rule whenever possible.*

Induced Game and Learning. The information design $\mathcal{M}$ induces a multi-stage game $\Gamma_{\mathcal{M}}$, where at each stage/period t, the seller observes the message m_t, posts a price p_t, and the consumer either accepts ($a_t = 1$) or rejects ($a_t = 0$). A perfect Bayesian equilibrium of $\Gamma_{\mathcal{M}}$ is a strategy profile and a system of beliefs satisfying Assumptions 1, 3, and 4.

An equilibrium of $\Gamma_{\mathcal{M}}$ and the distributions F and λ induce a stochastic process over

valuations, characteristics, messages, prices, and purchase decisions, and hence the platform's posterior over the valuation profile after every on-path history.

Definition 1 (Platform learning)**.** *Fix $\varepsilon > 0$ and sets $V_c^* \subseteq V$ for every $c \in \mathcal{C}$. The platform* ***can ε-learn*** *valuations in $\prod_{c\in\mathcal{C}} V_c^*$ if there exist an information design $\mathcal{M}$ and a perfect Bayesian equilibrium of $\Gamma_{\mathcal{M}}$ such that, for every realized valuation profile $\mathbf{v}^* \in \prod_{c\in\mathcal{C}} V_c^*$ and every characteristic $c \in \mathcal{C}$, the platform's posterior probability that $v_c \in [v_c^* - \varepsilon, v_c^* + \varepsilon] \cap V$ converges to one almost surely.*

With short-lived consumers, an accepted offer at price p gives a lower bound p on the valuation, and a rejected offer at price p gives an upper bound p. The platform learns by inducing prices that bring these bounds sufficiently close.[5] If the disclosure policy is probabilistic, different sequences of messages may induce different sequences of prices. The definition requires that for almost all such sequences, the prices and purchasing decisions (almost) reveal the consumer's valuation.[6] Finally, the definition requires that the same design and equilibrium make it possible to learn about all valuations in a given range.

Discussion. We now comment on some of the assumptions. First, the assumption that sellers are short-lived eliminates their incentives to experiment for future profit and allows us to study the extent to which information design alone can induce experimentation. Second, we assume that consumers arrive one at a time, with characteristics independently and identically drawn across periods. This assumption simplifies the notation, since each seller posts only one price in each period. All our results extend to multiple consumers arriving in each period and to some persistence in their characteristics.[7]

Third, much of the analysis extends immediately to environments in which sellers also observe the characteristics of newly arriving consumers. The only change is that the platform

[5]When consumers are forward-looking, consumers may reject prices even though their valuations exceed prices. Yet, as we will see, the same logic will apply.

[6]More precisely, the set of terminal histories at which the platform's posterior assigns probability one to $[v_c^* - \varepsilon, v_c^* + \varepsilon] \cap V$ has probability one, when the valuation is v_c^* – the probability measure is the one induced by the equilibrium.

[7]For example, the characteristics of newly arriving consumers can follow an irreducible and aperiodic Markov chain. The core requirement is that consumers with each characteristic enter the market infinitely often.

loses its ability to pool consumers' characteristics and this reduces the range of valuations that the platform can learn.[8]

Fourth, we assume that consumers with the same characteristic share a common valuation. More generally, characteristics could predict willingness to pay while allowing for heterogeneity within consumers with the same characteristic. In that case, Nature would draw a valuation distribution for each characteristic, and the platform would hold a prior over these distributions. The valuations of the successive consumers with a given characteristic would be independent draws from the relevant distribution. Repeated offers at a given price would then allow the platform to estimate demand at that price. Our model is the special case in which every possible within-group distribution concentrates on a single valuation.

Finally, the main analysis abstracts from price competition between sellers. In Online Appendix OA.2, we extend the model to Hotelling price competition with one characteristic and short-lived consumers.

3. Learning with short-lived consumers

With short-lived consumers, the platform can influence the optimal prices of a seller only by changing her beliefs about consumers' valuations. These beliefs must be mixtures of the valuation distributions associated with different histories and characteristics. Our main result characterizes the most the platform can learn in this way: valuations below a common threshold can be learned arbitrarily precisely, while those above it cannot be distinguished. We first define the threshold, then state the result and explain the intuition.

[8] In footnote 9, we explain how our main result, Theorem 1, is modified.

3.1. The learning threshold

Set $p^0 = 0$ and, for $k \geq 1$, define

$$p_c^k := \arg\max_{p \geq p^{k-1}} p \Pr\left(v_c \geq p \,\middle|\, v_{c'} \geq p^{k-1} \text{ for every } c' \in \mathcal{C}\right), \tag{2}$$

$$p^k := \max_{c \in \mathcal{C}} p_c^k. \tag{3}$$

At each step, we calculate the monopoly price for each characteristic conditional on all valuations being at least the preceding threshold, and take the highest of these prices as the next threshold. Since $p^{k-1} \leq p_c^k \leq p^k \leq 1$, all characteristic-specific price sequences converge to a common limit:

$$p^\infty := \lim_{k \to \infty} p^k = \lim_{k \to \infty} p_c^k \qquad \text{for every } c \in \mathcal{C}. \tag{4}$$

In words, p^∞ is the lowest price that sellers find optimal to offer to any consumer when they know that all consumers' valuations are at least p^∞. With one characteristic c, p^∞ equals the static monopoly price p_c^1. With independently drawn valuations, $p^\infty = \max_c p_c^1$. Affiliation can raise p^∞ further, because high valuations for other characteristics make a high valuation for the current consumer more likely.

Example 1 (Correlated valuations)**.** *Suppose* $|\mathcal{C}| = 2$ *and, for* $K > 0$*, let*

$$f(v_1, v_2) \propto \exp[-K \max\{v_1 - v_2, 0\}], \qquad (v_1, v_2) \in [0,1]^2.$$

The density satisfies Assumption 2, with higher K *implying stronger affiliation. Figure 1 shows how the two static monopoly prices* p_1^1 *and* p_2^1 *and the learning threshold* p^∞ *generated by (2)–(4) vary with* K*. As* $K \downarrow 0$*,* v_1 *and* v_2 *become independent uniform draws on* $[0,1]$*, so all three prices converge to* $1/2$*. As* $K \to \infty$*, the joint distribution of* (v_1, v_2) *converges to the uniform distribution over the triangle* $0 \leq v_1 \leq v_2 \leq 1$*, so the static monopoly prices converge to* $1/3$ *and* $1/\sqrt{3}$*, respectively, while* p^∞ *converges to one.*

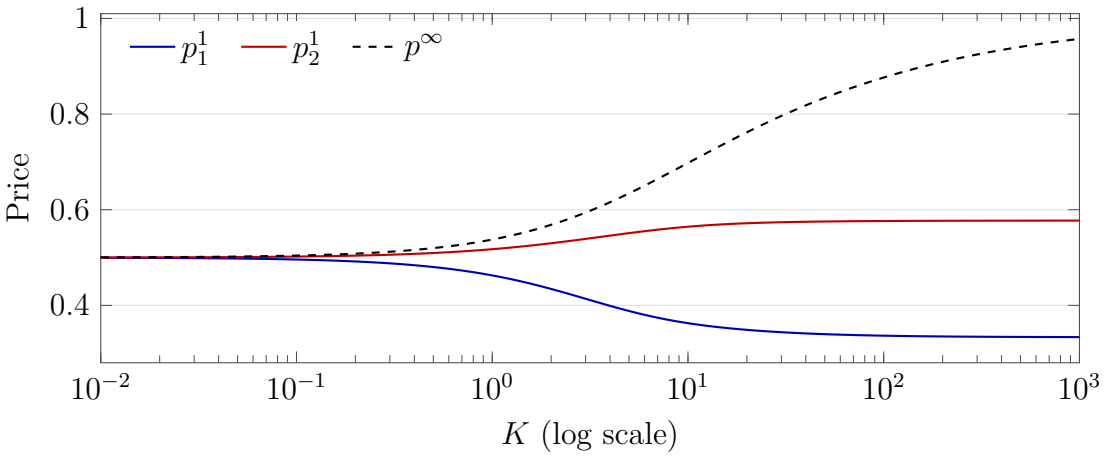


Figure 1: Static monopoly prices and the learning threshold.

3.2. Main result and intuition

The following theorem characterizes the scope for learning through information design when consumers are short-lived.

Theorem 1 (Learning threshold)**.**

(i) ***No learning above*** p^∞***.*** *For every information design, every induced equilibrium, and every characteristic c, any two valuation profiles that coincide in all coordinates except c, with both values of coordinate c in $[p^\infty, 1]$, induce the same distribution over platform histories.*

(ii) ***Learning below*** p^∞***.*** *For every $\varepsilon > 0$, the platform can ε-learn all valuations in $[0, p^\infty - \varepsilon)^C$.*

In the remainder of this section, we develop the intuition behind Theorem 1.[9] With one characteristic, we explain why learning is limited above the threshold and how concealing histories permits arbitrarily precise learning below it. We then show how concealing characteristics and affiliation each expand the range of valuations that can be learned.

[9]If sellers observe consumers' characteristics, an analogous statement of Theorem 1 applies, with the only difference that the learning threshold is now defined for each $c \in \mathcal{C}$ as $p_c^\infty := \lim_{k\to\infty} \tilde{p}_c^k$, where $\tilde{p}_c^0 = 0$ and, for $k \geq 1$, $\tilde{p}_c^k := \arg\max_p p \Pr\big(v_c \geq p \,\big|\, v_{c'} \geq \tilde{p}_{c'}^{k-1} \text{ for every } c' \in \mathcal{C}\big)$.

3.2.a. *One characteristic.* Suppose all consumers have the same characteristic c and the prior over their common valuation is uniform on V. The learning threshold is therefore $p^\infty = 1/2$. The first seller charges $1/2$, which consumers accept if $v^* \geq 1/2$. In period 2, the platform can only induce mixtures of two conditional valuation distributions: uniform on $[0, 1/2]$ following rejection and uniform on $[1/2, 1]$ following acceptance. Consider the following information design. The platform sends $(R, 1/2)$ whenever the first offer is rejected. When it is accepted, the platform sends $(A, 1/2)$ with probability ϕ and $(R, 1/2)$ with probability $1 - \phi$.

A seller receiving $(A, 1/2)$ has the most optimistic available posterior, uniform on $[1/2, 1]$, yet still charges $1/2$. This remains an upper bound on prices in subsequent periods, so all valuations above it generate identical purchasing decisions and cannot be learned.

A seller receiving $(R, 1/2)$ assigns probability $1/(2 - \phi)$ to rejection and probability $(1 - \phi)/(2 - \phi)$ to concealed acceptance. Her posterior is the corresponding mixture of the two uniform distributions, and her optimal price is

$$p_R(\phi) = \frac{2 - \phi}{4}.$$

As ϕ rises from zero to one, this price falls continuously from $1/2$ to $1/4$. By selecting ϕ, the platform can therefore induce any price in this interval.

By gradually increasing ϕ over time, the platform can make prices fall in arbitrarily small steps. An intermediate acceptance does not stop the design: the platform continues to use that history in subsequent pooling. Once prices reach $1/4$, the platform adapts the design by pooling histories in which the initial offer of $1/2$ was rejected and the offer of $1/4$ was accepted with histories in which all offers down to $1/4$ were rejected. This corresponds to garbling uniform valuation distributions on $[1/4, 1/2]$ and $[0, 1/4]$, respectively. Applying the same design allows prices to fall from $1/4$ to $1/8$ in small steps. The first acceptance following rejection locates the valuation between two successive prices; continued rejections eventually locate it below an arbitrarily small price. For any $\varepsilon > 0$, this design allows the

platform to ε-learn every valuation below $1/2$ within a finite number of periods.

3.2.b. *Two independent characteristics.* The platform can extend learning beyond characteristic-specific static monopoly prices by concealing characteristics. To see this, suppose that, in addition to characteristic c, there is another characteristic c', whose valuation is drawn independently of v_c from a prior with a static monopoly price of $3/4$. With characteristics disclosed, the preceding design allows the platform to learn valuations below $1/2$ for c and below $3/4$ for c'. Concealing characteristics extends learning for c into $[1/2, 3/4)$.

Consider a message $(c', A, 3/4)$, stating that the current consumer has characteristic c' and that the initial offer of $3/4$ to that characteristic was accepted. The platform sends this message with probability ϕ when these statements are true. It also sends it with probability ψ when the current consumer has characteristic c and the initial offer of $1/2$ to that characteristic was accepted. Otherwise, it reveals the true characteristic and purchase history.

A seller receiving $(c', A, 3/4)$ therefore faces a mixture of two valuation distributions: uniform on $[1/2, 1]$ for characteristic c, and the distribution of $v_{c'}$ conditional on $v_{c'} \geq 3/4$. Let θ be the seller's posterior probability that the current consumer has characteristic c. For $\theta > 0$, her optimal price is[10]

$$p(\theta) = \min\left\{\frac{3}{4}, \frac{1+\theta}{4\theta}\right\}.$$

By adjusting ϕ and ψ, the platform can vary θ and induce any price between $1/2$ and $3/4$.

To learn valuations in this interval, the platform repeatedly cycles through finitely many pairs (ϕ, ψ), chosen to induce prices between $1/2$ and $3/4$ with gaps smaller than ε. Each grid price is eventually offered to consumers with characteristic c with probability one. Their responses locate every $v_c^* \in [1/2, 3/4)$ between an accepted and a rejected price less than ε apart.

[10] For $p \in [1/2, 3/4]$, expected profit is $p[(1-\theta) + 2\theta(1-p)]$. The first-order condition $1 + \theta - 4\theta p = 0$ gives $p^* = (1+\theta)/(4\theta)$. For $0 < \theta \leq 1/2$, profit is increasing on this interval, so the optimal price is $3/4$.

3.2.c. *Affiliated valuations.* Now allow the valuations of c and c' to be affiliated, while their static monopoly prices remain $1/2$ and $3/4$. The platform can still use the preceding designs to learn valuations below $3/4$ for both characteristics. Affiliation can extend learning further because purchases by consumers with one characteristic provide information about the valuation associated with the other.

To do so, the platform reveals the current consumer's characteristic and sends a message $(\text{All}, A, 3/4)$, stating that earlier offers of $3/4$ were accepted by consumers of both characteristics. It sends this message whenever both acceptances occurred and with a small probability ϕ otherwise. When it does not send this message, it reveals the actual purchase history.

For small ϕ, a seller receiving $(\text{All}, A, 3/4)$ assigns probability close to one to both valuations being at least $3/4$. Her optimal price is therefore close to the corresponding conditional monopoly price, p_c^2 or $p_{c'}^2$, defined by (2) for $k = 2$. Affiliation can make these prices strictly higher than $3/4$. The platform then applies the preceding learning design to each characteristic and conceals characteristics to extend learning for both toward $p^2 = \max\{p_c^2, p_{c'}^2\}$.

Repeating this construction with messages (All, A, p^k) makes the common learning threshold approach $p^\infty = \lim_{k\to\infty} p^k$. At this threshold, even a seller who knows that both valuations are at least p^∞ finds it optimal to charge p^∞, whichever characteristic she faces. This remains an upper bound on prices, so valuations above it cannot be distinguished.

3.3. Comparison with full disclosure

To illustrate the role of information design, we compare the learning outcome in Theorem 1 with what the platform learns under full disclosure. Under full disclosure, each seller observes the current consumer's characteristic and the complete history of consumers' characteristics, offers, and purchasing decisions. Proposition 6 in Online Appendix OA.1 provides a full characterization of learning under this benchmark.

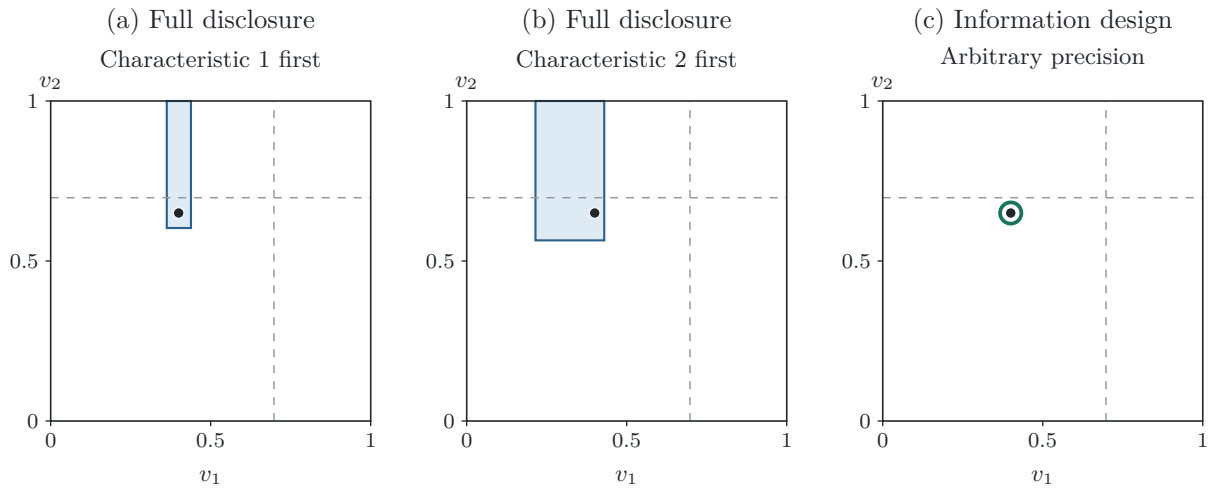


Figure 2: Full disclosure and information design in Example 1 for $K = 10$. Shaded rectangles show the limiting posterior supports: $[0.3628, 0.4384) \times [0.6030, 1]$ in Panel (a), when characteristic 1 arrives first, and $[0.2149, 0.4298) \times [0.5645, 1]$ in Panel (b), when characteristic 2 arrives first. Dots mark the true profile $(0.40, 0.65)$, and dashed lines mark $p^\infty \simeq 0.698$. Panel (c) shows the limit as learning precision improves under information design.

Consider first one characteristic, with valuation distribution F. Set $\widehat{p}_0 = 1$ and define

$$\widehat{p}_t = \underset{p \in [0, \widehat{p}_{t-1}]}{\arg\max}\, p\Big[F(\widehat{p}_{t-1}) - F(p)\Big], \qquad t \geq 1. \tag{5}$$

Here, $\widehat{p}_1$ is the static monopoly price, and $\widehat{p}_t$ is the optimal price when the seller knows that $v < \widehat{p}_{t-1}$. Under full disclosure, for any $v > 0$, prices fall along this sequence until the first acceptance. If $\widehat{p}_t$ is the first accepted offer, the platform's posterior is F conditional on $[\widehat{p}_t, \widehat{p}_{t-1})$. Every subsequent seller charges $\widehat{p}_t$, so learning stops. Our design can refine these intervals by inducing smaller price reductions after rejection, allowing arbitrarily precise learning of every valuation below $\widehat{p}_1$.

With affiliated valuations, information about one characteristic also affects prices offered to the other, so the learning outcome under the full disclosure policy can depend on the order of arrivals. For a given valuation profile, full disclosure therefore generates a probability distribution over learning outcomes.

We illustrate this by returning to Example 1 with $K = 10$, for which $p^\infty \simeq 0.698$. We

fix the realized valuations to $(v_1^*, v_2^*) = (0.40, 0.65)$ and set $\lambda_1 = \lambda_2 = 1/2$. Under full disclosure, the platform's limiting posterior is uniform on one of the two rectangles shown in Panels (a)–(b) of Figure 2, depending on which characteristic arrives first. Each outcome occurs with probability 1/2. In either rectangle, the monopoly price for each characteristic equals the lower endpoint of its valuation interval. All subsequent offers are therefore accepted, and learning stops. Since both realized valuations lie below p^∞, our design can instead learn both arbitrarily precisely, as illustrated in Panel (c).

4. Long-lived consumers

We now consider long-lived consumers, while retaining the assumption that sellers are short-lived. With short-lived consumers, information design can influence demand only by changing sellers' beliefs about valuations. With long-lived consumers, it can also change consumers' willingness to pay today by shaping the future offers they expect after rejecting the current offer. We show that this additional channel allows the platform to learn every valuation to any desired precision.

We focus on the case of one characteristic and suppress the characteristic subscript. A single consumer is active in each period and, after accepting an offer, exits and is replaced by a new consumer with the same valuation. Consumers observe the public history and the platform's messages and discount future payoffs by $\delta_B \in (0, 1)$.[11] That is, a consumer who enters in period t obtains a payoff of $\delta_B^{t'-t}(v^* - p_{t'})$ if he purchases in period $t' \geq t$ at price $p_{t'}$, and zero if he never purchases. We maintain Assumption 2 throughout.

We follow the literature on bargaining with incomplete information and restrict attention to equilibria in which consumers' strategies are stationary and monotone; see Fudenberg et al. (1985) and Ausubel et al. (2002, Sections 3–4). These restrictions also make the models with short- and long-lived consumers comparable, because short-lived consumers' optimal behavior

[11] All our results also hold if consumers do not observe the platform's messages, with only minor modifications to the constructions for learning with long-lived consumers.

is stationary and monotone.

Fix an information design and a belief system. Write $\mu_t(\cdot \mid m_t)$ for the seller's marginal posterior over v after receiving message m_t; we refer to this as the public belief. For a consumer entering in period t_e, let $e = (h^{t_e-1}, m_{t_e})$ record his entry history and initial message; we assume that consumers observe histories. Denote by $\beta_{t_e}(h^{t-1}, m_t, p, v)$ the probability that the consumer with value v and entering in period t_e accepts price p. The strategy $\beta_{t_e}(h^{t-1}, m_t, p_t, v)$ is *stationary and monotone* if there exists a function $\overline{\beta} : \mathbb{R}_+ \times [0,1] \to [0,1]$ such that $\beta_{t_e}(h^{t-1}, m_t, p_t, v) = \overline{\beta}(p_t, v)$ for all (h^{t-1}, m_t, p_t, v) such that v is in the support of the public belief $\mu_t(\cdot \mid m_t)$, and $\overline{\beta}(p_t, v)$ is decreasing in p_t.

As in Gul, Sonnenschein, and Wilson (1986), stationarity is not imposed at all profiles $(v, (h^{t-1}, m_t, p_t))$. As we explain in Online Appendix OB.1, imposing it at all profiles is a strong requirement, often inconsistent with sequential rationality.

Definition 2 (Monotone stationary equilibrium)**.** *For a fixed information design, a monotone stationary equilibrium is a strategy profile and a belief system such that:*

(i) Each player's strategy is sequentially rational.

(ii) Beliefs satisfy Bayes' rule whenever possible.

(iii) Every consumer's strategy is stationary and monotone, and any two consumers entering when the public belief is the same have the same stationary rule.

We restrict attention to weak perfect Bayesian equilibria with stationary, monotone consumer strategies. Consumers who enter with the same public belief must use the same rule; consumers who enter with different public beliefs may use different rules. This is a weaker requirement than requiring all consumers to adopt the same strategy.[12] We impose no additional restrictions on off-path beliefs. However, our constructions do not rely on "unreasonable" off-path beliefs. In particular, the belief system supporting the construction

[12]The stronger requirement would be too demanding. To see this, suppose that the platform fully discloses the realized history. The very first consumer rejects the initial price p_1 even when their valuation v is strictly larger than p_1 since they expect lower offers in the future. Now, consider the history where the current seller believes that the consumer's valuation is in $[p_1, 1]$. Since there is full disclosure, no sellers will ever price below p_1 in all subsequent histories. A consumer with valuation at least p_1 therefore accepts an offer of p_1 at this history, even though the first consumer rejects that same price.

of Theorem 2 arises as the limit of Bayesian beliefs generated by common perturbations of consumers' and sellers' strategies. Online Appendix OB.4 establishes this result and discusses its relation to standard equilibrium refinements.

We first examine what the platform can learn under full disclosure (Section 4.1). We then show how information design can improve on this benchmark (Section 4.2).

4.1. Full disclosure

This analysis draws on the one-sided incomplete-information bargaining literature, particularly Fudenberg et al. (1985), Ausubel et al. (2002), and Sobel and Takahashi (1983). A central result is the Coase conjecture: prices converge to zero as consumers become perfectly patient. This might suggest little scope for learning and seller profits. This intuition is not correct. First, although prices before the first purchase converge to zero as $\delta_B \to 1$, the valuation of the indifferent consumer need not, so purchasing decisions remain informative. Second, a purchase creates an information externality: it reveals to the next seller that the next consumer's valuation is above the valuation of an indifferent consumer at that price, allowing her to charge a higher price. The purchasing consumer does not internalize this effect because he exits the market. These sellers can therefore earn positive profits even as consumers become perfectly patient.[13]

Proposition 1 (Full disclosure)**.** *There exists a stationary equilibrium with the following properties:*

1. ***Prices and purchasing decisions.*** *On path, if all previous offers have been rejected, seller t posts*

$$p_t^* = (1-\delta_B)\sum_{k=0}^{\infty}\delta_B^k b_{t+k}^*, \tag{6}$$

where b_t^* is the valuation of a consumer indifferent between accepting the current offer and waiting to buy next period.[14] *A consumer with valuation v accepts if and only if $v \geq b_t^*$.*

[13]The proof of Proposition 1 relies on arguments developed in the literature on bargaining models with one-sided incomplete information. For completeness, we provide the proof in Online Appendix OB.2.

[14]That is, $b_t^* - p_t^* = \delta_B(b_t^* - p_{t+1}^*)$.

2. *If p_t^* is the first accepted offer, every seller $t' > t$ posts a price b_t^*, and every consumer accepts. Furthermore, the sequence $(b_t^*)_{t\geq 0}$ starts at $b_0^* = 1$, is strictly decreasing, and converges to zero as $t \to \infty$.*

3. ***Learning.*** *At the end of period t, the platform's posterior is F conditional on $v \in [0, b_t^*)$ if no offer has been accepted. If the first acceptance occurred in period $1 \leq t' \leq t$, its posterior is F conditional on $v \in [b_{t'}^*, b_{t'-1}^*)$ and remains unchanged thereafter.*

Before the first purchase, the equilibrium exhibits successive skimming: each rejection lowers the posterior's upper bound, so the valuations of the marginal consumers and the offers both fall. A first purchase in period t instead reveals $v^* \in [b_t^*, b_{t-1}^*)$, and consumers purchase at a price of b_t^* in every subsequent period.

Learning is determined by the sequence of indifferent consumers (b_t^*). Patient consumers require a lower price to accept today because they can wait for future offers. Inducing the same indifferent consumer (and so the same demand) therefore requires a lower price than in the myopic benchmark. If this price reduction is proportional across valuations, the seller's revenue falls by the same proportion for every choice of demand, leaving her optimal choice unchanged. Otherwise, patience can change which valuations the seller finds most profitable to serve.

With the uniform distribution, the price reduction is proportional across valuations. Hence, the sequence of indifferent consumers is independent of δ_B and equals $\{1/2, 1/4, 1/8, \ldots, 1/2^t, \ldots\}$.[15] For example, if $v^* \in [1/4, 1/2)$, the first consumer rejects the first offer and accepts the second, whose price converges to zero as $\delta_B \to 1$. Every seller entering in period $t > 2$ then knows that the consumer's valuation lies in $[1/4, 1/2)$ and charges $1/4$. Thus, even as the first purchase price approaches zero, the purchase remains informative and subsequent sellers charge a positive price.

[15] This independence of the cutoffs extends to power priors $F(v) = v^\alpha$, $\alpha > 0$, and already follows from the characterization in Sobel and Takahashi (1983, Section 4). Online Appendix OB.2 also illustrates that nontrivial learning persists when consumers are very patient for nonuniform priors.

4.2. Learning beyond full disclosure

Theorem 2 establishes two points. First, as in the case with short-lived consumers, information design expands what the platform can learn relative to full disclosure. Second, it allows valuations above the static monopoly price to be distinguished, which is impossible with short-lived consumers.

Theorem 2 (Learning with patient consumers)**.** *For every $\varepsilon > 0$, there exists a threshold $\underline{\delta}_B < 1$ such that, for every $\delta_B \in [\underline{\delta}_B, 1)$, the platform can ε-learn all valuations in finite time.*

Figure 3(a) illustrates the equilibrium path. Sellers charge increasing prices $p_1 < p_2 < \cdots < p_{T(\delta_B)}$ until an offer is rejected or the final offer is accepted. The blue points represent these offers, while the red points represent the valuations b_t of consumers indifferent between accepting p_t and waiting. If p_t is the first rejected price, the platform learns that $v^* \in [b_{t-1}, b_t)$ when $t \geq 2$, and that $v^* \in [0, b_1)$ when $t = 1$. In Figure 3(a), the first rejected price is p_5, which reveals that $v^* \in [b_4, b_5)$. The information design behind Theorem 2 shapes demand so that it is highly price insensitive below p_t but drops sharply when the price exceeds p_t. Figure 3(b) illustrates the demand faced by a seller entering in period 5 after all prices $p_1, \ldots, p_4$ have been accepted. This shape of demand makes p_t optimal for the seller. The discontinuous drop in demand when the price rises above p_t is not essential; the information design can be modified to generate continuous demand.

We now explain the information design that makes these prices p_t incentive compatible. The construction uses a low continuation price b_0, induced by the message that the first offer p_1 was rejected. Following acceptance of $(p_1, \ldots, p_{t-1})$, a new consumer enters and the platform reveals that $v^* > b_{t-1}$. The platform now wants the seller to charge a price p_t that is accepted if and only if $v^* \geq b_t$, where $b_t > b_{t-1}$ is chosen so that $b_t - b_{t-1} < \varepsilon$. Rejection would then locate the valuation within an interval shorter than ε. If p_t is rejected, the platform reveals that $v^* \in [b_{t-1}, b_t)$, and the next seller charges $r(b_{t-1})$. If this offer is also rejected, the platform sends the message "the offer p_1 was rejected," conditional on which the next seller charges b_0. Writing $r(v) = (1 - \delta_B)v + \delta_B b_0$ for $v \geq b_0$, the price $r(b_{t-1})$ makes a

consumer with valuation b_{t-1} indifferent between accepting and waiting another period to buy at b_0. In turn, p_t must make a consumer with valuation b_t indifferent between accepting today and buying at $r(b_{t-1})$ next period, i.e., $b_t - p_t = \delta_B[b_t - r(b_{t-1})]$.

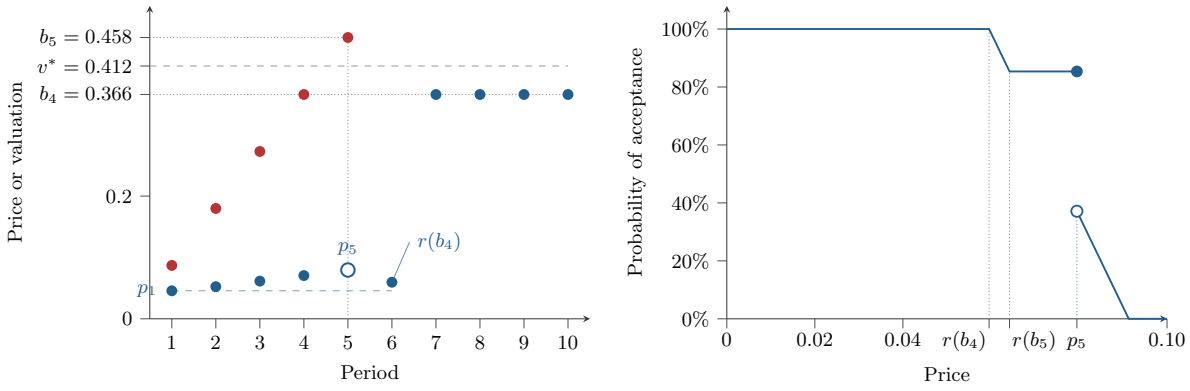
To make p_t optimal, the platform specifies what sellers learn following deviations. Suppose the seller charges $p > p_t$. If this offer is rejected, the platform conceals the history and sends the next seller the message "the offer p_1 was rejected." Conditional on this message, she charges b_0. The prospect of buying at b_0 makes waiting more attractive, causing demand to drop sharply when the price rises above p_t. This deters the deviation.

Now consider a lower price $p \in [r(b_t), p_t)$. Following rejection, the platform sends the next seller the message "the offer p_1 was rejected" with probability ϕ and reports that $v^* \in [b_{t-1}, b_t)$ with probability $1 - \phi$. Conditional on these messages, the seller charges b_0 and $r(b_{t-1})$, respectively. By choosing ϕ so that

$$b_t - p = \delta_B\Big[b_t - \phi b_0 - (1-\phi)r(b_{t-1})\Big],$$

the platform ensures that the indifferent consumer remains b_t, so the seller serves the same consumers at a lower price and earns less.

If the seller charges a price below $r(b_t)$ and the consumer rejects, the platform sends subsequent sellers the message "the offer p_1 was rejected," conditional on which they charge b_0. This means that in period t, when consumers receive an offer $p < r(b_t)$, they have the same willingness to pay as in period 1. Seller t faces the corresponding demand conditional on $v^* > b_{t-1}$. Among these deviations, the seller does best by charging $r(b_{t-1})$, which every remaining consumer accepts. Choosing b_t sufficiently close to b_{t-1} ensures that almost every remaining consumer accepts p_t. Since $p_t > r(b_{t-1})$, the seller's expected profit can then be made greater than $r(b_{t-1})$. The platform can therefore deter these deviations while ensuring $b_t - b_{t-1} < \varepsilon$.



(a) Equilibrium path.

(b) Demand at $t = 5$ when $(p_1, \ldots, p_4)$ have been accepted.

Figure 3: **Equilibrium path and demand.** Uniform prior, $v^* = 0.412$, $\delta_B = 0.95$, and an information design that generates ε-learning with $\varepsilon = 0.1$. Panel (a) shows the first ten periods. Filled blue circles mark accepted prices, the open blue circle marks the first rejected offer p_5, and red circles mark the marginal valuations b_t. Rejection identifies $v^* \in [b_4, b_5)$, where $b_4 \simeq 0.366$ and $b_5 \simeq 0.458$. The price is $r(b_4)$ in period 6 and b_4 thereafter. Panel (b) shows the probability of acceptance at each price, conditional on $v^* \geq b_4$.

Repeating this construction after each acceptance yields a finite sequence of offers that partitions the valuation range into intervals shorter than ε. The first rejection, or acceptance of the final offer, identifies the interval containing v^*. The resulting equilibrium has stationary and monotone consumer strategies. From some finite date onward, each seller earns at least $v^* - \varepsilon$. Section 5 shows how to choose the learning phase so that its discounted cost also vanishes as patience increases.

4.2.a. *Moderately patient consumers.* Theorem 2 establishes learning throughout the support when consumers are sufficiently patient. In Online Appendix OB.3, we extend the result to every $\delta_B \in (0, 1)$: for every $\varepsilon > 0$, there exists a finite date by which the platform can ε-learn all valuations in $[0, 1]$.

In the construction above, a high δ_B ensures that the initial interval $[0, b_1)$ is shorter than ε. For less patient consumers, we add an initial stage of decreasing offers. Prices and acceptance cutoffs are chosen jointly so that each offer is optimal for the seller, and the final cutoff is below ε. If the initial consumer rejects every offer, the valuation is already learned.

Otherwise, his first purchase identifies an interval between two consecutive cutoffs, and a new consumer enters. Within that interval, we can use the construction used in Theorem 2 to obtain ε-learning.

5. Consumer and producer surplus

In this section, we characterize the combinations of consumer and producer surplus that the platform can attain in the one-characteristic environment. The platform discounts future surplus with factor $\delta \in (0,1)$, and we focus on the attainable outcomes as $\delta \uparrow 1$. We analyze environments with short-lived and long-lived consumers, setting $\delta_B = \delta$ in the latter.

Fix an information design $\mathcal{M}$ and an induced equilibrium. Let a_t indicate whether trade occurs in period t, and let p_t denote the price. Discounted producer and consumer surplus are

$$\mathrm{PS}_\delta(\mathcal{M}) = (1-\delta)\sum_{t\geq 1}\delta^{t-1}\mathbb{E}[a_t p_t],$$
$$\mathrm{CS}_\delta(\mathcal{M}) = (1-\delta)\sum_{t\geq 1}\delta^{t-1}\mathbb{E}[a_t (v-p_t)].$$

Expectations are taken over v, the information design, and any equilibrium randomization. The factor $1-\delta$ expresses welfare as discounted average surplus per period. For a fixed induced payoff sequence, discounted average surplus $X_\delta(\mathcal{M})$, $X \in \{\mathrm{PS}, \mathrm{CS}\}$, converges to the sequence's long-run (Cesàro) average as $\delta \uparrow 1$, whenever that average exists. Because marginal cost is zero, the maximum ex ante expected total surplus in any period is $\mathbb{E}[v]$, attained when trade always occurs.

Before play, the platform may publicly randomize over finitely many information designs, independently of v. Fix an induced equilibrium for each design, restricted to be stationary when consumers are long lived.

Definition 3 (Attainable outcomes)**.** *An outcome (P, C) consists of producer surplus P and consumer surplus C. An outcome (P, C) is* asymptotically attainable *if there is a family of*

finite lotteries over information designs $\{\pi_\delta\}$ *satisfying*

$$\lim_{\delta \uparrow 1} \mathbb{E}|_{\mathcal{M} \sim \pi_\delta} \left[\left(\mathrm{PS}_\delta(\mathcal{M}), \mathrm{CS}_\delta(\mathcal{M}) \right) \right] = (P, C).$$

A useful benchmark is the set of attainable outcomes with static information design when the platform observes the realized valuation v^* from the outset; surplus is still evaluated ex ante over v. Let $\hat{p}$ be the static monopoly price, and let $\mathrm{PS}(\hat{p})$ and $\mathrm{CS}(\hat{p})$ be the corresponding static producer and consumer surplus. In this informed-platform benchmark, Bergemann et al. (2015) show that every nonnegative surplus pair with total surplus no greater than $\mathbb{E}[v]$ and producer surplus at least $\mathrm{PS}(\hat{p})$ is attainable. This set is region $A + B$ in Figure 4.

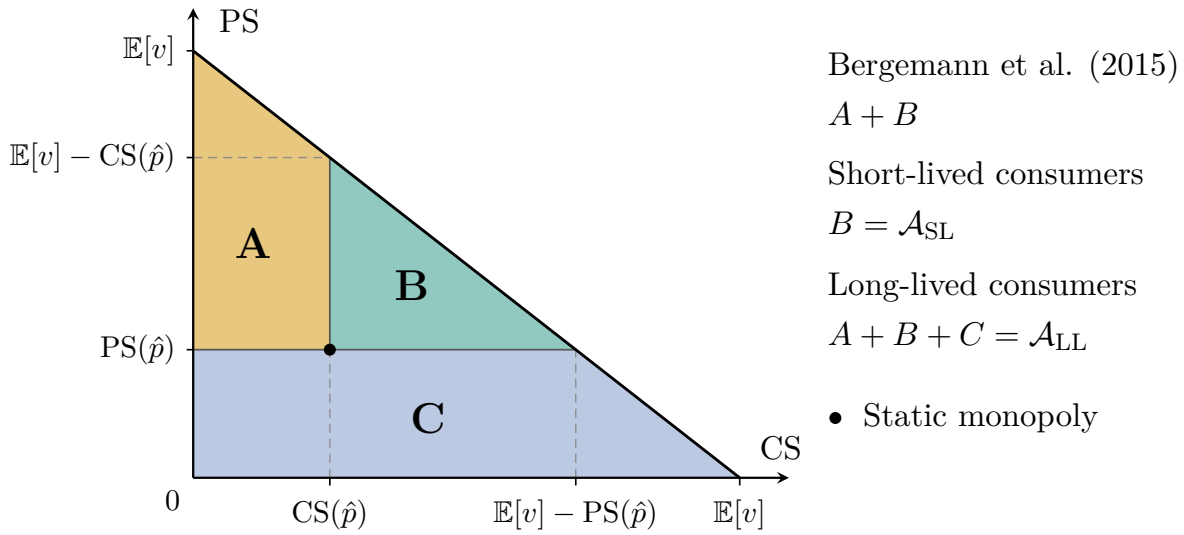


Figure 4: **Attainable consumer and producer surplus.** Schematic comparison with the informed-platform benchmark of Bergemann et al. (2015). The black dot is the static monopoly outcome, and the diagonal is $\mathrm{PS} + \mathrm{CS} = \mathbb{E}[v]$. The dynamic sets are limiting outcomes as $\delta \uparrow 1$, with $\delta_B = \delta$ for long-lived consumers.

The next result characterizes attainable outcomes when the platform uses information design to learn v^*. Let $\mathcal{A}_{\mathrm{SL}}$ and $\mathcal{A}_{\mathrm{LL}}$ denote the sets of asymptotically attainable outcomes with short-lived and long-lived consumers, respectively.

Theorem 3 (Attainable surplus outcomes)**.**

(i) ***Short-lived consumers.***

$$\mathcal{A}_{\text{SL}} = \{(P, C) : P \geq \text{PS}(\hat{p}), \quad C \geq \text{CS}(\hat{p}), \quad P + C \leq \mathbb{E}[v]\}\,.$$

(ii) ***Long-lived consumers.***

$$\mathcal{A}_{\text{LL}} = \{(P, C) : P \geq 0, \quad C \geq 0, \quad P + C \leq \mathbb{E}[v]\}\,.$$

With short-lived consumers, every information design leaves each side at least its static monopoly surplus. Sellers can ignore the platform's message and charge $\hat{p}$, guaranteeing expected profit $\text{PS}(\hat{p})$. Consumers retain at least $\text{CS}(\hat{p})$ because no design can induce a price above $\hat{p}$. Together with the total-surplus bound, these restrictions define the triangle in part (i). This is region B in Figure 4.

Three constructions attain its vertices. Concealing the history gives the static outcome $(\text{PS}(\hat{p}), \text{CS}(\hat{p}))$.

To maximize producer surplus, the platform first learns valuations below $\hat{p}$ arbitrarily precisely and then reveals the learned intervals. Sellers charge close to the valuation below $\hat{p}$, while higher valuations remain pooled and pay $\hat{p}$. Since learning below $\hat{p}$ takes finite time for any given precision, this yields $(\mathbb{E}[v] - \text{CS}(\hat{p}), \text{CS}(\hat{p}))$ in the limit.

To maximize consumer surplus, the platform first learns valuations below $\hat{p}$ and then uses the learned distribution to induce efficient trade while leaving sellers their static monopoly profit. The design extends the uniform-profit segmentation of Bergemann et al. (2015) and yields $(\text{PS}(\hat{p}), \mathbb{E}[v] - \text{PS}(\hat{p}))$. Public randomization over these three designs attains every point in the triangle. Region A in Figure 4 thus identifies the surplus outcomes lost when the platform must learn valuations from short-lived consumers' purchases.

With long-lived consumers, neither side's static monopoly surplus remains a lower bound. Only nonnegativity and the total-surplus bound remain, defining the triangle in part (ii).

This is region $A + B + C$ in Figure 4.

To see why, consider first the no-disclosure policy of Proposition 4 in Appendix C.2. There is a stationary equilibrium in which sellers charge a constant price over time. Consumers can reject a higher offer and wait for the equilibrium price from the next seller. As consumers become patient, this option makes price increases less profitable, and the lowest possible equilibrium price converges to zero, yielding $(0, \mathbb{E}[v])$.

Second, the platform can use the information design constructed in Theorem 2 to ε-learn v^*. For any fixed ε and sufficiently high δ, there is a finite date $T_\varepsilon(\delta)$ after which every seller earns at least $v^* - \varepsilon$. Although this date may increase as $\delta \uparrow 1$, the proof of Theorem 3(ii) in Appendix C.2 shows that it grows at most logarithmically in $1/(1-\delta)$ for each fixed precision. The learning phase is therefore short enough that its discounted cost vanishes. Letting ε decrease sufficiently slowly as $\delta \uparrow 1$ yields $(\mathbb{E}[v], 0)$.

Finally, the platform can make both surpluses vanish by sustaining a constant price that approaches one. It conceals histories along the equilibrium path. Following rejection of a lower offer, it sometimes switches to full disclosure, with a probability that depends on the rejected offer. The prospect of low future prices makes consumers excluded by the prescribed price unwilling to accept moderate discounts, while larger discounts leave sellers too little profit. Sellers therefore prefer the prescribed price. As consumers become patient, almost all valuations are excluded from trade, and both surpluses converge to zero. The formal construction is given in Proposition 5 in Appendix C.2.

Public randomization over these three designs can yield, in expectation, any outcome in the triangle characterized in part (ii). With long-lived consumers, information design thus recovers region A and expands the attainable set beyond the static informed-platform benchmark to include region C. By shaping the option to wait, the platform can attain any division of efficient surplus, as well as any lower total surplus and its division between the two sides.[16]

[16] Makris and Renou (2023, Section 5.3, Proposition 1) obtain a related surplus characterization with a fully informed mediator using multistage information design. In their bilateral-trade model, the buyer is

6. Concluding remarks

We show how a platform's disclosure choices determine what it can learn from consumers' purchases. With short-lived consumers, learning is limited by sellers' pricing incentives. With long-lived consumers, disclosure also shapes the value of waiting and permits learning throughout the valuation range. These differences determine the surplus divisions the platform can attain. We now position our paper in the literature.

6.1. Dynamic Persuasion

Dynamic persuasion studies how a designer discloses information about an unknown state ω to influence agents' actions over time. The designer starts from a prior, observes actions a, and receives signals whose distribution $\rho(\cdot \mid a, \omega)$ may depend on those actions. An information design specifies, after each history, recommendations that Bayesian agents are willing to follow.

Most of the dynamic persuasion literature focuses on two scenarios. In the first one, the designer receives a fully informative signal about the state at the outset. In the second scenario, signals arrive over time but do not depend on the agent's actions, i.e., $\rho(\cdot \mid a, \omega) = \rho(\cdot \mid \omega)$. Examples include Ely (2017), Arieli et al. (2024), and Renault et al. (2017), who study myopic agents, as well as Ball (2023), Ely and Szydlowski (2020), Koh and Sanguanmoo (2022), Orlov et al. (2020), Smolin (2021), and Zhao et al. (2024), who study long-lived agents.[17]

In our model, the designer is a platform that learns about demand by influencing sellers' prices. Its disclosures induce prices that make consumers' purchase decisions more informative, allowing the platform to learn more and shape later prices more effectively. Several papers also study the interaction between disclosure and information generation. Che and Hörner

initially uninformed about his valuation; the mediator can inform the seller before the offer and the buyer after observing it. Their surplus triangle approaches the one above as the lowest possible valuation tends to zero.

[17] Relatedly, Renou and Venel (2025) compare information structures in discounted decision problems, including problems in which actions affect future information.

(2018), Immorlica et al. (2020) and Kremer et al. (2014) study how recommendations induce a sequence of myopic agents to experiment. Koh and Li (2025) study how disclosure promotes social learning across actions when agents have heterogeneous preferences, while Knoepfle and Salmi (2024) study how the timing of evidence disclosure affects forward-looking agents' adoption decisions. Our paper focuses on the limits to learning consumers' valuations when a platform influences prices through information disclosure.[18] Despite the different applications and methods of analysis, these papers all examine how disclosure affects the generation of new information. We expect future research to further explore these feedback loops between persuasion and data collection.

6.2. Price Discrimination and Endogenous Information

Closer to applications, our work is inspired by research on how the information available before trade can be used to shape market outcomes (see, among others, Bergemann et al., 2015, Roesler and Szentes, 2017, Armstrong and Zhou, 2022, Elliott et al., 2026, Ichihashi and Smolin, 2025, Yang, 2022).[19] We endogenize the intermediary's information acquisition and characterize how its interaction with disclosure shapes attainable consumer and producer surplus.

Other papers examine consumers' strategic incentives when their behavior or disclosures reveal information to sellers or platforms, as well as the design of institutions that govern its collection and aggregation (see, among others, Conitzer et al., 2012, Bonatti and Cisternas, 2020, Ichihashi, 2023, Bergemann et al., 2022, Doval and Skreta, 2026). We complement this literature by characterizing how much a platform can learn from purchase decisions when it controls the information disclosed to sellers, and how consumers' option to wait changes both

[18] See Fudenberg and Villas-Boas (2006) for a survey of behavior-based price discrimination and customer recognition. For direct price experimentation by a monopolist facing demand uncertainty, see Rothschild (1974).

[19] A separate branch of the literature on consumer privacy examines additional welfare implications of platforms' capacity to record users' online activity—implications that go beyond price discrimination. Examples include Acemoglu et al. (2025) and Fainmesser et al. (2023); see Goldfarb and Que (2023) for a survey.

learning and attainable surplus.

In this literature, Ali et al. (2023) study whether consumers benefit from disclosing verifiable information about their valuations when sellers can price discriminate. They show the existence of an equilibrium that replicates our full-disclosure outcome with one characteristic and short-lived consumers (see the discussion in Section 3.3). Both their construction and ours use intervals of valuations within which the seller prefers to serve every consumer.[20] In their model, this prompts consumers in each interval to disclose truthfully, while disclosing lower valuations is discouraged by appropriate off-equilibrium beliefs.

References

D. Acemoglu, A. Makhdoumi, A. Malekian, and A. Ozdaglar. When big data enables behavioral manipulation. *American Economic Review: Insights*, 7(1):19–38, 2025. doi: 10.1257/aeri.20230589.

S. N. Ali, G. Lewis, and S. Vasserman. Voluntary disclosure and personalized pricing. *The Review of Economic Studies*, 90(2):538–571, 2023. doi: 10.1093/restud/rdac033.

I. Arieli, Y. Babichenko, D. Shaiderman, and X. Shi. Persuading while learning. arXiv preprint arXiv:2407.13964, 2024.

M. Armstrong and J. Zhou. Consumer information and the limits to competition. *American Economic Review*, 112(2):534–577, 2022. doi: 10.1257/aer.20210083.

L. M. Ausubel, P. Cramton, and R. J. Deneckere. Bargaining with incomplete information. In R. J. Aumann and S. Hart, editors, *Handbook of Game Theory with Economic Applications*, volume 3, chapter 50, pages 1897–1945. Elsevier Science B.V., Amsterdam, 2002.

[20] Specifically, these are the intervals $[\widehat{p}_t, \widehat{p}_{t-1}]$, with $\widehat{p}_t$ defined in (5).

I. Ball. Dynamic information provision: Rewarding the past and guiding the future. *Econometrica*, 91(4):1363–1391, 2023.

S. Bar-Gill, E. Brynjolfsson, and N. Hak. Helping small businesses become more data-driven: A field experiment on eBay. *Management Science*, 70(11):7345–7372, 2024. doi: 10.1287/mnsc.2021.02026.

D. Bergemann, B. Brooks, and S. Morris. The limits of price discrimination. *American Economic Review*, 105(3):921–957, 2015. doi: 10.1257/aer.20130848.

D. Bergemann, A. Bonatti, and T. Gan. The economics of social data. *The RAND Journal of Economics*, 53(2):263–296, 2022.

A. Bonatti and G. Cisternas. Consumer scores and price discrimination. *The Review of Economic Studies*, 87(2):750–791, 2020.

Y.-K. Che and J. Hörner. Recommender systems as mechanisms for social learning. *The Quarterly Journal of Economics*, 133(2):871–925, 2018.

Y. Chen, X. Cui, A. Li, B. Wu, and L. Yang. The role of digital platforms in data markets: How platform-shared market data through advanced analytics empower firm innovation. *Management Science*, 2026. doi: 10.1287/mnsc.2024.06633. Articles in Advance.

V. Conitzer, C. R. Taylor, and L. Wagman. Hide and seek: Costly consumer privacy in a market with repeat purchases. *Marketing Science*, 31(2):277–292, 2012.

L. Doval and V. Skreta. Purchase history and product personalization. *The RAND Journal of Economics*, 57(1):78–102, 2026. doi: 10.1111/1756-2171.70007.

M. Elliott, A. Galeotti, A. Koh, and W. Li. Market segmentation through information. *The Review of Economic Studies*, art. rdag067, 2026. doi: 10.1093/restud/rdag067.

J. C. Ely. Beeps. *American Economic Review*, 107(1):31–53, 2017.

J. C. Ely and M. Szydlowski. Moving the goalposts. *Journal of Political Economy*, 128(2): 468–506, 2020.

I. P. Fainmesser, A. Galeotti, and R. Momot. Digital privacy. *Management Science*, 69(6): 3157–3173, 2023.

D. Fudenberg and J. M. Villas-Boas. Behavior-based price discrimination and customer recognition. *Handbook on economics and information systems*, 1:377–436, 2006.

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

A. Goldfarb and V. F. Que. The economics of digital privacy. *Annual Review of Economics*, 15:267–286, 2023.

F. Gul, H. Sonnenschein, and R. Wilson. Foundations of dynamic monopoly and the Coase conjecture. *Journal of Economic Theory*, 39(1):155–190, 1986. doi: 10.1016/0022-0531(86)90024-4.

S. Ichihashi. Dynamic privacy choices. *American Economic Journal: Microeconomics*, 15(2): 1–40, 2023.

S. Ichihashi and A. Smolin. Data provision to an informed seller. *Games and Economic Behavior*, 153:131–144, 2025. doi: 10.1016/j.geb.2025.06.002.

N. Immorlica, J. Mao, A. Slivkins, and Z. S. Wu. Incentivizing exploration with selective data disclosure. In *Proceedings of the 21st ACM Conference on Economics and Computation*, pages 647–648. Association for Computing Machinery, 2020. doi: 10.1145/3391403.3399487. URL https://arxiv.org/abs/1811.06026.

S. Karlin and Y. Rinott. Classes of orderings of measures and related correlation inequalities. i.

multivariate totally positive distributions. *Journal of Multivariate Analysis*, 10(4):467–498, 1980. doi: 10.1016/0047-259X(80)90065-2.

J. Knoepfle and J. Salmi. Dynamic evidence disclosure: Delay the good to accelerate the bad. arXiv preprint arXiv:2406.11728, 2024.

A. Koh and R. Li. Balanced social learning. Mimeo, 2025.

A. Koh and S. Sanguanmoo. Attention capture. arXiv preprint arXiv:2209.05570, 2022.

I. Kremer, Y. Mansour, and M. Perry. Implementing the "wisdom of the crowd". *Journal of Political Economy*, 122(5):988–1012, 2014.

M. Makris and L. Renou. Information design in multistage games. *Theoretical Economics*, 18 (4):1475–1509, 2023. doi: 10.3982/TE4769. URL https://onlinelibrary.wiley.com/doi/abs/10.3982/TE4769.

R. B. Myerson and P. J. Reny. Perfect conditional ε-equilibria of multi-stage games with infinite sets of signals and actions. *Econometrica*, 88(2):495–531, 2020. doi: 10.3982/ECTA13426. URL https://doi.org/10.3982/ECTA13426.

D. Orlov, A. Skrzypacz, and P. Zryumov. Persuading the principal to wait. *Journal of Political Economy*, 128(7):2542–2578, 2020.

J. Renault, E. Solan, and N. Vieille. Optimal dynamic information provision. *Games and Economic Behavior*, 104:329–349, 2017.

L. Renou and X. Venel. Comparing experiments in discounted problems, 2025. URL https://arxiv.org/abs/2405.16458.

A.-K. Roesler and B. Szentes. Buyer-optimal learning and monopoly pricing. *American Economic Review*, 107(7):2072–2080, 2017.

M. Rothschild. A two-armed bandit theory of market pricing. *Journal of Economic Theory*, 9(2):185–202, 1974. doi: 10.1016/0022-0531(74)90066-0.

A. Smolin. Dynamic evaluation design. *American Economic Journal: Microeconomics*, 13(4): 300–331, 2021.

J. Sobel and I. Takahashi. A multistage model of bargaining. *The Review of Economic Studies*, 50(3):411–426, 1983. doi: 10.2307/2297673.

K. H. Yang. Selling consumer data for profit: Optimal market-segmentation design and its consequences. *American Economic Review*, 112(4):1364–1393, 2022.

W. Zhao, C. Mezzetti, L. Renou, and T. Tomala. Contracting over persistent information. *Theoretical Economics*, 19(2):917–974, 2024.

Appendix

A. Short-lived consumers

A.1. Proof of Theorem 1

Proof of part (i) of Theorem 1

We start with a preliminary observation, stated in Lemma 1, that we use to prove part (i). Under affiliation, raising the lower and upper endpoints of the conditioning valuation intervals weakly increases the optimal price for every characteristic. Let (I'_c, I_c) be two intervals in $[0,1]$. We write $I'_c \succeq I_c$ if $\sup I'_c \geq \sup I_c$ and $\inf I'_c \geq \inf I_c$.

Lemma 1. *If $I'_{\tilde{c}} \succeq I_{\tilde{c}}$ for all $\tilde{c} \in \mathcal{C}$, then for all $c \in \mathcal{C}$,*

$$p'_c := \arg\max_{\tilde{p}_c} \tilde{p}_c \mathbb{P}_f(v_c \geq \tilde{p}_c \mid v_{\tilde{c}} \in I'_{\tilde{c}}, \forall \tilde{c}) \geq p_c := \arg\max_{\tilde{p}_c} \tilde{p}_c \mathbb{P}_f(v_c \geq \tilde{p}_c \mid v_{\tilde{c}} \in I_{\tilde{c}}, \forall \tilde{c}).$$

Recall that the arg max are unique as a consequence of the log-concavity of f.

Proof. Set $a := \inf I_c$, $b := \sup I_c$, $a' := \inf I'_c$, and $b' := \sup I'_c$. Since $I'_c \succeq I_c$, we have $a' \geq a$ and $b' \geq b$. The optimal prices lie in $[a, b)$ and $[a', b')$, respectively. Thus, if $a' \geq b$, the result is immediate. Suppose $a' < b$.

Step 1: Affiliation orders the conditional densities. Let $I_{-c} := \prod_{\tilde{c} \neq c} I_{\tilde{c}}$ and define I'_{-c} analogously. Define the unnormalized conditional densities

$$f(v_c, I_{-c}) := \int_{I_{-c}} f(v_c, v_{-c})\, dv_{-c}, \qquad f(v_c, I'_{-c}) := \int_{I'_{-c}} f(v_c, v_{-c})\, dv_{-c}.$$

Note that the map $(v_{\tilde{c}}, I_{\tilde{c}}) \mapsto \mathbb{1}_{I_{\tilde{c}}}(v_{\tilde{c}})$ is log-supermodular with respect to the orders $(\geq, \succeq)$. Since f is affiliated, closure under products and integration implies that the map $(v_c, I_{-c}) \mapsto f(v_c, I_{-c})$ is also log-supermodular (Karlin and Rinott, 1980). Therefore, $I'_{-c} \succeq I_{-c}$ implies that $r(v_c) := f(v_c, I'_{-c})/f(v_c, I_{-c})$ is increasing in v_c.

Step 2: The ratio of demands is increasing in price. For $p \in [a', b)$, the ratio

of conditional demands equals a positive constant independent of p times $R(p) := \int_p^{b'} f(v_c, I'_{-c})\,dv_c \big/ \int_p^b f(v_c, I_{-c})\,dv_c$. The sign of $R'(p)$ is the sign of

$$\int_p^b [r(v_c) - r(p)] f(v_c, I_{-c})\,dv_c + \int_b^{b'} r(v_c) f(v_c, I_{-c})\,dv_c,$$

which is nonnegative because r is increasing. Hence R is increasing.

Step 3: The optimal price is higher under the intervals $(I'_{\tilde{c}})_{\tilde{c}}$. Let $\Pi(p)$ and $\Pi'(p)$ denote the seller's profits conditional on the original and primed intervals, respectively. For $p \in [a', b)$, we have $\Pi'(p) = KR(p)\Pi(p)$, where $K > 0$ is independent of p.

Suppose, toward a contradiction, that $p'_c < p_c$. Both prices lie in $[a', b)$, and uniqueness of the optimum under the original intervals gives $\Pi(p_c) > \Pi(p'_c) > 0$. Since R is increasing,

$$\frac{\Pi'(p_c)}{\Pi'(p'_c)} = \frac{R(p_c)}{R(p'_c)} \frac{\Pi(p_c)}{\Pi(p'_c)} > 1,$$

contradicting the optimality of p'_c. Therefore, $p'_c \geq p_c$.

□

We now prove part (i). For every history h^{t-1}, let

$$I_c(h^{t-1}) := \Big[\sup\{p_s : a_s = 1,\ c_s = c\}, \inf\{p_s : a_s = 0,\ c_s = c\}\Big) \cap [0, 1]$$

be the interval of valuations v_c consistent with h^{t-1}, with the conventions that the supremum of the empty set is $-\infty$ and the infimum of the empty set is $+\infty$.

We first establish a property of the upper orthant $[p^\infty, 1]^C$. For every characteristic c, define

$$\Pi_c^\infty(p) := p\mathbb{P}_f(v_c \geq p \,|\, v_{c'} \geq p^\infty \text{ for every } c' \in \mathcal{C})\,.$$

By (2) and (4), continuity of conditional demand and the maximum theorem imply that $p^\infty = \arg\max_{p \geq p^\infty} \Pi_c^\infty(p)$ for every $c \in \mathcal{C}$. Conditional on the upper orthant, demand equals one at every $p \leq p^\infty$. Hence, p^∞ also maximizes Π_c^∞ over all $p \in [0, 1]$.

We now prove by induction that no seller posts a price above p^∞. Suppose that all prices posted before period t are weakly below p^∞. This holds when $t = 1$ because no prices have yet been posted. It follows that, for every history h^{t-1} and every characteristic c',

$[p^\infty, 1] \succeq I_{c'}(h^{t-1})$.

For almost every message m_t under the induced message distribution, Bayesian conditioning gives the following representation of the seller's expected profit from posting p

$$\Pi_t(p \mid m_t) = \sum_{c \in \mathcal{C}} \int_{H^{t-1}} p\mathbb{P}_{\mathcal{M}}\Big(v_c \geq p \,\Big|\, h^{t-1}, c_t = c, m_t\Big)\; d\mathbb{P}_{\mathcal{M}}\Big(h^{t-1}, c_t = c \,\Big|\, m_t\Big). \tag{7}$$

For every pair (c, h^{t-1}) in the support of the conditional distribution in (7), Bayesian consistency gives

$$\mathbb{P}_{\mathcal{M}}\Big(v_c \geq p \,\Big|\, h^{t-1}, c_t = c, m_t\Big) = \mathbb{P}_f\Big(v_c \geq p \,\Big|\, v_{c'} \in I_{c'}(h^{t-1}) \text{ for every } c'\Big).$$

Lemma 1 and $[p^\infty, 1] \succeq I_{c'}(h^{t-1})$ for every c' therefore imply that

$$\arg\max_p p\mathbb{P}_{\mathcal{M}}\Big(v_c \geq p \,\Big|\, h^{t-1}, c_t = c, m_t\Big) \leq \arg\max_p \Pi_c^\infty(p) = p^\infty.$$

Log-concavity implies that each component profit $p \mapsto p\mathbb{P}_{\mathcal{M}}(v_c \geq p \,|\, h^{t-1}, c_t = c, m_t)$ in (7) is single-peaked and has a unique maximizer. Since that maximizer is weakly below p^∞, each component is nonincreasing above p^∞ and strictly decreasing wherever its demand is positive. Consequently, for every $p > p^\infty$ such that $\Pi_t(p \mid m_t) > 0$, it holds that $\Pi_t(p \mid m_t) < \Pi_t(p^\infty \mid m_t)$.

If instead $\Pi_t(p \mid m_t) = 0$, then p cannot be optimal because the seller can obtain strictly positive profit at some lower price. Thus, no price strictly above p^∞ can be optimal. The argument applies to every message m_t with positive probability and completes the induction.

It follows that consumers are never offered a price above p^∞. Fix $c \in \mathcal{C}$ and two valuation profiles $\mathbf{v}^*, \hat{\mathbf{v}}^* \in V^C$ such that $v^*_{c'} = \hat{v}^*_{c'}$ for every $c' \neq c$, $v^*_c, \hat{v}^*_c \in [p^\infty, 1]$, and $v^*_c \neq \hat{v}^*_c$. Under both valuation profiles, consumers with characteristic c accept every price, while consumers with every other characteristic make the same purchase decision because their valuations are unchanged. The two profiles therefore induce the same distribution over histories under every information design.

Proof of Part (ii) of Theorem 1

We provide a roadmap of the proof.

1. We introduce some notation and establish two preliminary lemmas, Lemmas 2 and 3, that are used in the subsequent constructions.
2. We consider a single characteristic c and construct an information design $\mathcal{M}_c^*$ that guarantees ε-learning of all valuations in $[0, p_c^1)$ in finite time, where p_c^1 is the initial monopoly price defined in (2) of the main text. This result is summarized in Proposition 2.
3. We consider multiple characteristics and construct an information design $\mathcal{M}_{I,c^*}^{**}$ that extends the learning result from one characteristic to multiple characteristics. This result is summarized in Proposition 3.
4. We construct a grand mechanism that uses these two information designs over finitely many blocks, following the increasing threshold sequence (p^k) defined in (2)–(3) of the main text, and repeats this sequence of blocks.
5. We show that the grand mechanism achieves the learning claimed in part (ii) of Theorem 1.

1. Notation and Preliminary Results Throughout, we write $\mathbb{P}_f$ for the probability over the Borel sets of $[0,1]^C$ induced by f, and $\mathbb{P}$ for the probability over histories induced by the mechanism and the strategies of the short-lived sellers and consumers. We write I for a product of non-empty intervals $\prod_{c' \in \mathcal{C}} I_{c'}$ and let $F_I := F(\cdot \mid \mathbf{v} \in I)$. Fix $c \in \mathcal{C}$ and let $F_{I,c}$ and $f_{I,c}$ denote, respectively, the marginal distribution and density of v_c under F_I. When $I = [0,1]^C$, we have $F_I = F$.

To start with, observe that the unnormalized density

$$v_c \longmapsto \int_{\prod_{c' \neq c} I_{c'}} f(v_c, v_{-c})\, dv_{-c}$$

is log-concave by the Prékopa–Leindler inequality. It follows that its restriction to I_c and $f_{I,c}$ are also log-concave.

We let $D_{I,c}(p)$ be the demand at p: $D_{I,c}(p) = 0$ for $p \geq \sup I_c$, $D_{I,c}(p) = 1$ for $p \leq \inf I_c$, and, for $p \in I_c$,

$$D_{I,c}(p) = \frac{F_{I,c}(\sup I_c) - F_{I,c}(p)}{F_{I,c}(\sup I_c) - F_{I,c}(\inf I_c)}.$$

Notice that $\arg\max_{p\in[0,1]} pD_{I,c}(p)$ is single-valued (because of log-concavity of $f_{I,c}$).

Set $p^0_{I,c} = 1$ and define inductively for all $k \geq 1$:

$$p^k_{I,c} := \arg\max_{p\in[0,1]} pD_{I\cap[0,p^{k-1}_{I,c}],c}(p), \tag{8}$$

where we slightly abuse notation and write $I \cap [0, p^{k-1}_{I,c}]$ for the product of intervals $\prod_{c'\neq c} I_{c'} \times (I_c \cap [0, p^{k-1}_{I,c}])$. Once $p^k_{I,c} = \inf I_c$, keep the sequence constant. Again, log-concavity implies that the maximizers are unique. Clearly, $(p^k_{I,c})_k$ is a decreasing sequence of elements in $[\inf I_c, \sup I_c]$, hence admits a limit. We establish two lemmas concerning the conditional marginal distribution $F_{I,c}$. Lemma 2 is used in the proof of Proposition 2. Lemma 3 is used both in that proof and in the construction of the grand mechanism.

Lemma 2. *The limit* $\lim_{k\to\infty} p^k_{I,c} = \inf I_c$.

Proof of Lemma 2. By contradiction, suppose that $\lim_{k\to\infty} p^k_{I,c} = \ell > \inf I_c$. Choose $p \in (\inf I_c, \ell)$. From the optimality of $p^k_{I,c}$, we have

$$p^k_{I,c} D_{I\cap[0,p^{k-1}_{I,c}],c}(p^k_{I,c}) \geq pD_{I\cap[0,p^{k-1}_{I,c}],c}(p),$$

for all k. Taking limits on both sides, we have that $0 \geq pD_{I\cap[0,\ell],c}(p)$. Since p and f are strictly positive, $pD_{I\cap[0,\ell],c}(p) > 0$, hence a contradiction. □

Lemma 3. *For each k such that $p^k_{I,c} > \inf I_c$ and $\widehat{p} \in [p^{k+1}_{I,c}, p^k_{I,c}]$, there exists $\mu \in \left[\frac{\mathbb{P}_{f_{I,c}}(I_c\cap[0,p^k_{I,c}])}{\mathbb{P}_{f_{I,c}}(I_c\cap[0,p^k_{I,c}])+\mathbb{P}_{f_{I,c}}(I_c\cap[p^k_{I,c},p^{k-1}_{I,c}])}, 1\right]$ such that*

$$\{\widehat{p}\} = \arg\max_{p\in[0,1]} p\left[\mu D_{I\cap[0,p^k_{I,c}],c}(p) + (1-\mu) D_{I\cap[p^k_{I,c},p^{k-1}_{I,c}],c}(p)\right].$$

Proof of Lemma 3. Observe that it is never optimal to price above $p^k_{I,c}$ since $p^k_{I,c}$ is the optimal price when facing the demand $D_{I\cap[0,p^{k-1}_{I,c}],c}$ and, therefore, when facing the de-

mand $D_{I\cap[p^k_{I,c},p^{k-1}_{I,c}],c}$. Thus, we can restrict attention to prices $p \in [0, p^k_{I,c}]$. At these prices, $D_{I\cap[p^k_{I,c},p^{k-1}_{I,c}],c}(p) = 1$. For $\inf I_c < p < p^k_{I,c}$, the derivative is:

$$1 - \frac{\mu}{F_{I,c}(p^k_{I,c})}[F_{I,c}(p) + pf_{I,c}(p)]$$
$$= 1 - \frac{\mu}{F_{I,c}(p^k_{I,c})} + \frac{\mu}{F_{I,c}(p^k_{I,c})}\left[(1 - F_{I,c}(p))\left(1 - \frac{pf_{I,c}(p)}{1 - F_{I,c}(p)}\right)\right].$$

Its left-hand limit at $p^k_{I,c}$ is equal to $1 - \mu F_{I,c}(p^{k-1}_{I,c})/F_{I,c}(p^k_{I,c})$ since $F_{I,c}(p^{k-1}_{I,c}) - F_{I,c}(p^k_{I,c}) - p^k_{I,c}f_{I,c}(p^k_{I,c}) = 0$ – the first-order condition for the optimality of $p^k_{I,c}$. Thus, at $\mu = \frac{F_{I,c}(p^k_{I,c})}{F_{I,c}(p^{k-1}_{I,c})}$, its value is zero and it is strictly negative for larger values of μ.

For $\inf I_c < p \leq p^k_{I,c} \leq p^1_{I,c}$,

$$1 - \frac{pf_{I,c}(p)}{1 - F_{I,c}(p)} \geq 0.$$

By log-concavity of $f_{I,c}$, the derivative is strictly decreasing on $(\inf I_c, p^k_{I,c})$. Below $\inf I_c$, profit equals p. Thus, there is a unique maximizer for each μ, possibly at the lower endpoint $\inf I_c$. From Berge's maximum theorem, the maximizer is continuous in μ, taking value $p^{k+1}_{I,c}$ (resp., $p^k_{I,c}$) at $\mu = 1$ (resp., $\mu = \frac{F_{I,c}(p^k_{I,c})}{F_{I,c}(p^{k-1}_{I,c})}$). □

2. Single Characteristic and Design $\mathcal{M}^*_c$. We now consider an environment with a single characteristic c and valuation distribution $F_{I,c}$. We construct an information design $\mathcal{M}^*_c$ that learns valuations below the initial monopoly price $p^1_{I,c}$ with arbitrary precision in finite time. The design $\mathcal{M}^*_c$ induces sellers to offer a finite grid of prices $\widehat{p}_1 < \cdots < \widehat{p}_N < p^1_{I,c}$. Recall that $(p^k_{I,c})_k$, defined in (8), is the decreasing sequence of monopoly prices following successive rejections. To induce prices in $[p^{k+1}_{I,c}, p^k_{I,c}]$, the platform mixes the posteriors following acceptance and rejection of $p^k_{I,c}$, conditional on rejection of $p^{k-1}_{I,c}$. The higher the probability assigned to acceptance, the closer the induced price is to $p^k_{I,c}$.

Proposition 2. *Assume there is a single characteristic c, $\lambda_c = 1$, whose valuation has distribution $F_{I,c}$. There exists a mechanism $\mathcal{M}^*_c$ such that the platform ε-learns all valuations in $[\inf I_c, p^1_{I,c})$ in at most $t^*_c < \infty$ periods.*

Proof. The proof is constructive. Fix $\varepsilon > 0$. Lemma 2 guarantees that there exists a positive integer K such that

$$p_{I,c}^K - \inf I_c < \varepsilon.$$

If $K = 1$, rejection of the initial offer $p_{I,c}^1$ locates the valuation in $[\inf I_c, p_{I,c}^1)$, whose length is smaller than ε. Thus, the initial offer suffices to establish the result. We therefore consider $K \geq 2$. For each $k = 1, \ldots, K-1$, choose a finite grid

$$p_{I,c}^{k+1} = \widehat{p}_{k,0} < \widehat{p}_{k,1} < \cdots < \widehat{p}_{k,N_k} = p_{I,c}^k$$

with mesh smaller than ε. From Lemma 3, for each (k,n), there exists $\mu_{k,n} \in (0,1]$ such that $\widehat{p}_{k,n}$ is the unique maximizer of

$$p\left[\mu_{k,n} D_{I\cap[0,p_{I,c}^k],c}(p) + (1-\mu_{k,n}) D_{I\cap[p_{I,c}^k, p_{I,c}^{k-1}],c}(p)\right].$$

Description of design $\mathcal{M}_c^*$. As a convention, we set $\mu_{0,0} = 0$.

Initialization. At period $t = 1$, $\widehat{M}_1 = \{\varnothing, m_\ell, m_r\}$ and $\mathcal{M}_1(\emptyset)[m_\ell] = 1$.

Block 1: Learning in $[p_{I,c}^2, p_{I,c}^1]$. The block lasts N_1 periods. At period $t = 2$, $\widehat{M}_2 = \{\varnothing, m_\ell, m_r\}$. We set $\mathcal{M}_2((m_\ell, p_{I,c}^1, R))[m_\ell] = 1$, $\mathcal{M}_2((m_\ell, p_{I,c}^1, A))[m_\ell] = \phi_{1,N_1-1}$, $\mathcal{M}_2((m_\ell, p_{I,c}^1, A))[m_r] = 1 - \phi_{1,N_1-1}$, and, at all other histories h^1, $\mathcal{M}_2(h^1)[\varnothing] = 1$, where

$$\phi_{1,N_1-1} = \frac{(1-\mu_{1,N_1-1})\mathbb{P}((m_\ell, p_{I,c}^1, R))}{\mu_{1,N_1-1}\mathbb{P}((m_\ell, p_{I,c}^1, A))}.$$

This is well-defined since

$$\mu_{1,N_1-1} \geq \frac{\mathbb{P}((m_\ell, p_{I,c}^1, R))}{\mathbb{P}((m_\ell, p_{I,c}^1, R)) + \mathbb{P}((m_\ell, p_{I,c}^1, A))} = \frac{\mathbb{P}_{f_{I,c}}(I_c \cap [0, p_{I,c}^1])}{\mathbb{P}_{f_{I,c}}(I_c \cap [0, p_{I,c}^1]) + \mathbb{P}_{f_{I,c}}(I_c \cap [p_{I,c}^1, 1])}$$

by Lemma 3. Thus, conditional on the message m_ℓ, the seller assigns probability μ_{1,N_1-1} to the history $(m_\ell, p_{I,c}^1, R)$ and prices at $\widehat{p}_{1,N_1-1}$.

For periods $2 < t \leq N_1 + 1$, the set of messages is $\widehat{M}_t = \{\varnothing, m_\ell, m_r\}$, and

$$\mathcal{M}_t(h^{t-1})[m_t] = \begin{cases} 1 & h_1^{t-1} = (m_\ell, p_{I,c}^1, R) \;\&\; m_t = m_\ell, \\ \phi_{1,N_1+1-t} & h_1^{t-1} = (m_\ell, p_{I,c}^1, A) \;\&\; m_t = m_\ell, \\ 1 - \phi_{1,N_1+1-t} & h_1^{t-1} = (m_\ell, p_{I,c}^1, A) \;\&\; m_t = m_r, \\ 1 & h_1^{t-1} \notin \{(m_\ell, p_{I,c}^1, A), (m_\ell, p_{I,c}^1, R)\} \;\&\; m_t = \varnothing, \end{cases}$$

where h_1^{t-1} is the first coordinate of h^{t-1} and ϕ_{1,N_1+1-t} is defined as above.

Block $k \in \{2, \ldots, K-1\}$: Learning in $[p_{I,c}^{k+1}, p_{I,c}^k]$. The block lasts N_k periods, from $t = \sum_{k'=1}^{k-1} N_{k'} + 2$ to $t = \sum_{k'=1}^{k} N_{k'} + 1$. The construction is nearly identical to the one for Block 1. We partition histories into three subsets. Membership is determined by the history up to the end of Block $k-1$ and is unaffected by messages and purchase decisions within Block k. The first subset H_R^t includes all histories where the message m_ℓ was always sent through the end of Block $k-1$, the prices offered up to the last period of Block $k-1$ were $p_{I,c}^k < \cdots < \widehat{p}_{k',n'} < \cdots < \widehat{p}_{1,N_1} = p_{I,c}^1$, and none were accepted. The second subset H_A^t includes all histories where the message m_ℓ was always sent through the end of Block $k-1$, the prices offered up to the last period of Block $k-1$ were $p_{I,c}^k < \cdots < \widehat{p}_{k',n'} < \cdots < \widehat{p}_{1,N_1} = p_{I,c}^1$, the offer $p_{I,c}^{k-1}$ was rejected, the offer $p_{I,c}^k$ was accepted, and there is no pair of prices such that the higher price was accepted and the lower price rejected – these histories are consistent. The third subset includes all other histories.

For each grid index n, define

$$\phi_{k,n} = \frac{1 - \mu_{k,n}}{\mu_{k,n}} \frac{F_{I,c}(p_{I,c}^k)}{F_{I,c}(p_{I,c}^{k-1}) - F_{I,c}(p_{I,c}^k)}.$$

By Lemma 3, $\phi_{k,n} \in [0, 1]$. These probabilities are calibrated using the two classes at the beginning of the block.

At period t of Block k, let $n_t = \sum_{k'=1}^{k} N_{k'} + 1 - t$ denote the current grid index. At histories in H_R^t, the message m_ℓ is sent with probability 1. At histories in H_A^t, the message m_ℓ is sent with probability ϕ_{k,n_t} and the message m_r with the complementary probability.

At all other histories, the dummy message $\varnothing$ is sent with probability one. This completes the description of the mechanism.

Learning implication of design $\mathcal{M}_c^*$. By construction, conditional on m_ℓ at period t, a seller assigns probability μ_{k,n_t} to valuations being in $I_c \cap [0, p_{I,c}^k]$ and the complementary probability to $I_c \cap [p_{I,c}^k, p_{I,c}^{k-1}]$. Therefore, conditional on message m_ℓ, the seller posts the price $\widehat{p}_{k,n_t}$. Suppose then that $v_c^* \in [\widehat{p}_{k,n}, \widehat{p}_{k',n'})$ for two consecutive prices in the grid. The first consumer rejects the offer $p_{I,c}^1$, thus inducing the message m_ℓ and the price $\widehat{p}_{1,N_1-1}$, with probability one. Each subsequent offer above v_c^* is rejected, and the next price in the prescribed grid is then offered with probability one. After the last rejected price $\widehat{p}_{k',n'}$, the price $\widehat{p}_{k,n}$ is offered and accepted. The platform has thus learned that v_c^* is in $[\widehat{p}_{k,n}, \widehat{p}_{k',n'})$, an interval of length smaller than ε. The history $(m_\ell, p_{I,c}^1, R, \ldots, m_\ell, \widehat{p}_{k,n}, A)$ has probability one, conditional on $v_c \in [\widehat{p}_{k,n}, \widehat{p}_{k',n'})$. If instead $v_c^* \in [\inf I_c, p_{I,c}^K)$, all offers up to and including $p_{I,c}^K$ are rejected. At that point, the platform knows that $v_c^* \in [\inf I_c, p_{I,c}^K)$, an interval of length smaller than ε. Including the initial offer, the construction takes at most $t_c^* = 1 + \sum_{k=1}^{K-1} N_k < \infty$ periods, where the sum is empty when $K = 1$. □

3. Multiple Characteristics and Design $\mathcal{M}_{I,c^*}^{**}$. We now construct an information design $\mathcal{M}_{I,c^*}^{**}$ that ε-learns valuations $v_c^* \in [\inf I_c, p_{I,c^*}^1)$ for every characteristic c, where $c^* \in \arg\max_{c' \in \mathcal{C}} p_{I,c'}^1$ (Proposition 3). The main idea is to induce sellers, upon receiving a particular message, to assign arbitrarily high probability to facing a consumer with characteristic c^*, regardless of the actual characteristic. Sellers therefore price nearly as if they were facing such consumers. Learning is asymptotic.

The design repeats two phases. In Phase I, the platform reveals the current consumer's characteristic and runs $\mathcal{M}_c^*$ separately for each characteristic, using only the history of consumers with that characteristic. In Phase II, the platform uses the interval identified for c^* on the event that all Phase I consumers have that characteristic. By occasionally sending these reports when the current consumer has another characteristic, the platform induces prices that extend learning for the other characteristics to valuations below p_{I,c^*}^1.

To run the separate mechanisms in Phase I, sellers must know how many times the mechanism for the current characteristic has already been run. Since consumers' arrival times are random, calendar time alone does not provide this information. We therefore index messages by the local period s, writing $\widehat{M}_s = \{\varnothing_s, m_{\ell,s}, m_{r,s}\}$. If t_c consumers with characteristic c have arrived before calendar period t, the message for a new consumer with that characteristic belongs to $\widehat{M}_{t_c+1}$. We denote by $\Gamma_{\mathcal{M}_c^*}$ the game induced by $\mathcal{M}_c^*$, and by $\widehat{H}_c^s$ its set of histories after s periods.

Proposition 3. *For any distribution F_I, choose $c^* \in \arg\max_{c\in\mathcal{C}} p^1_{I,c}$. Assume that $\inf I_{c^*} \leq p^1_{I,c}$ for all $c \in \mathcal{C}$. There exists a design $\mathcal{M}^{**}_{I,c^*}$ such that the platform ε-learns valuations in $[\inf I_c, p^1_{I,c^*})$ for all $c \in \mathcal{C}$.*

Proof. The proof is constructive.

Description of design $\mathcal{M}^{**}_{I,c^*}$. We repeat the following design infinitely often, using independent randomization across repetitions. Each repetition consists of two distinct phases.

Phase I. Fix $\varepsilon > 0$. Phase I lasts t^* calendar periods. For each c, let t_c^* be the bound from Proposition 2 on the number of arrivals required to $\varepsilon/5$-learn all valuations in $[\inf I_c, p^1_{I,c})$ with the mechanism $\mathcal{M}_c^*$ and set $t^* = \max_c t_c^*$. In the notation of Proposition 2, $t_c^* = 1 + \sum_{k=1}^{K-1} N_k$, where K and N_k are chosen for characteristic c and accuracy $\varepsilon/5$, with the sum understood to be zero when $K = 1$. At each period t of Phase I, the set of messages $M_t = \mathcal{C} \times \bigcup_{s\leq t} \widehat{M}_s$.

We define $\rho_c^{t-1} : H^{t-1} \to \bigcup_{s=0}^{s=t-1} \widehat{H}_c^s$ by $\rho_c^{t-1}((c_s, (\widehat{c}_s, \widehat{m}_s), p_s, a_s)_{s=1}^{t-1}) = (\widehat{m}_s, p_s, a_s)_{s:c_s=\widehat{c}_s=c}$. With the history h^{t-1}, we associate the (possibly empty) history $\widehat{h}_c^{t_c}$ of $\Gamma_{\mathcal{M}_c^*}$, where the first coordinate of $\widehat{h}_c^{t_c}$ corresponds to the first occurrence of $(c, (c, \widehat{m}_s), p_s, a_s)$, the second coordinate to the second occurrence, etc. For all histories h^{t-1} during Phase I with $t_c < t_c^*$, $\mathcal{M}_t(h^{t-1}, c)(c, \widehat{m}) = \mathcal{M}_c^*(\rho_c^{t-1}(h^{t-1}))[\widehat{m}]$.[21] If $t_c \geq t_c^*$, the design sends $(c, \varnothing_{t_c+1})$ with probability one. That is, the mechanism treats each characteristic independently, reveals the

[21] We slightly abuse notation in that $\mathcal{M}_c^*(\rho_c^{t-1}(h^{t-1}))[\widehat{m}]$ is defined only for $\widehat{m} \in \widehat{M}_{t_c+1}$, where $t_c \leq t-1$ is the number of times a consumer with characteristic c entered the market by time $t-1$ at history h^{t-1}. We extend it to all other messages by assigning a zero probability to them.

characteristics of newly entered consumers to the current sellers, and runs $\mathcal{M}_c^*$, when the characteristic is c.

PHASE II. Let $H_{c^*}^{t^*}$ be the subset of histories at the end of Phase I such that $c_s = \widehat{c}_s = c^*$ for all $s = 1, \ldots, t^*$. The probability of this subset is $\lambda_{c^*}^{t^*} > 0$. Write $\mathcal{I}_{c^*}$ for the set of intervals $\{I_{c^*}(h_{c^*}^{t^*}) : h_{c^*}^{t^*} \in \text{supp } \mathbb{P}(\cdot|H_{c^*}^{t^*})\}$. These intervals form a finite partition of the original support I_{c^*}. They consist of the intervals between consecutive prices in the c^* construction, together with the bottom interval $I_{c^*} \cap [0, p_{I,c^*}^K)$, when nonempty, and the top interval $I_{c^*} \cap [p_{I,c^*}^1, 1]$. Conditional on $H_{c^*}^{t^*}$, the identified interval is the unique member of this partition containing $v_{c^*}^*$. Subsequent Phase I offers, including those after completion of the local procedure, do not refine it. Let ν be an arbitrary distribution fully supported on $\mathcal{I}_{c^*}$. For a reported interval $J \in \mathcal{I}_{c^*}$, write $I \cap J$ for the rectangle obtained from I by restricting its c^* coordinate to J.

Throughout Phase II, let

$$M_t = \Big(\mathcal{C} \times (\mathcal{I}_{c^*} \cup \{[0,1]\})\Big) \cup \{\varnothing\}, \qquad t^* + 1 \leq t \leq (C+1)t^*.$$

Let $h_{c^*}^{t^*} \in H_{c^*}^{t^*}$ be a history at the end of Phase I. At the first period of Phase II, $t = t^* + 1$, the history is simply $h^{t-1} = h_{c^*}^{t^*}$. At any subsequent period $t \leq (C+1)t^*$, consider any history extending $h_{c^*}^{t^*}$, that is, $h^{t-1} = \Big(h_{c^*}^{t^*}, (c_s, m_s, p_s, a_s)_{s=t^*+1}^{t-1}\Big)$. At each of these histories, the platform uses the same rule: (i) $\mathcal{M}_t(h^{t-1}, c)[c^*, J] = \phi\nu(J)$, (ii) $\mathcal{M}_t(h^{t-1}, c)[c, [0,1]] = 1 - \phi$, and (iii) $\mathcal{M}_t(h^{t-1}, c^*)[c^*, I_{c^*}(h_{c^*}^{t^*})] = 1$. Rules (i)–(ii) apply to $c \neq c^*$, with rule (i) applying to every $J \in \mathcal{I}_{c^*}$. Rule (iii) applies when the current consumer has characteristic c^*; the platform reports the interval identified at the end of Phase I.

At all histories that do not extend a history in $H_{c^*}^{t^*}$, the platform sends $\varnothing$ with probability one. Phase II lasts Ct^* periods. This completes the description of the design.

IMPLICATION OF THE DESIGN $\mathcal{M}_{I,c^*}^{**}$.

Write $\mathbb{P}_{I,f}$ for probability under F_I. Conditional on $H_{c^*}^{t^*}$, a consumer with characteristic c^* generates the message (c^*, J) exactly when $v_{c^*} \in J$. A consumer with characteristic $c \neq c^*$ generates the same message with probability $\phi\nu(J)$, independently of valuations. The

event $H_{c^*}^{t^*}$ depends only on Phase I arrivals, so its probability $\lambda_{c^*}^{t^*}$ cancels from Bayes' rule. Consequently, upon observing (c^*, J), the seller chooses p to maximize

$$p\left[\lambda_{c^*}\mathbb{P}_{I,f}(v_{c^*} \in J)D_{I\cap J,c^*}(p) + \phi\nu(J)\sum_{c\neq c^*}\lambda_c D_{I,c}(p)\right].$$

To obtain demand conditional on this message, we divide the expression in brackets by $Z = \lambda_{c^*}\mathbb{P}_{I,f}(v_{c^*} \in J) + \phi\nu(J)\sum_{c\neq c^*}\lambda_c$, the probability of receiving the message conditional on $H_{c^*}^{t^*}$. Since Z does not depend on p, this division leaves the optimal price unchanged. Conditional demand is therefore a mixture of $D_{I\cap J,c^*}$ and demand from the other characteristics, with weight $\lambda_{c^*}\mathbb{P}_{I,f}(v_{c^*} \in J)/Z$ on the former. This weight approaches one as $\phi \to 0$. We now show that, for sufficiently small ϕ, the seller prices arbitrarily close to $\arg\max_p pD_{I\cap J,c^*}(p)$. We need the following result.

Lemma 4. *Consider two continuous demand functions* $\widetilde{D}, D^* : [0,1] \to [0,1]$ *such that* $\{p^*\} := \arg\max_{p\in[0,1]} pD^*(p)$. *For* $\alpha \in [0,1]$, *define* $p(\alpha) := \min\arg\max_{p\in[0,1]} p\left((1-\alpha)\widetilde{D}(p) + \alpha D^*(p)\right)$. *As* $\alpha \to 1$, $p(\alpha) \to p^*$.

Proof of Lemma 4. We have the following:

$$\left|p\left((1-\alpha)\widetilde{D}(p) + \alpha D^*(p)\right) - pD^*(p)\right| \leq (1-\alpha)D^*(p) + (1-\alpha)\widetilde{D}(p) \leq 2(1-\alpha),$$

where we use the fact that prices and demands are in $[0,1]$. Thus, the convergence is uniform. Since $p \mapsto pD^*(p)$ has the unique maximizer p^*, the technical Lemma 5 stated and proved below implies that $p(\alpha) \to p^*$. The same argument shows that every maximizer is close to p^* when α is sufficiently close to one. □

From now on, K, N_k, and $\widehat{p}_{k,n}$ refer to the construction for c^*. Apart from the bottom and top intervals, the elements of $\mathcal{I}_{c^*}$ are intervals between consecutive prices in that construction.

Consider $J = [\widehat{p}_{k,n}, \widehat{p}_{k,n+1})$, with $0 \leq n < N_k$. Its conditional monopoly price is unique by log-concavity and lies in its closure. By Lemma 4, for sufficiently small ϕ, the price $\widehat{p}_t$ induced by the message (c^*, J) is within $\varepsilon/5$ of that monopoly price. Since the interval has length smaller than $\varepsilon/5$, it follows that $|\widehat{p}_t - \widehat{p}_{k,n}| < 2\varepsilon/5$.

For the neighboring interval $J' = [\widehat{p}_{k,n+1}, \widehat{p}_{k,n+2})$, with $n+2 \leq N_k$, the same argument gives $|\widehat{p}'_t - \widehat{p}_{k,n+1}| < 2\varepsilon/5$. Consequently,

$$|\widehat{p}_t - \widehat{p}'_t| \leq |\widehat{p}_t - \widehat{p}_{k,n}| + |\widehat{p}_{k,n} - \widehat{p}_{k,n+1}| + |\widehat{p}_{k,n+1} - \widehat{p}'_t| < \varepsilon.$$

The same bounds apply to consecutive intervals across block boundaries and to the bottom interval $[\inf I_{c^*}, p^K_{I,c^*})$, when nonempty: its length is smaller than $\varepsilon/5$, and its conditional monopoly price lies in its closure. This includes the case $K = 1$.

For the top interval of the Phase I partition, the induced price is exactly p^1_{I,c^*} for sufficiently small ϕ. Indeed, below p^1_{I,c^*} the unnormalized profit contributed by characteristic c^* is linear with strictly positive slope; the other characteristics' profit terms have bounded slopes and are multiplied by ϕ. Above p^1_{I,c^*}, every component's profit is decreasing, since $p^1_{I,c} \leq p^1_{I,c^*}$ for all c. The platform can send this top-interval message even when the current consumer has characteristic $c \neq c^*$. The seller then charges p^1_{I,c^*}, which the consumer rejects whenever $v^*_c < p^1_{I,c^*}$. Observing this rejection together with acceptance of a lower price less than ε away allows the platform to locate the valuation in an interval of length smaller than ε. This ensures that valuations close to the upper endpoint p^1_{I,c^*} can also be learned.

Note that the choice of $\phi > 0$ depends on the message (c^*, J). However, since there are finitely many intervals in $\mathcal{I}_{c^*}$, we can choose a uniform ϕ so that, after sorting the induced prices and adjoining the initial price $p^1_{I,c}$, every realized valuation in $[p^1_{I,c}, p^1_{I,c^*})$ is bracketed by prices less than ε apart. The initial price $p^1_{I,c}$ can be observed in Phase I of a different repetition.

Finally, suppose that $v^*_c \in [\inf I_c, p^1_{I,c})$. By construction of the mechanism $\mathcal{M}^*_c$ and the choice of $t^* \geq t^*_c$, there is a strictly positive probability that, at the end of Phase I, the platform's posterior over v_c assigns probability one to $[v^*_c - \varepsilon, v^*_c + \varepsilon]$. Indeed, the event that all t^* consumers arriving during Phase I have characteristic c has probability $\lambda_c^{t^*} > 0$.

Suppose that $v^*_c \in [p^1_{I,c}, p^1_{I,c^*})$, so $c \neq c^*$. Recall that $\inf I_{c^*} \leq p^1_{I,c}$. At each period t in Phase II, conditional on any on-path successor h^{t-1} of $h^{t^*}_{c^*} \in H^{t^*}_{c^*}$, there is a strictly positive

probability $\lambda_c \phi \nu(J) > 0$ that the newly entered consumer has characteristic c and the message sent to the seller is (c^*, J). This probability is independent of $v^*_{c'}$ for all $c' \in \mathcal{C}$. Conditional on the message (c^*, J), the seller offers the price $\widehat{p}_t$. For each message, we select the same optimal price in every period and repetition. Recall that $|\widehat{p}_t - \widehat{p}_{k,n}| < 2\varepsilon/5$ when $J = [\widehat{p}_{k,n}, \widehat{p}_{k',n'})$ is an interval between consecutive construction prices.

Next, if this offer is accepted and $t < (C+1)t^*$, conditional on the history $(h^{t-1}, (c, (c^*, J), \widehat{p}_t, A))$, there is a strictly positive probability $\lambda_c \phi \nu(J') > 0$ that the newly entered consumer also has characteristic c and the message sent to the seller is (c^*, J'). Conditional on this message, the seller offers the price $\widehat{p}_{t+1}$. By the preceding grid argument, whenever v^*_c lies between two Phase II prices, the messages can be chosen so that $\widehat{p}_t \leq v^*_c < \widehat{p}_{t+1}$ and $\widehat{p}_{t+1} - \widehat{p}_t < \varepsilon$. Thus, with strictly positive probability, the platform observes a consumer with characteristic c purchasing at $\widehat{p}_t$ and another consumer with that characteristic rejecting the offer $\widehat{p}_{t+1}$. It therefore learns that $v^*_c \in [\widehat{p}_t, \widehat{p}_{t+1})$.

A similar argument applies after the history $(h^{t-1}, (c, (c^*, J), \widehat{p}_t, R))$, with rejection of the higher price followed by acceptance of the lower price. If instead the two prices supplied by the grid argument are $p^1_{I,c}$ and a Phase II price $\widehat{p}_{t+1}$, acceptance of $p^1_{I,c}$ is observed when the first Phase I consumer has characteristic c, an event of probability $\lambda_c > 0$. Rejection of $\widehat{p}_{t+1}$ can be observed in another repetition. Together, these observations reveal that $v^*_c \in [p^1_{I,c}, \widehat{p}_{t+1})$, an interval of length smaller than ε.

Finally, the set of histories $H^{t^*}_{c^*}$ has strictly positive probability $\lambda^{t^*}_{c^*}$: all we need is for consumers with characteristic c^* to arrive in all t^* periods of Phase I. There are at least two periods in Phase II when $c \neq c^*$. Therefore, each pair of Phase II offers described above can be observed with strictly positive probability in a repetition. When an observation from Phase I is also required, it likewise occurs with strictly positive probability. Conditional on any fixed realized valuation profile, repetitions are independent and these probabilities are the same in every repetition. The second Borel–Cantelli lemma therefore implies that all the required observations eventually occur almost surely. Since the platform retains these observations

and there are finitely many characteristics, it ε-learns every $v_c^* \in [\inf I_c, p^1_{I,c^*})$. $\square$

We conclude with Lemma 5, which is a technical lemma that we used to prove Lemma 4 above.

Lemma 5. *Let $(\varphi_n)_n$ be a sequence of real-valued functions $\varphi_n : [0,1] \to \mathbb{R}$, each attaining its maximum. Suppose that it converges uniformly to a continuous function φ having a unique maximizer x^*. For every $\varepsilon > 0$, there exists n^* such that, for every $n \geq n^*$, all maximizers of φ_n lie within ε of x^*.*

Proof of Lemma 5. Fix $\varepsilon > 0$. If $[0,1] \setminus (x^* - \varepsilon, x^* + \varepsilon)$ is empty, the conclusion is immediate; hence suppose this set is nonempty. Let $\delta := \varphi(x^*) - \sup_{x \in [0,1] \setminus (x^* - \varepsilon, x^* + \varepsilon)} \varphi(x) > 0$, where the strict inequality follows from continuity of φ, compactness of the set over which the supremum is taken, and uniqueness of the maximizer. Choose n^* such that $\|\varphi_n - \varphi\|_\infty \leq \delta/3$ for every $n \geq n^*$. For all $n \geq n^*$, for all $x \in [0,1] \setminus (x^* - \varepsilon, x^* + \varepsilon)$, we have

$$\varphi_n(x) \leq \varphi(x) + \frac{\delta}{3} \leq \varphi(x^*) - \frac{2\delta}{3} < \varphi(x^*) - \frac{\delta}{3} \leq \varphi_n(x^*),$$

hence no maximizer lies in $[0,1] \setminus (x^* - \varepsilon, x^* + \varepsilon)$. $\square$

4. The Grand Mechanism

We now construct a grand mechanism that repeats a finite sequence of blocks, following the increasing threshold sequence (p^k) defined in (2)–(3) of the main text. In Block k, the platform sends messages that induce sellers to price approximately as if all valuations lay in $[p^{k-1}, 1]^C$, regardless of the realized history. Each block follows the two phases described in Proposition 3. This gives a positive probability of extending learning to valuations below a threshold close to p^k. We choose enough blocks to bring the final threshold close to p^∞. Repeating the sequence ensures that the purchase and rejection observations needed for learning eventually occur almost surely.

Choose $k^* < \infty$ such that $p^{k^*-1} < p^{k^*} \leq p^\infty$ and $p^{k^*} \geq p^\infty - \varepsilon/5$. Here, p^k and p^∞ are defined in (3) and (4), respectively. We write c^k for an element of $\arg\max_{c \in \mathcal{C}} p^k_c$. If there

are several maximizers, choose among them the characteristic with the highest actual initial price in Block k, defined below.

Let $\mathcal{I}$ denote the set of rectangles in $[0,1]^C$, i.e., products of intervals. We say that the history h^t is *consistent* if $\sup\{p_s : a_s = 1,\ c_s = c\} < \inf\{p_s : a_s = 0,\ c_s = c\}$ for all c, that is, if the largest price accepted by a consumer with characteristic c is smaller than the smallest price rejected. For a history to be inconsistent, we need a consumer to buy at a price above their valuation or to reject a price no greater than their valuation. With a consistent history $h^t := (c_s, m_s, p_s, a_s)_{s=1}^{s=t}$, we associate the rectangle $I(h^t)$ given by

$$\prod_{c\in\mathcal{C}}\Big([\sup\{p_s : a_s = 1,\ c_s = c\}, \inf\{p_s : a_s = 0,\ c_s = c\}) \cap [0,1]\Big),$$

with the convention that $I(\emptyset) = [0,1]^C$.

Description of the Grand Design. The information design repeats the sequence of k^* blocks defined below, using independent randomization across repetitions. We describe one repetition, with time, histories, and arrival counts measured from its beginning. In what follows, the prices $\hat{p}^k \in (0,1)$ are taken as given. We define them later on, along with all the relevant parameters.

Write $\mathcal{M}^*_{I,c}$ for the single-characteristic construction of Proposition 2 under $F_{I,c}$, and $t^*_c(I)$ for its local time bound at the accuracy used in the block. Set $\hat{p}^0 = 0$. For Block k, let $\tau_k = \max_{c\in\mathcal{C}} t^*_c(I^{k-1})$ and $L_k = (C+1)\tau_k$, where the rectangles I^{k-1} are defined below.

Block 1. Let $I^0 = [0,1]^C$. Block 1 starts at $T_1 = 1$ and runs $\mathcal{M}^{**}_{I^0,c^1}$ for L_1 periods.

Block $k = 2, \ldots, k^*$. Let $I^{k-1} = [\hat{p}^{k-1}, 1]^C$. Block k starts at period $T_k = 1 + \sum_{j=1}^{k-1} L_j$. Let $\widehat{H}^{T_k-1} := \{h^{T_k-1} \in H^{T_k-1} : I(h^{T_k-1}) = I^{k-1}\}$. In Block k, the platform sends a message that makes sellers assign arbitrarily high probability to $\widehat{H}^{T_k-1}$, regardless of the actual consistent history. The message provides no further information about play in earlier blocks. Following this message, we run $\mathcal{M}^{**}_{I^{k-1},c^k}$ for L_k periods. Thus, Block k has two distinct phases. Phase I lasts τ_k periods and Phase II lasts $C\tau_k$ periods.

More formally, in Phase I, the set of messages is $M_t = \mathcal{I} \times \mathcal{C} \times \bigcup_{s=1}^{t-T_k+1} \widehat{M}_s$. Define $\rho_c^{k,t-1} : H^{t-1} \to \bigcup_{s=0}^{s=t-T_k} \widehat{H}_c^s$ by

$$\rho_c^{k,t-1}(h^{T_k-1}, (c_s, \underbrace{(\widehat{I}_s, \widehat{c}_s, \widehat{m}_s)}_{m_s}, p_s, a_s)_{s=T_k}^{t-1}) = (\widehat{m}_s, p_s, a_s)_{s\in\{T_k,\ldots,t-1\}:c_s=\widehat{c}_s=c},$$

with the convention that $\rho_c^{k,T_k-1}(h^{T_k-1}) = \emptyset$. Let $t_{k,c}$ denote the length of $\rho_c^{k,t-1}(h^{t-1})$, the number of retained observations of characteristic c in the current block. The next local message belongs to $\widehat{M}_{t_{k,c}+1}$; the local mechanism assigns zero probability to other messages.

At $t = T_k$, for all consistent histories h^{t-1}, for all c, we let

$$\mathcal{M}_t(h^{t-1}, c)[I(h^{T_k-1}), c, \widehat{m}] = (1-\phi)\mathcal{M}^*_{I^{k-1},c}(\rho_c^{k,t-1}(h^{t-1}))[\widehat{m}],$$
$$\mathcal{M}_t(h^{t-1}, c)[I^{k-1}, c, \widehat{m}] = \phi\mathcal{M}^*_{I^{k-1},c}(\rho_c^{k,t-1}(h^{t-1}))[\widehat{m}],$$

for all $(I, c, \widehat{m})$.[22] For all inconsistent histories h^{t-1}, $\mathcal{M}_t(h^{t-1}, c)$ assigns probability one to $\left([0,1]^C, c, \varnothing_{t_{k,c}+1}\right)$. In words, Phase I treats each characteristic separately and runs $\mathcal{M}^*_{I^{k-1},c}$ on each of them, independently. The only difference from one block to the next is that sellers assign different "prior" probabilities to the valuations.

The rectangle included in the message is chosen only at the beginning of the block and remains the same throughout both phases. Conditional on this rectangle, follow Propositions 2 and 3, using the actual prices to define the histories and intervals, and the probabilities of these histories to calculate the probabilities of sending each message. The probability used to pool characteristics in Proposition 3 is separate from ϕ and is chosen after fixing ϕ and the Phase I construction.

The same reporting device is used in Phase II. Sellers again assign arbitrarily high probability to $\widehat{H}^{T_k-1}$, regardless of the actual consistent history. The platform then uses the Phase I histories to induce learning in $[\hat{p}_c^k, \hat{p}^k)$.

[22]Notice that, by construction, for all consistent successors of histories in $\widehat{H}^{T_k-1}$, the sellers observe the message I^{k-1} with probability one.

5. The Design Guarantees ε-learning

We choose the parameters of the grand mechanism and then derive the implication for learning.

CHOOSING THE PARAMETERS. Under the conditional prior $F_{I^{k-1}}$, Propositions 2 and 3 provide a design that ε-learns all valuations in $[\inf I_c^{k-1}, \max_{c'} p^1_{I^{k-1},c'})$ for each characteristic c. In particular, if $I^{k-1} = [p^{k-1}, 1]^C$, this gives learning on $[p^{k-1}, p^k)^C$.

In the actual design, sellers are only nearly convinced that valuations lie in $I^{k-1} = [\hat{p}^{k-1}, 1]^C$. Write $\hat{p}^k_c$ for the initial price induced in Phase I for characteristic c, conditional on the report I^{k-1}, and set $\hat{p}^k := \hat{p}^k_{c^k}$ and $I^k = [\hat{p}^k, 1]^C$. The parameter choice must therefore control two approximations. The first approximation is of the benchmark threshold p^k by the induced initial price $\hat{p}^k$. Lemma 6 shows how to control this approximation.

Let $p^0 = 0$ and define inductively for $k \geq 1$:

$$p^k(\alpha_1, \ldots, \alpha_k) = \min \arg \max_{p \in [0,1]} p\Big[\alpha_k D_{[p^{k-1}(\alpha_1,\ldots,\alpha_{k-1}),1]^C, c^k}(p) + (1-\alpha_k)\widetilde{D}_k(p)\Big],$$

with $\alpha_k \in [0,1]$ and $\widetilde{D}_k : [0,1] \to [0,1]$ a continuous demand function. The posterior-mixture demands used below are continuous, since the history-conditioned valuation distributions are atomless.

Lemma 6. *For all $k \geq 1$, $p^k(\alpha_1, \ldots, \alpha_k) \to p^k$ as $\alpha_\ell \to 1$ for all $\ell \leq k$.*

Proof of Lemma 6. The proof is by induction on k. For $k = 1$, since prices and demands lie in $[0,1]$, we have, for every $p \in [0,1]$,

$$\left| p\left(\alpha_1 D_{[0,1]^C, c^1}(p) + (1-\alpha_1)\widetilde{D}_1(p)\right) - pD_{[0,1]^C, c^1}(p)\right| \leq 1 - \alpha_1.$$

Thus, the profit functions converge uniformly to $pD_{[0,1]^C, c^1}(p)$, whose unique maximizer is p^1. Lemma 5 therefore gives $p^1(\alpha_1) \to p^1$.

For $k \geq 2$, the induction hypothesis and uniform continuity of F imply that

$D_{[p^{k-1}(\alpha_1,\ldots,\alpha_{k-1}),1]^C,c^k}$ converges uniformly to $D_{[p^{k-1},1]^C,c^k}$, since the limiting conditioning rectangle has positive probability. Moreover, for every $p \in [0,1]$,

$$\begin{aligned}&\left|p\Big(\alpha_k D_{[p^{k-1}(\alpha_1,\ldots,\alpha_{k-1}),1]^C,c^k}(p) + (1-\alpha_k)\widetilde{D}_k(p)\Big) - pD_{[p^{k-1},1]^C,c^k}(p)\right| \\ &\leq 1-\alpha_k + \left|D_{[p^{k-1}(\alpha_1,\ldots,\alpha_{k-1}),1]^C,c^k}(p) - D_{[p^{k-1},1]^C,c^k}(p)\right|.\end{aligned}$$

The right-hand side converges uniformly to zero as $\alpha_\ell \to 1$ for all $\ell \leq k$. Since the limiting profit function has the unique maximizer p^k, Lemma 5 gives $p^k(\alpha_1,\ldots,\alpha_k) \to p^k$, completing the induction. □

Lemma 6 controls the first approximation. If $\hat{p}^{k-1}$ is sufficiently close to p^{k-1}, we can choose the parameters so that the induced initial price $\hat{p}^k$ is arbitrarily close to p^k. Here $p^k = p^1_{[p^{k-1},1]^C,c^k}$ is the monopoly price for characteristic c^k conditional on all valuations lying in $[p^{k-1},1]^C$.

The second approximation is of the price grids constructed under $F_{I^{k-1}}$ by the prices induced under sellers' actual beliefs. The following arguments show how to control this approximation so that the gaps between consecutive induced prices remain small enough to ensure ε-learning. Fix a block $k \leq k^*$ and its rectangle $I^{k-1} = [\hat{p}^{k-1},1]^C$. By our choice of k^*, when $\hat{p}^{k-1}$ is sufficiently close to p^{k-1}, we have $p^1_{I^{k-1},c^k} > \hat{p}^{k-1}$. Thus, this monopoly price is interior and satisfies the first-order condition.

For each characteristic c, use the decreasing price sequence $(p^j_{I^{k-1},c})_j$ defined in (8). As in Proposition 2, index its steps by $j \geq 1$ and the grid prices by n, suppressing dependence on c in the grid prices and mixing parameters. At each step with $p^j_{I^{k-1},c} > \inf I^{k-1}_c$, Lemma 3 provides parameters $\mu^k_{j,n} \in (0,1]$ that induce a grid of prices $p^k_{j,n}$ in $[p^{j+1}_{I^{k-1},c}, p^j_{I^{k-1},c}]$ with mesh $\varepsilon/15$.

Choose the probabilities block by block, keeping earlier blocks fixed. Suppose that $\mathbb{P}(\widehat{H}^{T_k-1}) > 0$. Each history in $\widehat{H}^{T_k-1}$ has the same probability conditional on any valuation

profile in I^{k-1}. Conditional on the message specifying I^{k-1} at the beginning of the block, the distribution of valuations is therefore

$$\frac{(1-\phi)\mathbb{P}(\widehat{H}^{T_k-1})F_{I^{k-1}}+\phi F}{(1-\phi)\mathbb{P}(\widehat{H}^{T_k-1})+\phi}.$$

With earlier blocks fixed, the weight on F converges to zero as ϕ decreases to zero.

Apply Proposition 2 using the actual prices in the history conditions and the probabilities of these histories to calculate the probabilities of sending each message. These probabilities remain in $[0,1]$: for every offered grid price with $\mu^k_{j,n}<1$, the corresponding probability under $F_{I^{k-1}}$ is strictly below one, while $\mu^k_{j,n}=1$ requires probability zero. If preceding prices and probabilities converge to those under $F_{I^{k-1}}$, the probabilities of the histories satisfying these conditions also converge. The demands conditional on these histories consequently converge uniformly, provided the corresponding histories under $F_{I^{k-1}}$ have positive probability. Lemma 5 then implies convergence of the next optimal price. Induction over the finitely many offers therefore gives $|\hat{p}^k_{j,n}-p^k_{j,n}|<\varepsilon/15$, so consecutive induced prices are less than $\varepsilon/5$ apart. Phase II uses the intervals defined by these actual prices. After fixing ϕ and the Phase I construction, choose the probability used to pool characteristics sufficiently small. The same argument then gives the price approximations used in Proposition 3.

We also verify that the price induced by the top-interval message in Phase II is exactly $\hat{p}^k$. For sufficiently small ϕ, the seller's profit for each characteristic under the displayed distribution is single-peaked. In a neighborhood of an interior monopoly price under $F_{I^{k-1}}$, log-concavity implies that the second derivative of profit is bounded above by a strictly negative constant almost everywhere. Under F, the second derivative of profit is bounded in absolute value in this neighborhood. Thus, for sufficiently small ϕ, the profit derivative under the mixture remains strictly decreasing there. Outside this neighborhood, it retains the sign of the profit derivative under $F_{I^{k-1}}$. If the conditional monopoly price is the lower support endpoint, Lemma 1 implies that the prior monopoly price is no higher, so both profit functions decrease above that endpoint.

Choose ϕ small enough that c^k has the highest actual initial price, $\hat{p}^k$. Subsequent Phase I messages do not distinguish valuations above $\hat{p}^k$. Above this price, profit is proportional to profit under the characteristic's initial belief and is decreasing. The top interval for c^k therefore remains $[\hat{p}^k, 1]$. The argument for the top interval in Proposition 3 then implies that its Phase II price is exactly $\hat{p}^k$.

Finally, starting from a history in $\widehat{H}^{T_k-1}$, consider all Phase I arrivals having characteristic c^k and accepting, followed by the top-interval message in every Phase II period, with every characteristic arriving at least once and accepting. This event has positive probability. No offer exceeds $\hat{p}^k$, and every characteristic accepts $\hat{p}^k$, so the resulting rectangle is exactly $I^k = [\hat{p}^k, 1]^C$. Hence $\mathbb{P}(\widehat{H}^{T_{k+1}-1}) > 0$. Starting from $I^0 = [0,1]^C$, this verifies the condition for every subsequent block.

From the construction in the proof of Proposition 3, it follows that we guarantee ε-learning of all valuations in $[\hat{p}^{k-1}, \hat{p}^k)$ with strictly positive probability at the end of the block k, regardless of the realized histories at the end of block $k-1$.

Learning. In each block k, regardless of the preceding consistent history, there is a strictly positive probability that sufficiently many consumers of each characteristic arrive and receive the offers needed for the platform to ε-learn their valuations in $[\hat{p}^{k-1}, \hat{p}^k)$. Therefore, after the k^* blocks, there is a strictly positive probability that the platform has ε-learned valuations in $[0, \hat{p}^{k^*})$. Repetitions are independent conditional on $\mathbf{v}^*$, so the second Borel–Cantelli lemma implies that this learning event eventually occurs almost surely.

B. Long-lived consumers

B.1. Information design and ε-learning

Proof of Theorem 2.

Construction.

Throughout this proof, write $\delta = \delta_B$. Before giving the construction, we describe its *equilibrium path*. The grid $0 < b_0 < b_1 < \cdots < b_N < 1$, the prices $p_1, \ldots, p_N$, and the function r will be defined below. If $v \in [0, b_1)$, the first offer p_1 is rejected and every later offer is b_0, which is accepted if and only if $v \geq b_0$. If $v \in [b_n, b_{n+1})$ for some $n = 1, \ldots, N-1$, the offers $p_1, \ldots, p_n$ are accepted and p_{n+1} is rejected. The same consumer then purchases at $r(b_n)$ and exits. His replacement, and every subsequent consumer, purchases at b_n. Finally, if $v \in [b_N, 1]$, the offers $p_1, \ldots, p_N$ are accepted, after which b_N is offered and accepted forever. Thus, for every v, the price and the posterior become constant after finitely many periods. If $v \in [b_n, b_{n+1})$, the limiting posterior is F conditional on that interval, and the price is b_n. We now describe the construction.

Defining the grid. Let $\varepsilon > 0$. For δ sufficiently close to one, we construct a finite grid

$$0 < b_0 < b_1 < \cdots < b_N < 1,$$

of thresholds, where $b_1 < \varepsilon$, $b_{n+1} - b_n < \varepsilon$ for $n = 1, \ldots, N-1$, and $b_N > 1 - \varepsilon$.

Define the function $r : [0,1] \to [0,1]$ as

$$r(v) := \begin{cases} v, & 0 \leq v < b_0, \\ (1-\delta)v + \delta b_0, & b_0 \leq v \leq 1, \end{cases} \tag{9}$$

and note that r is increasing, concave, and continuous.

For $a, b \in [0,1]$, write $p(a,b) := (1-\delta)b + \delta r(a)$ and, for $n = 1, \ldots, N$, define $p_n := p(b_{n-1}, b_n)$.

For every $b \in (0,1]$, let $\widehat{p}(b)$ be the solution to

$$\max_{p \in [0,b]} p \frac{F(b) - F(p)}{F(b) - F(0)}.$$

Write $p^1 = \widehat{p}(1)$ for the static monopoly price, i.e., the maximizer of $p[1 - F(p)]$.

Lemma 7. *For every $\varepsilon > 0$, there exists $\underline{\delta} < 1$ such that, for all $\delta \in [\underline{\delta}, 1)$, we can choose $b_1 \in (0, p^1)$ to satisfy*

$$\Big[(1-\delta)b_1 + \delta \widehat{p}(b_1)\Big] f(b_1) = (1-\delta)[1 - F(b_1)] \qquad \textit{and} \qquad b_1 < \varepsilon.$$

Proof of Lemma 7. We already know that $\widehat{p}(b)$ is unique, interior, and continuous as a function of b. Its first-order condition is

$$F(b) - F(\widehat{p}(b)) = \widehat{p}(b) f(\widehat{p}(b)).$$

We first show that

$$\lim_{b \downarrow 0} \frac{\widehat{p}(b)}{b} = \frac{1}{2}. \tag{10}$$

Let $b_k \downarrow 0$ and take a subsequence along which $\widehat{p}(b_k)/b_k$ converges to $q \in [0,1]$. Since $0 < \widehat{p}(b_k)/b_k < 1$, convergence of subsequences is guaranteed by compactness. The mean-value theorem implies the existence of $\xi_k \in (\widehat{p}(b_k), b_k)$ such that

$$F(b_k) - F(\widehat{p}(b_k)) = f(\xi_k)[b_k - \widehat{p}(b_k)].$$

Combining this identity with the first-order condition and dividing by b_k gives

$$\left(1 - \frac{\widehat{p}(b_k)}{b_k}\right) f(\xi_k) = \frac{\widehat{p}(b_k)}{b_k} f(\widehat{p}(b_k)).$$

Because $0 < \widehat{p}(b_k) < \xi_k < b_k \to 0$, right-continuity of f at zero implies that $f(\widehat{p}(b_k))$ and

$f(\xi_k)$ both converge to $f(0)$. Taking limits therefore gives $(1-q)f(0) = qf(0)$. Thus every convergent subsequence has limit $q = 1/2$, which proves (10).

Next, consider

$$\Phi_\delta(b) := \Big[(1-\delta)b + \delta\widehat{p}(b)\Big]\frac{f(b)}{1-F(b)}.$$

For every constant $c > 0$, we have

$$\frac{\Phi_\delta(c(1-\delta))}{1-\delta} \longrightarrow \frac{cf(0)}{2} \qquad \text{as } \delta \uparrow 1.$$

Choose $0 < c_- < 2/f(0) < c_+$. For all δ sufficiently close to one, we then have

$$\Phi_\delta(c_-(1-\delta)) < 1-\delta < \Phi_\delta(c_+(1-\delta)).$$

Continuity of Φ_δ yields a solution $b_1 \in (c_-(1-\delta), c_+(1-\delta))$. We can thus choose δ sufficiently close to one to ensure that $c_+(1-\delta) < \min\{p^1, \varepsilon\}$. Note that $b_1 = O(1-\delta)$. This proves the result. □

Choose b_1 as in Lemma 7 and let $b_0 = \widehat{p}(b_1)$. Note that $b_0 = O(1-\delta)$. We now define b_2 to b_N. Choose $\bar{b} \in (\max(b_1, 1-\varepsilon), 1)$. Since p, r, and F are continuous and $1 - F(a) \geq 1 - F(\bar{b}) > 0$ for $a \leq \bar{b}$, the function

$$(a,b) \longmapsto p(a,b)\frac{1-F(b)}{1-F(a)} - r(a)$$

is continuous on the compact set $[b_1, \bar{b}]^2$ and, hence, uniformly continuous. Uniform continuity lets us choose a single bound ρ on the gap between consecutive cutoffs that works for every $a \in [b_1, \bar{b}]$. On the diagonal, the value of the function is

$$p(a,a) - r(a) = (1-\delta)[a - r(a)] = (1-\delta)\delta(a - b_0) \geq (1-\delta)\delta(b_1 - b_0) > 0.$$

By uniform continuity, there exists $\rho > 0$ such that, whenever $b_1 \leq a \leq b \leq \bar{b}$ and

$b - a < \rho$, the absolute difference between the value at (a, b) and its value at (a, a) is less than $\frac{1}{2}(1-\delta)\delta(b_1 - b_0)$. Therefore,

$$p(a,b)\frac{1-F(b)}{1-F(a)} > r(a) \tag{11}$$

under the same conditions.

Choose an integer $N \geq 2$ sufficiently large that

$$d := \frac{\bar{b} - b_1}{N-1} < \min(\varepsilon, \rho),$$

and set

$$b_n = b_1 + (n-1)d, \qquad n = 1, \ldots, N.$$

Then, $b_N = \bar{b} > 1 - \varepsilon$, every successive gap equals $d < \varepsilon$, and (11) gives, for $n = 1, \ldots, N-1$,

$$p_{n+1}\frac{1-F(b_{n+1})}{1-F(b_n)} > r(b_n). \tag{12}$$

Write $I_0 = [0, b_1)$, $I_n = [b_n, b_{n+1})$ for $1 \leq n < N$, and $I_N = [b_N, 1]$. For an interval I, write F_I for F conditional on I.

Implication of the choice of b_1. The function r is increasing and concave, hence log-concave on $(0, 1]$. Since $1 - F$ is log-concave, the function $v \mapsto r(v)[1 - F(v)]$ is log-concave. For $v > b_0$, its derivative is

$$(1-\delta)[1 - F(v)] - r(v)f(v).$$

Lemma 7 says that this derivative is zero at b_1. Therefore, $v \mapsto r(v)[1 - F(v)]$ is maximized at b_1 and is decreasing on $[b_1, 1]$. We repeatedly use this observation below.

Information design. The messages that can be sent depend on the period: $m_{p,n}$ can be sent only in period n, for $n = 1, \ldots, N$. The message $\varnothing$ can be sent from period two onward, and the message $m_{r,n+1}$ from period $n + 1$ onward, for $n = 2, \ldots, N$. We also use messages $m_{b,n}$,

for $n = 1, \ldots, N$. If $n < N$, $m_{b,n}$ can be sent from period $n+3$ onward; $m_{b,N}$ can be sent from period $N+1$ onward. Under the sellers' strategies, $m_{p,n}$ induces price p_n, and $m_{b,n}$ induces price b_n. The message $m_{r,n+1}$ induces price $r(b_{n-1})$, while $\varnothing$ induces price b_0. No other messages can be sent.

For $n = 2, \ldots, N$, note that

$$p_n - r(b_n) = \delta(1-\delta)(b_{n-1} - b_0) > 0.$$

For $p \in [r(b_n), p_n]$, define

$$\lambda_n(p) := \frac{p - r(b_n)}{p_n - r(b_n)} \in [0,1]. \tag{13}$$

The following rules specify the information design at every history.

1. At the first period (empty history), send $m_{p,1}$ with probability one.
2. Consider a consumer who entered at period n when the message is $m_{p,n}$, where $2 \leq n \leq N$. At every history at which he remains active, if he rejects an offer $p \in [r(b_n), p_n]$, send $m_{r,n+1}$ next period with probability $\lambda_n(p)$ and $\varnothing$ with probability $1 - \lambda_n(p)$. If he rejects any other offer, send $\varnothing$ with probability one.
3. After any rejection by the first consumer, i.e., the one entering when the message is $m_{p,1}$, or by a consumer who entered at $\varnothing$, send $\varnothing$ next period with probability one. This rule depends on the consumer's entry message, not the current message. A consumer who entered at $m_{p,n}$, with $n \geq 2$, remains subject to Rule 2 even when the current message is $\varnothing$.
4. After either decision by a consumer who entered at $m_{b,n}$, send $m_{b,n}$ again.
5. If the current message is $m_{p,n}$ and p_n is accepted, send $m_{p,n+1}$ when $n < N$, and send $m_{b,N}$ when $n = N$.
6. If the current message is $m_{r,n+1}$ and $r(b_{n-1})$ is accepted, send $m_{b,n-1}$.
7. After every other acceptance, send $\varnothing$.

Sellers. A seller who observes $\varnothing$ offers b_0. A seller who observes $m_{p,n}$ offers p_n. A seller who observes $m_{r,n+1}$ offers $r(b_{n-1})$. A seller who observes $m_{b,n}$ offers b_n.

Consumers. A consumer who enters at $m_{p,n}$ accepts p if and only if $p \leq R_n(v)$, where

$$R_n(v) := \begin{cases} r(v), & v < b_n, \\ \max(r(v), p_n), & v \geq b_n. \end{cases} \tag{14}$$

Since $p_1 = r(b_1)$, we have $R_1 = r$. For $n \geq 2$, every price in $[r(b_n), p_n]$ is accepted precisely when $v \geq b_n$. At every price outside that interval, acceptance is equivalent to $p \leq r(v)$.

A consumer who enters at $\varnothing$ accepts if and only if $p \leq r(v)$.

A consumer who enters at $m_{b,n}$ accepts if and only if $p \leq S_n(v)$, where

$$S_n(v) := \begin{cases} v, & 0 \leq v < b_n, \\ (1-\delta)v + \delta b_n, & b_n \leq v \leq 1. \end{cases} \tag{15}$$

Thus, the strategy a consumer follows depends only on his valuation, the current price and the message at which he entered the market. As we will see, two different messages at which a consumer enters the market correspond to two different public beliefs. The public belief is $F_{[b_{n-1},1]}$ when the message is $m_{p,n}$ and $n > 1$, $F_{[b_n,b_{n+1})}$ when the message is $m_{b,n}$ (no consumers enter when the message is $m_{r,n+1}$) and $F_{[0,b_1)}$ when the message is $\varnothing$.[23] Clearly, all acceptance rules are monotone in prices. Thus, the strategies satisfy all the postulated properties.

Optimality and Beliefs.

We first repeat the description of the equilibrium path. If $v \in I_0$, the initial consumer rejects p_1 and the next message is $\varnothing$. Every later offer is b_0, which is accepted if and only if $v \geq b_0$. If $v \in I_n$, for $1 \leq n < N$, the offers $p_1, \ldots, p_n$ are accepted and p_{n+1} is rejected. Since $\lambda_{n+1}(p_{n+1}) = 1$, the next message is $m_{r,n+2}$. The same consumer purchases at $r(b_n)$. His

[23] We let $b_{N+1} = 1$.

replacement then receives $m_{b,n}$, and every subsequent offer is b_n and is accepted. If $v \in I_N$, all N offers are accepted and $m_{b,N}$ is sent forever, with price b_N.

Under the prescribed strategies, $\varnothing$ occurs with positive probability only after p_1 has been rejected in period 1. Following that rejection, the platform sends $\varnothing$ in every subsequent period. It also sends $\varnothing$ at other histories, which have zero equilibrium probability. Observing $\varnothing$ therefore induces the public belief $F_{[0,b_1)}$. Similarly, when $n > 1$, the message $m_{p,n}$ has positive probability only when the prices $(p_1, \ldots, p_{n-1})$ have been offered and all have been accepted. Thus, upon observing $m_{p,n}$, the public belief is $F_{[b_{n-1},1)}$. The message $m_{r,n}$ has positive probability only when the prices $(p_1, \ldots, p_{n-1})$ have been offered, the first $n-2$ prices have been accepted and p_{n-1} is rejected. Thus, upon observing $m_{r,n}$, the public belief is $F_{[b_{n-2},b_{n-1})}$. Finally, when $n < N$, the message $m_{b,n}$ has positive probability only when the prices $(p_1, \ldots, p_n, p_{n+1}, r(b_n))$ have been offered, the first n prices have been accepted, the price p_{n+1} is rejected and $r(b_n)$ is accepted. Thus, upon observing $m_{b,n}$, the public belief is $F_{[b_n,b_{n+1})}$. Upon observing $m_{b,N}$, the posterior is $F_{[b_N,1]}$.

We now verify the optimality of the sellers' strategies. Upon observing the message $\varnothing$ at period t, a seller assigns probability one to consumers using the acceptance rule r – see Rule 3. Clearly, if the seller offers a price $p \geq r(1)$, its expected demand is zero. For any remaining price, let $\widehat{v}$ satisfy $p = r(\widehat{v})$. If $\widehat{v} \leq b_0$, the seller's expected profit, up to the positive multiplicative constant $F(b_1) - F(0)$, is

$$\widehat{v}[F(b_1) - F(\widehat{v})],$$

which is maximized at b_0 by the definition of b_0. If $b_0 \leq \widehat{v} \leq b_1$, the expected profit is

$$r(\widehat{v})[F(b_1) - F(\widehat{v})] \leq \widehat{v}[F(b_1) - F(\widehat{v})] \leq b_0[F(b_1) - F(b_0)].$$

If $\widehat{v} > b_1$, the expected demand is zero – since beliefs are $F_{[0,b_1)}$. Thus, b_0 is optimal at $\varnothing$.

At $m_{p,1}$, the consumer uses $R_1 = r$ and the seller has prior F. Since $p_1 = r(b_1)$ and b_1

maximizes $r(v)[1-F(v)]$, the initial price p_1 is optimal.

At $m_{p,n}$, for $n=2,\ldots,N$, the posterior is $F_{[b_{n-1},1]}$ and the consumer uses R_n. The price p_n gives expected profit

$$p_n \frac{1-F(b_n)}{1-F(b_{n-1})} > r(b_{n-1})$$

by (12). Every price in $[r(b_n), p_n]$ sells to exactly the same valuations as p_n, at a weakly lower price. At every price outside this interval, the demand is governed by r. Again, a price $p > r(1)$ has zero demand. For any remaining price, let $\widehat{v}$ satisfy $p = r(\widehat{v})$. If $\widehat{v} \leq b_{n-1}$, revenue is at most $r(b_{n-1})$. If $\widehat{v} > b_{n-1}$, revenue is

$$r(\widehat{v}) \frac{1-F(\widehat{v})}{1-F(b_{n-1})} \leq r(b_{n-1}),$$

because $r(v)[1-F(v)]$ is decreasing on $[b_1, 1]$. Thus, p_n is optimal.

At $m_{r,n+1}$, the posterior is $F_{[b_{n-1},b_n)}$ and the active consumer uses R_n. Since the seller assigns probability one to $[v < b_n]$, the expected demand is induced by the acceptance rule r. Let $\widehat{v} \in [b_{n-1}, b_n)$ satisfy $p = r(\widehat{v})$. The expected profit, up to the positive multiplicative constant $F(b_n) - F(b_{n-1})$, is:

$$r(\widehat{v})[F(b_n) - F(\widehat{v})].$$

Its derivative satisfies

$$\begin{aligned}(1-\delta)[F(b_n)-F(\widehat{v})] - r(\widehat{v})f(\widehat{v}) &= \underbrace{\left[(1-\delta)[1-F(\widehat{v})] - r(\widehat{v})f(\widehat{v})\right]}_{\leq 0} - (1-\delta)[1-F(b_n)] \\ &\leq -(1-\delta)[1-F(b_n)] < 0,\end{aligned}$$

where the first term is nonpositive because $\widehat{v} \geq b_1$.

Charging less than $r(b_{n-1})$ sells to all consumers, but at a lower price, and pricing above $r(b_n)$ has zero expected demand. Hence, it is optimal to offer $r(b_{n-1})$.

At $m_{b,n}$, the consumer uses S_n, and the posterior is $F_{[b_n,b_{n+1})}$. (We define $b_{N+1} = 1$.) Let

$\widehat{v} \in [b_n, b_{n+1})$ satisfy $p = S_n(\widehat{v})$. The expected profit, up to a positive multiplicative constant, is

$$S_n(\widehat{v})[F(b_{n+1}) - F(\widehat{v})].$$

Since $S_n(\widehat{v}) = r(\widehat{v}) + \delta(b_n - b_0)$ on this range, its derivative is

$$\begin{aligned}&(1-\delta)[F(b_{n+1}) - F(\widehat{v})] - S_n(\widehat{v})f(\widehat{v}) \\ &\quad = \underbrace{\Big[(1-\delta)[1 - F(\widehat{v})] - r(\widehat{v})f(\widehat{v})\Big]}_{\leq 0} - (1-\delta)[1 - F(b_{n+1})] - \delta(b_n - b_0)f(\widehat{v}) < 0.\end{aligned}$$

The expression in brackets is nonpositive because $\widehat{v} \geq b_n \geq b_1$. Prices at or above $S_n(b_{n+1})$ have zero demand, apart from a possible endpoint of zero probability. Therefore, $S_n(b_n) = b_n$ is optimal.

It remains to verify the optimality of the consumers' strategies. A consumer who entered at $\varnothing$ receives b_0 after every rejection. For $v \geq b_0$, buying rather than rejecting changes his payoff by

$$(v - p) - \delta(v - b_0) = r(v) - p.$$

For $v < b_0$, rejection gives zero and the difference is again $v - p = r(v) - p$. The same argument applies to the initial consumer, whose rule is $R_1 = r$.

Consider a consumer who entered at $m_{p,n}$, for $2 \leq n \leq N$. At a price outside $[r(b_n), p_n]$, rejection leads next period to b_0; thus, using the acceptance rule r is optimal. Now consider $p \in [r(b_n), p_n]$. The definition of $\lambda_n(p)$ gives

$$p = \lambda_n(p)p_n + (1 - \lambda_n(p))r(b_n) \tag{16}$$

$$= (1-\delta)b_n + \delta\Big[\lambda_n(p)r(b_{n-1}) + (1 - \lambda_n(p))b_0\Big]. \tag{17}$$

For $v \geq b_{n-1}$, if the consumer rejects p, he purchases at either continuation price, i.e., b_0 or $r(b_{n-1})$. To see this, note that if the consumer rejects $r(b_{n-1})$ after having rejected p, the

offer is b_0 next period (since $r(b_{n-1}) < r(b_n)$), and

$$[v - r(b_{n-1})] - \delta(v - b_0) = (1 - \delta)(v - b_{n-1}) \geq 0.$$

Thus, the consumer has no incentive to reject $r(b_{n-1})$. Similarly, he has no incentive to reject b_0. Thus, buying at p rather than rejecting changes his payoff by

$$v - p - \delta\Big[\lambda_n(p)(v - r(b_{n-1})) + (1 - \lambda_n(p))(v - b_0)\Big] = (1 - \delta)(v - b_n).$$

It is optimal to accept precisely when $v \geq b_n$ among valuations in $[b_{n-1}, 1]$.

If $b_0 \leq v < b_{n-1}$, the consumer rejects $r(b_{n-1})$ after having rejected p and then purchases at b_0. The difference between accepting now and following this continuation is

$$\begin{aligned} v - p - \delta^2\lambda_n(p)(v - b_0) - \delta(1 - \lambda_n(p))(v - b_0) &= v - p - \delta[1 - \lambda_n(p)(1 - \delta)](v - b_0) \\ &= (1 - \delta)\Big[v - b_n + \delta\lambda_n(p)(v - b_{n-1})\Big] \\ &< 0. \end{aligned}$$

If $v < b_0$, then $p > b_0 > v$, and every continuation offer before purchase is at least b_0, so rejection is optimal. This verifies R_n after every history, including histories at which the consumer faces $\varnothing$.

Finally, a consumer who entered at $m_{b,n}$ receives b_n after every rejection. If $v < b_n$, waiting gives zero, so he accepts precisely when $p \leq v$. If $v \geq b_n$, he purchases next period at b_n, and the acceptance condition is

$$v - p \geq \delta(v - b_n) \quad \Longleftrightarrow \quad p \leq (1 - \delta)v + \delta b_n.$$

This is exactly S_n, including for valuations outside the seller's posterior support.

Finally, if $v \in I_n$ for $1 \leq n < N$, the platform identifies that interval after the rejection in

period $n+1$. If $v \in I_0$, it identifies I_0 after the first rejection; the response to b_0 in period two further distinguishes $[0, b_0)$ from $[b_0, b_1)$. If $v \in I_N$, it identifies I_N after N acceptances. All these intervals have length less than ε, so learning is complete after N periods. Every rejection lottery on the equilibrium path is degenerate, since $\lambda_n(p_n) = 1$. □

C. Consumer and producer surplus

C.1. Short-lived consumer

Proof of Theorem 3(i). We construct the three vertices of the attainable set and then combine them through public randomization.

The producer-optimal vertex $(\mathbb{E}[v] - \mathrm{CS}(\hat{p}), \mathrm{CS}(\hat{p}))$. Fix $\varepsilon > 0$. After a finite time $T(\varepsilon)$ the platform identifies either an interval $I = [a, b) \subseteq [0, \hat{p})$ with $b - a \leq \varepsilon$, or $[\hat{p}, 1]$. Thereafter, let the platform reveal the identified interval to every seller. On I, charging a guarantees revenue $a \geq \mathbb{E}[v \mid v \in I] - \varepsilon$. Optimal expected revenue is therefore at least this amount, and conditional consumer surplus is at most ε. On $[\hat{p}, 1]$, the monopoly price remains $\hat{p}$. Let $\mathrm{PS}_t(\mathcal{M})$ and $\mathrm{CS}_t(\mathcal{M})$ denote expected producer and consumer surplus in period t, respectively. Averaging across intervals, for every $t > T(\varepsilon)$,

$$\mathbb{E}[v] - \mathrm{CS}(\hat{p}) - \varepsilon \leq \mathrm{PS}_t(\mathcal{M}) \leq \mathbb{E}[v] - \mathrm{CS}(\hat{p}),$$
$$\mathrm{CS}(\hat{p}) \leq \mathrm{CS}_t(\mathcal{M}) \leq \mathrm{CS}(\hat{p}) + \varepsilon.$$

For fixed ε, the learning phase has discounted weight $1 - \delta^{T(\varepsilon)} \to 0$ as $\delta \uparrow 1$. Letting ε decrease sufficiently slowly as $\delta \uparrow 1$ gives the claimed vertex.

The consumer-optimal vertex $(\mathrm{PS}(\hat{p}), \mathbb{E}[v] - \mathrm{PS}(\hat{p}))$. Fix $\varepsilon > 0$. We first choose small intervals near zero and $\hat{p}$ whose combined contribution to total surplus is less than ε. Choose

$0 < 2\eta < b < \hat{p}$, with b sufficiently close to $\hat{p}$, such that

$$\int_0^{2\eta} v\, dF(v) + \int_b^{\hat{p}} v\, dF(v) < \varepsilon.$$

For $\gamma \in (0, \eta)$, use the construction in the proof of Proposition 2 in Appendix A.1, choosing a mesh strictly smaller than γ and including b as a gridpoint. Let a be the smallest gridpoint weakly above η. Then we obtain a partition:

$$[0, a), \quad K_i = [\ell_i, u_i)\ (i = 1, \ldots, N), \quad [b, \hat{p}), \quad [\hat{p}, 1],$$

where $a \in [\eta, \eta + \gamma)$, the K_i partition $[a, b)$, and $\max_i(u_i - \ell_i) \leq \gamma$. Without loss of generality, index the intervals so that $\ell_1 < \cdots < \ell_N$. This construction ensures that the platform identifies the cell containing v^* in finite time.

Uniform-profit segmentation of Bergemann et al. (2015) in an auxiliary economy.

We construct a finite auxiliary market by placing the probability mass of each cell K_i at its lower endpoint ℓ_i, and the mass of $[\hat{p}, 1]$ at $\hat{p}$.[24] The corresponding masses are

$$\alpha_i := F(u_i) - F(\ell_i), \qquad h := 1 - F(\hat{p}).$$

Demand in the auxiliary market is weakly lower at every price and unchanged at $\hat{p}$. Hence, $\hat{p}$ remains a monopoly price, yielding revenue $\mathrm{PS}(\hat{p})$.

We use the direct uniform-profit segmentation of Bergemann et al. (2015). Define

$$B_i := h + \sum_{k>i} \alpha_k, \qquad \theta_i := \frac{\ell_i \alpha_i}{\mathrm{PS}(\hat{p}) - \ell_i B_i}, \qquad s_1 := 1, \qquad s_{i+1} := s_i(1 - \theta_i).$$

Because $\hat{p}$ is the unique monopoly price under the prior and $\ell_i < b < \hat{p}$, we have $\mathrm{PS}(\hat{p}) > \ell_i(\alpha_i + B_i)$, so $\theta_i \in (0, 1)$ and $s_i > 0$. For auxiliary valuation ℓ_k, send message m_i with

[24] The auxiliary market retains the original weights on $[a, b) \cup [\hat{p}, 1]$, so its total mass is less than one. Normalizing these weights leaves optimal prices and the segmentation unchanged.

probability $s_i\theta_i$ for each $i < k$ and m_k with probability s_k. For auxiliary valuation $\hat{p}$, send m_i with probability $s_i\theta_i$ for every i and m_H with probability s_{N+1}. These probabilities sum to one because

$$\sum_{i<k} s_i\theta_i + s_k = 1, \qquad \sum_{i=1}^{N} s_i\theta_i + s_{N+1} = 1.$$

In the auxiliary economy, the seller is indifferent between ℓ_i and $\hat{p}$ after m_i; select ℓ_i. After m_H, select the optimal price $\hat{p}$. These prescribed prices achieve efficient trade and aggregate producer surplus $\mathrm{PS}(\hat{p})$. We next verify that they remain optimal in the original economy.

Implementation in the original economy. After learning, apply the same message probabilities to each $v \in K_k$ as to auxiliary valuation ℓ_k, and to each $v \in H := [\hat{p}, 1]$ as to $\hat{p}$. Reveal each omitted region through a separate message, inducing its conditional monopoly price.

After message m_i, the seller's unnormalized posterior assigns the prior measure to K_i, θ_i times the prior measure to $\bigcup_{k>i} K_k \cup H$, and zero measure elsewhere. We recommend price ℓ_i, which yields

$$\ell_i(\alpha_i + \theta_i B_i) = \theta_i \mathrm{PS}(\hat{p}).$$

A lower price serves the same consumers at a lower margin. At any price $p \geq u_i$, the K_i component generates no demand, so profit is at most

$$\theta_i p[1 - F(p)] \leq \theta_i \mathrm{PS}(\hat{p}).$$

It remains to consider $p \in (\ell_i, u_i)$. On this interval, profit is

$$\pi_i(p) = p\{F(u_i) - F(p) + \theta_i B_i\}.$$

Since $[\eta, b]$ is bounded away from the unique monopoly price,

$$\kappa := \min_{p \in [\eta, b]} \{\mathrm{PS}(\hat{p}) - p[1 - F(p)]\} > 0.$$

Moreover,

$$\mathrm{PS}(\hat{p}) - \ell_i B_i \geq \mathrm{PS}(\hat{p}) - \ell_i[1 - F(\ell_i)] \geq \kappa, \qquad \theta_i B_i \leq \frac{b}{\kappa}\alpha_i.$$

Because $\alpha_i = O(\gamma)$ uniformly and $\min_{p\in[\eta,b]} pf(p) > 0$, for every $p \in (\ell_i, u_i)$,

$$\pi_i'(p) \leq \alpha_i + \theta_i B_i - \min_{x\in[\eta,b]} xf(x) < 0$$

when γ is sufficiently small. Thus, ℓ_i is optimal after m_i. After m_H, the posterior is $F_{[\hat{p},1]}$, under which $\hat{p}$ is optimal.

Every valuation in $\bigcup_i K_i \cup H$ trades, and seller profit from these messages is

$$\mathrm{PS}(\hat{p}) \left(\sum_{i=1}^{N} s_i \theta_i + s_{N+1} \right) = \mathrm{PS}(\hat{p}).$$

The omitted regions contribute nonnegative surplus totaling at most

$$\int_0^a v \, dF(v) + \int_b^{\hat{p}} v \, dF(v) < \varepsilon.$$

Therefore, in every period after learning,

$$\mathrm{PS}(\hat{p}) \leq \mathrm{PS}_t(\mathcal{M}) \leq \mathrm{PS}(\hat{p}) + \varepsilon,$$
$$\mathbb{E}[v] - \mathrm{PS}(\hat{p}) - \varepsilon \leq \mathrm{CS}_t(\mathcal{M}) \leq \mathbb{E}[v] - \mathrm{PS}(\hat{p}).$$

For fixed ε, learning ends in finite time, so its discounted weight vanishes as $\delta \uparrow 1$. Letting ε decrease sufficiently slowly as $\delta \uparrow 1$ therefore yields the consumer-optimal vertex $(\mathrm{PS}(\hat{p}), \mathbb{E}[v] - \mathrm{PS}(\hat{p}))$.

The static vertex $(\mathrm{PS}(\hat{p}), \mathrm{CS}(\hat{p}))$. If the platform conceals the entire history, every seller charges $\hat{p}$. The surplus pair is $(\mathrm{PS}(\hat{p}), \mathrm{CS}(\hat{p}))$ in every period and hence for every δ.

The attainable set. We have shown that the producer-optimal, consumer-optimal, and static vertices are asymptotically attainable in the sense of Definition 3. Public randomization

before play, independently of v, across the three associated information designs therefore attains every point in the triangle $\mathcal{A}_{\mathrm{SL}}$ stated in Theorem 3(i).

Conversely, a seller can ignore the platform's message and charge $\hat{p}$, so expected producer surplus in every period is at least $\mathrm{PS}(\hat{p})$. By the price bound underlying Theorem 1(i), no design induces a price above $\hat{p}$. A short-lived consumer's surplus is therefore at least $\max\{0, v - \hat{p}\}$, yielding expected consumer surplus at least $\mathrm{CS}(\hat{p})$. Total surplus cannot exceed $\mathbb{E}[v]$. Hence, for every design and induced equilibrium,

$$\mathrm{PS}_\delta(\mathcal{M}) \geq \mathrm{PS}(\hat{p}), \qquad \mathrm{CS}_\delta(\mathcal{M}) \geq \mathrm{CS}(\hat{p}), \qquad \mathrm{PS}_\delta(\mathcal{M}) + \mathrm{CS}_\delta(\mathcal{M}) \leq \mathbb{E}[v].$$

These inequalities are preserved by public randomization and by limits as $\delta \uparrow 1$, excluding every outcome outside the triangle. This proves part (i). □

C.2. Long-lived consumer

Proof of Theorem 3(ii). Set $\delta_B = \delta$. We establish the three vertices of the attainable set and then combine them through public randomization.

The producer-optimal vertex $(\mathbb{E}[v], 0)$. Fix $\varepsilon \in (0, 1)$ and use the construction of Theorem 2. We show that the time $T(\varepsilon, \delta)$ after which every seller earns at least $v - \varepsilon$ can be chosen to grow at most logarithmically in $1/(1-\delta)$. We establish this bound using a geometric grid: each step increases the current cutoff by a fixed proportion, except for the final step.

Lemma 7 and (10) imply $b_1 \sim 2(1-\delta)/f(0)$ and $b_0/b_1 \to 1/2$. For $b_1 \leq a \leq b \leq 1 - \varepsilon/2$, the margin in (11), evaluated at $b = a$, is

$$p(a, a) - r(a) = (1-\delta)\delta(a - b_0) \geq (1-\delta)a/8$$

for δ sufficiently close to one. Before setting $b = a$, the margin in the grid inequality has a derivative with respect to b that is uniformly $O_\varepsilon(1-\delta)$, because $p(a, b) = O(1-\delta)$, f is bounded, and $1 - F(a)$ is bounded away from zero. Thus the inequality remains strict

whenever $b - a \leq \eta a$, for some sufficiently small $\eta \in (0, \varepsilon/2)$ independent of δ. We may therefore use the grid

$$b_{t+1} = \min\{(1+\eta)b_t, 1 - \varepsilon/2\},$$

stopping at $b_N = 1 - \varepsilon/2$. All gaps are smaller than ε, and the equilibrium verification is unchanged. This grid has $N = O_\varepsilon(\log(1/(1-\delta)))$ points. Allowing two further periods for the permanent prices to take effect gives

$$T(\varepsilon, \delta) = N + 2 = O_\varepsilon\left(\log \frac{1}{1-\delta}\right).$$

From period $T(\varepsilon, \delta)$ onward, seller revenue is at least $v - \varepsilon$ for every valuation, including $v < b_0 < \varepsilon$, where revenue is zero. Revenue before that time is nonnegative, so

$$\mathrm{PS}_\delta(\mathcal{M}) \geq \delta^{T(\varepsilon,\delta)-1}(\mathbb{E}[v] - \varepsilon).$$

For fixed ε, the logarithmic bound implies $(1-\delta)T(\varepsilon, \delta) \to 0$, so the right-hand side converges to $\mathbb{E}[v] - \varepsilon$. Letting ε decrease sufficiently slowly as $\delta \uparrow 1$ yields producer surplus $\mathbb{E}[v]$ in the limit. Nonnegativity and the total-surplus bound then imply that consumer surplus converges to zero, establishing the vertex $(\mathbb{E}[v], 0)$.

The consumer-optimal vertex $(0, \mathbb{E}[v])$. Let $\mathcal{M}^{\mathrm{ND}}$ conceal the entire history by sending the same message after every history. We construct an equilibrium with a constant price: following a price increase, consumers can wait for the prescribed price next period. Define the generalized failure rate $g(p) := pf(p)/[1 - F(p)]$. Log-concavity implies that the hazard rate $f(p)/[1 - F(p)]$ is nondecreasing, so g is strictly increasing. Moreover, $g(0) = 0$, and the first-order condition for the unique static monopoly price gives $g(\hat{p}) = 1$. For each $\delta \in (0, 1)$, therefore, there is a unique $p_\delta^{\mathrm{ND}} \in (0, \hat{p})$ satisfying $g(p_\delta^{\mathrm{ND}}) = 1 - \delta$, and $p_\delta^{\mathrm{ND}} \downarrow 0$ as $\delta \uparrow 1$.

Proposition 4 (No disclosure)**.** *Consider the information design $\mathcal{M}^{\mathrm{ND}}$. For every $\delta \in (0, 1)$, there is a stationary equilibrium with constant on-path price p_δ^{ND}.*

Proof. If a seller charges $p \le p_\delta^{\mathrm{ND}}$, a consumer accepts if and only if $v^* \ge p$. If instead $p > p_\delta^{\mathrm{ND}}$, the consumer accepts if and only if

$$v^* \ge \tau_\delta(p) := \frac{p - \delta p_\delta^{\mathrm{ND}}}{1-\delta}.$$

This strategy is optimal when all future sellers charge p_δ^{ND}: the cutoff $\tau_\delta(p)$ equates the payoff from buying now with that from waiting one period, while a price weakly below the future price is accepted whenever it does not exceed the valuation.

Consider a seller's deviation. For $p \le p_\delta^{\mathrm{ND}}$, expected profit is $p[1-F(p)]$, which is increasing because $g(p) \le g(p_\delta^{\mathrm{ND}}) = 1-\delta < 1$. For $p > p_\delta^{\mathrm{ND}}$, write $x = \tau_\delta(p)$, so that $p = (1-\delta)x + \delta p_\delta^{\mathrm{ND}}$. Expected profit is

$$\Pi_\delta(x) = \left[(1-\delta)x + \delta p_\delta^{\mathrm{ND}}\right][1-F(x)],$$

with derivative

$$\frac{\Pi'_\delta(x)}{1-F(x)} = (1-\delta) - \left[(1-\delta)x + \delta p_\delta^{\mathrm{ND}}\right]\frac{f(x)}{1-F(x)}.$$

The second term is nondecreasing in x, and at $x = p_\delta^{\mathrm{ND}}$ it equals $g(p_\delta^{\mathrm{ND}}) = 1-\delta$. Hence Π_δ is nonincreasing for $x \ge p_\delta^{\mathrm{ND}}$; prices for which $\tau_\delta(p) > 1$ yield zero demand. Thus, charging p_δ^{ND} is optimal, establishing the stationary equilibrium. □

Expected consumer surplus is constant across periods in this equilibrium. Normalized discounted consumer surplus therefore satisfies

$$\mathrm{CS}_\delta(\mathcal{M}^{\mathrm{ND}}) = \int_{p_\delta^{\mathrm{ND}}}^1 (v - p_\delta^{\mathrm{ND}})\, dF(v) \longrightarrow \mathbb{E}[v] \qquad \text{as } \delta \uparrow 1.$$

Nonnegativity and the total-surplus bound then imply $\mathrm{PS}_\delta(\mathcal{M}^{\mathrm{ND}}) \to 0$. Thus this design and its selected equilibria attain $(0, \mathbb{E}[v])$.

The vertex $(0,0)$. The next proposition constructs an information design and a stationary equilibrium in which sellers charge a constant price along the equilibrium path. This price converges to one as $\delta \uparrow 1$, so trade vanishes in the limit.

Proposition 5 (Vanishing surplus)**.** *There exist $\bar{\delta} \in (0,1)$ and information designs $\mathcal{M}^0_\delta$, for $\delta \in (\bar{\delta}, 1)$, each admitting a stationary equilibrium in which sellers charge the same price p^*_δ in every period along the equilibrium path. Moreover, $p^*_\delta \to 1$ as $\delta \uparrow 1$.*

Proof. Construction of p^*_δ. For each δ, select a full-disclosure equilibrium from Proposition 1, including its continuation after arbitrary rejected offers. Write P_δ and Q_δ for the corresponding nondecreasing price and reservation-price functions in (27) of Online Appendix OB.2. Recall that Corollary 1 in Online Appendix OB.2 implies $P_\delta(1) \to 0$ as $\delta \uparrow 1$, and hence

$$Q_\delta(1) = 1 - \delta + \delta P_\delta(1) \longrightarrow 0.$$

For δ sufficiently close to one, define p^*_δ as the largest solution in $(0,1)$ of

$$p^*_\delta[1 - F(p^*_\delta)] = 2Q_\delta(1).$$

The static revenue function is continuous, positive on $(0,1)$, and zero at one. Since $Q_\delta(1) \to 0$, this solution exists for δ sufficiently close to one and satisfies $p^*_\delta \to 1$.

Information design. The platform sends a concealing message H initially and after every acceptance. Rejected offers may trigger a full-disclosure phase, which lasts until the current consumer purchases. The information design follows these rules:

(i) If an offer $p < Q_\delta(p^*_\delta)$ is rejected at H, initiate the selected full-disclosure continuation following rejection of that offer from prior F.

(ii) If an offer $p \in [Q_\delta(p^*_\delta), p^*_\delta]$ is rejected at H, initiate the full-disclosure continuation at posterior $F_{[0,p^*_\delta]}$ with probability

$$\theta(p) = \frac{p^*_\delta - p}{p^*_\delta - Q_\delta(p^*_\delta)} = \frac{p^*_\delta - p}{\delta[p^*_\delta - P_\delta(p^*_\delta)]},$$

and send H otherwise.

(iii) If an offer $p > p^*_\delta$ is rejected at H, send H.

(iv) Whenever disclosure is initiated under (i) or (ii), follow the selected full-disclosure continuation until the incumbent purchases, then send H to the next seller. Reveal the information defining the continuation's initial posterior and subsequent transactions, while keeping earlier transactions concealed.

These rules apply after every history. In particular, the first rule uses the complete disclosure continuation after the actual rejected offer, including any off-path randomization.

The lottery is well defined because monotonicity of Q_δ and the choice of p^*_δ give $Q_\delta(p^*_\delta) \leq Q_\delta(1) \leq p^*_\delta/2 < p^*_\delta$, so $\theta(p) \in [0,1]$ throughout the interval in rule (ii).

Seller strategies. A seller receiving H charges p^*_δ. After any message in a full-disclosure continuation, the seller uses the corresponding price strategy of the selected full-disclosure equilibrium.

Consumer strategies. Every consumer enters at H. At this message, a consumer with valuation v accepts an offer p if and only if $p \leq R(v)$, where

$$R(v) = \begin{cases} Q_\delta(v), & v \leq p^*_\delta, \\ (1-\delta)v + \delta p^*_\delta, & v > p^*_\delta. \end{cases}$$

During disclosure, the consumer follows the selected full-disclosure strategy on the posterior support and its sequentially optimal completion outside that support.

Optimality and beliefs. We first verify consumer incentives. At H, continuation play gives $U(v) = (v - p^*_\delta)_+$. Let $V(v)$ be the value before the first offer in the disclosure continuation at $F_{[0,p^*_\delta]}$; for $v \geq p^*_\delta$, it equals $v - P_\delta(p^*_\delta)$. For $p < Q_\delta(p^*_\delta)$, the inherited disclosure continuation makes the prescribed reservation rule optimal. For $p > p^*_\delta$, rejection returns to H, giving reservation price $(1-\delta)v + \delta p^*_\delta$ for $v > p^*_\delta$; valuations $v \leq p^*_\delta$ reject. For $p \in [Q_\delta(p^*_\delta), p^*_\delta]$, the gain from accepting rather than rejecting is

$$A(v,p) = v - p - \delta\{\theta(p)V(v) + [1-\theta(p)]U(v)\}.$$

For $v \leq p^*_\delta$, at $p = Q_\delta(p^*_\delta)$ rejection initiates disclosure with certainty. The consumer can buy at the next offer $P_\delta(p^*_\delta)$, so $V(v) \geq v - P_\delta(p^*_\delta)$. Using (27) in Online Appendix OB.2,

$$A\Big(v, Q_\delta(p^*_\delta)\Big) \leq (1-\delta)(v - p^*_\delta) \leq 0.$$

At $p = p^*_\delta$, the gain is $v - p^*_\delta \leq 0$. Since $A(v,p)$ is affine in p, consumers with these valuations reject throughout the interval, with ties resolved according to the reservation rule. For $v > p^*_\delta$, substitution gives $A(v,p) = (1-\delta)(v - p^*_\delta) > 0$. Thus the prescribed rule is optimal at every offer.

During disclosure, the supported valuations lie in $[0, p^*_\delta]$, where R agrees with the selected full-disclosure strategy. All consumers enter at H and use the same rule, with the specified sequentially optimal completion outside the posterior support. The construction therefore satisfies stationarity and monotonicity, including after a consumer's own deviation.

We next verify seller incentives. At H, an offer below $Q_\delta(p^*_\delta)$ yields revenue at most $Q_\delta(1)$, strictly less than $p^*_\delta[1 - F(p^*_\delta)] = 2Q_\delta(1)$. Offers in $[Q_\delta(p^*_\delta), p^*_\delta]$ serve the same upper tail and yield the highest revenue at p^*_δ. An offer $p > p^*_\delta$ induces cutoff $x = (p - \delta p^*_\delta)/(1-\delta) > p^*_\delta$. If $x \leq 1$, then $p < x$, and the choice of p^*_δ as the largest solution of the defining revenue equation gives

$$p[1 - F(x)] \leq x[1 - F(x)] \leq p^*_\delta[1 - F(p^*_\delta)].$$

If $x > 1$, revenue is zero. Thus p^*_δ is optimal. Sellers in disclosure continuations behave optimally by following the selected full-disclosure equilibrium.

At H, the seller's posterior is F, since every valuation generates H on path. Disclosure begins with the posterior induced by rejection of the offer; the lottery adds no information conditional on rejection. Within disclosure, including after zero-probability histories, use the belief system of the selected full-disclosure equilibrium.

Finally, $\theta(p^*_\delta) = 0$, so every on-path decision is followed by H. Thus every seller charges p^*_δ in every period along the equilibrium path, and $p^*_\delta \to 1$ as $\delta \uparrow 1$. □

In this equilibrium, trade occurs in every period when $v^* > p^*_\delta$ and never occurs otherwise. Expected flows are constant over time, so the resulting design $\mathcal{M}^0_\delta$ satisfies

$$\begin{aligned}
\mathrm{PS}_\delta(\mathcal{M}^0_\delta) &= p^*_\delta[1 - F(p^*_\delta)] \longrightarrow 0, \\
\mathrm{CS}_\delta(\mathcal{M}^0_\delta) &= \int_{p^*_\delta}^{1} (v - p^*_\delta)\, dF(v) \longrightarrow 0.
\end{aligned}$$

Thus $(0, 0)$ is asymptotically attainable.

The attainable set. Public randomization before play, independently of v, across the three information designs associated with the three vertices attains every point in $\mathcal{A}_{\mathrm{LL}}$. Conversely, nonnegativity and the total-surplus bound hold under every design and are preserved by public randomization and limits. Thus no outcome outside the triangle is attainable, proving part (ii). □

Online Appendix (Not for publication)

Section OA studies short-lived consumers: OA.1 analyzes learning under full disclosure, and OA.2 allows for competing sellers. Section OB studies long-lived consumers: OB.1 clarifies stationarity, OB.2 analyzes full disclosure, OB.3 establishes learning for arbitrarily patient consumers, and OB.4 examines belief consistency.

OA. Short-lived consumers

OA.1. Full disclosure with short-lived consumers

Under *full disclosure*, each new seller observes the current characteristic and the entire history of characteristics, prices, and purchase decisions, and maximizes current expected profit.

Accepted and rejected offers narrow the set of possible valuations. We characterize the limiting set and illustrate how the order of arrivals affects what the platform learns.

Recursive prices and learning. For a rectangle $R = \prod_{c\in\mathcal{C}}[\ell_c, u_c] \subseteq [0,1]^C$, with $\ell_c < u_c$ for every c, define

$$p_c(R) := \arg\max_{p\in[\ell_c,u_c]} p\,\mathbb{P}_F(v_c \geq p \mid \mathbf{v}\in R) = \arg\max_{p\in[\ell_c,u_c]} p\int_{R\cap\{\mathbf{v}:v_c\geq p\}} f(\mathbf{v})\,d\mathbf{v}. \tag{18}$$

This price is unique under the maintained assumptions. Fix a realized profile $\mathbf{v}^*$ and an arrival sequence $\mathbf{c} = (c_1, c_2, \ldots)$. Starting from $R_0(\mathbf{v}^*, \mathbf{c}) = [0,1]^C$, the period-$t$ price and the rectangle after the purchase decision are given recursively by

$$\begin{aligned} p_t(\mathbf{v}^*, \mathbf{c}) &= p_{c_t}\Big(R_{t-1}(\mathbf{v}^*, \mathbf{c})\Big), \\ R_t(\mathbf{v}^*, \mathbf{c}) &= R_{t-1}(\mathbf{v}^*, \mathbf{c}) \cap \begin{cases} \{\mathbf{v} : v_{c_t} \geq p_t(\mathbf{v}^*, \mathbf{c})\}, & \text{if } v^*_{c_t} \geq p_t(\mathbf{v}^*, \mathbf{c}), \\ \{\mathbf{v} : v_{c_t} < p_t(\mathbf{v}^*, \mathbf{c})\}, & \text{if } v^*_{c_t} < p_t(\mathbf{v}^*, \mathbf{c}). \end{cases} \end{aligned} \tag{19}$$

Both $p_t(\mathbf{v}^*, \mathbf{c})$ and $R_t(\mathbf{v}^*, \mathbf{c})$ depend only on the first t arrivals; $R_{t-1}(\mathbf{v}^*, \mathbf{c})$ depends only on the first $t-1$.

An acceptance raises the lower bound for the arriving characteristic; a rejection lowers its upper bound. The other coordinate bounds do not change, although their conditional distributions generally do. We include interval endpoints when writing $R_t(\mathbf{v}^*, \mathbf{c})$; because valuations have a density, this does not affect posterior beliefs. Write $R_t(\mathbf{v}^*, \mathbf{c}) = \prod_c [\ell_{c,t}(\mathbf{v}^*, \mathbf{c}), u_{c,t}(\mathbf{v}^*, \mathbf{c})]$ and define

$$R_\infty(\mathbf{v}^*, \mathbf{c}) := \bigcap_{t \geq 0} R_t(\mathbf{v}^*, \mathbf{c}) = \prod_c [\ell_{c,\infty}(\mathbf{v}^*, \mathbf{c}), u_{c,\infty}(\mathbf{v}^*, \mathbf{c})],$$

where

$$\ell_{c,\infty}(\mathbf{v}^*, \mathbf{c}) := \lim_t \ell_{c,t}(\mathbf{v}^*, \mathbf{c}), \qquad u_{c,\infty}(\mathbf{v}^*, \mathbf{c}) := \lim_t u_{c,t}(\mathbf{v}^*, \mathbf{c}).$$

These limits exist because the lower bounds increase, the upper bounds decrease, and all bounds lie in $[0, 1]$.

Proposition 6 (Learning under full disclosure)**.** *Fix $\mathbf{v}^* \in (0,1)^C$ and arrival probabilities $\lambda_c > 0$ for every c.*

(i) For almost every arrival sequence, full disclosure yields a limiting rectangle $R_\infty(\mathbf{v}^*, \mathbf{c})$ of strictly positive prior probability. The platform's posterior after period t is $F(\cdot \mid R_t(\mathbf{v}^*, \mathbf{c}))$ and converges to the prior conditional on the limiting rectangle $R_\infty(\mathbf{v}^*, \mathbf{c})$,

$$F(\cdot \mid R_\infty(\mathbf{v}^*, \mathbf{c})). \tag{20}$$

(ii) Conditional on $\mathbf{v}^*$, the learning outcome has a probability distribution over rectangles induced by the arrival process. For every measurable collection $\mathcal{A}$ of rectangles, this distribution is

$$\Pi_{\mathbf{v}^*,\lambda}(\mathcal{A}) = \mathbb{P}_\lambda(R_\infty(\mathbf{v}^*, \mathbf{c}) \in \mathcal{A}), \tag{21}$$

where $\mathbb{P}_\lambda$ is the iid law of the characteristic sequence and $R_\infty(\mathbf{v}^*, \mathbf{c})$ is determined by (18)–(19).

Proof of Proposition 6. Part (i). Every rectangle reached at a finite date has positive volume. Starting from $R_0(\mathbf{v}^*, \mathbf{c}) = [0,1]^C$, an acceptance leaves a positive-width interval: the chosen price is below the current upper bound, since a price at that bound yields zero profit whereas a lower price yields positive profit. A rejection also leaves a positive-width interval because the rejected price strictly exceeds v^*_{ct} and hence the current lower bound. Arrivals are independent of valuations, and each purchase decision adds exactly the inequality in (19). Successive Bayesian updating therefore gives $F(\cdot \mid R_t(\mathbf{v}^*, \mathbf{c}))$ at every finite date.

Positive probability of the limiting rectangle. We first show that only finitely many rejections occur. Each rejection reduces the relevant upper bound by a fixed proportion, so infinitely many rejections would contradict the strictly positive realized valuation.

The maintained density is bounded above and bounded away from zero on the cube. To see this, let $m > 0$ be its smallest value at a vertex. Log-concavity gives $f(\mathbf{v}) \geq m$ everywhere. With $z = (1/2, \ldots, 1/2)$, it also gives $f(z)^2 \geq f(\mathbf{v}) f(\mathbf{1} - \mathbf{v}) \geq m f(\mathbf{v})$. Hence $m \leq f \leq M := f(z)^2/m < \infty$.

For a rectangle $R = \prod_c [\ell_c, u_c]$, let $R_{-c} := \prod_{d \neq c} [\ell_d, u_d]$ denote its projection onto the other coordinates, and write

$$g_c(x) := \int_{R_{-c}} f(x, v_{-c})\, dv_{-c}, \qquad V_{-c} := \mathrm{vol}(R_{-c}).$$

Then $mV_{-c} \leq g_c(x) \leq MV_{-c}$ on $[\ell_c, u_c]$. If the optimal price $p = p_c(R)$ is rejected, then $p > v^*_c \geq \ell_c$, so it is an interior optimum. Its first-order condition, $p\, g_c(p) = \int_p^{u_c} g_c(x)\, dx$, implies $mp \leq M(u_c - p)$ and hence $p \leq q u_c$, where $q := M/(M+m) < 1$. Every rejection thus reduces that coordinate's upper bound by a factor of at most q. Infinitely many rejections would force its upper bound to zero, contradicting $v^*_c > 0$. There are therefore only finitely many rejections. Each final upper bound is strictly above v^*_c, while each limiting lower bound is at most v^*_c. Thus $R_\infty(\mathbf{v}^*, \mathbf{c})$ has positive volume and $F(R_\infty(\mathbf{v}^*, \mathbf{c})) > 0$.

Posterior convergence. Since $R_t(\mathbf{v}^*, \mathbf{c}) \downarrow R_\infty(\mathbf{v}^*, \mathbf{c})$, continuity of probability from above

implies $F(R_t(\mathbf{v}^*, \mathbf{c})) \downarrow F(R_\infty(\mathbf{v}^*, \mathbf{c})) > 0$. Consequently,

$$\begin{aligned}&\left\| F(\cdot \mid R_t(\mathbf{v}^*, \mathbf{c})) - F(\cdot \mid R_\infty(\mathbf{v}^*, \mathbf{c})) \right\|_{\mathrm{TV}} \\ &\quad = 1 - \frac{F(R_\infty(\mathbf{v}^*, \mathbf{c}))}{F(R_t(\mathbf{v}^*, \mathbf{c}))} \longrightarrow 0.\end{aligned}$$

Part (ii). Conditional on $\mathbf{v}^*$, the recursion determines $R_\infty(\mathbf{v}^*, \mathbf{c})$ from the arrival sequence. Its endpoints are limits of measurable functions of finite arrival histories and are therefore measurable. Its distribution under the iid arrival law $\mathbb{P}_\lambda$ is thus given by (21). □

For a fixed valuation profile, each realized arrival sequence therefore determines one limiting rectangle. Different arrival sequences can generate different rectangles because information about one characteristic changes the prices offered to other characteristics. Within the learned rectangle, the posterior retains the dependence in F; it is not generally a product of truncated marginal distributions.

The proposition describes limiting learning, not necessarily learning completed at a finite date. Only finitely many rejections occur for an interior valuation profile, but subsequent acceptances can continue to tighten lower bounds as information passes across characteristics.

Example: two characteristics below the learning boundary. We develop the example depicted in Figure 2 in Section 3.3, which compares full disclosure and information design. The prior is that of Example 1 in Section 3.1 with $K = 10$, namely

$$f(v_1, v_2) = Z^{-1} \exp[-10 \max\{v_1 - v_2, 0\}], \qquad (v_1, v_2) \in [0, 1]^2,$$

where Z is the normalizing constant. The recursion $q^0 = 0$ and $q^{k+1} = \max_{c=1,2} p_c([q^k, 1]^2)$ gives the information-design learning boundary $p^\infty = \lim_k q^k \simeq 0.6978$. Fix $\mathbf{v}^* = (0.40, 0.65)$ and $\lambda_1 = \lambda_2 = 1/2$. Both valuations lie below this boundary, so information design can learn both to arbitrary precision. Under full disclosure, however, the platform learns one of two rectangles, depending on which characteristic arrives first. All prices below are obtained from

(18); numerical values are rounded to four decimals.

Characteristic 1 arrives first. The initial offer to characteristic 1 is 0.3628, which is accepted. This acceptance makes sellers more optimistic about the valuation associated with characteristic 2. The next consumer with that characteristic is offered 0.6030, rather than the initial monopoly price of 0.5645, and accepts. Given these two acceptances, the next consumer with characteristic 1 is offered 0.4384 and rejects. The platform has therefore identified the rectangle

$$R_A = [0.3628, 0.4384) \times [0.6030, 1].$$

Characteristic 2 arrives first. The initial offer to characteristic 2 is 0.5645, which is accepted. The next offer to characteristic 1 is then 0.4298, which is rejected. The posterior support becomes $[0, 0.4298) \times [0.5645, 1]$. On this rectangle the density is constant, so the next seller facing characteristic 1 charges $0.4298/2 \simeq 0.2149$. This offer is accepted, leaving

$$R_B = [0.2149, 0.4298) \times [0.5645, 1].$$

Why these are the only two outcomes. Repeated arrivals between the informative observations above only delay learning. After an acceptance, the same characteristic faces the same price until news about the other characteristic changes beliefs. In the second case, any arrivals of characteristic 2 after the rejection of 0.4298 continue to face 0.5645 and accept. Thus the first arriving characteristic alone selects the eventual rectangle, and

$$\mathbb{P}_\lambda\Big(R_\infty(\mathbf{v}^*, \mathbf{c}) = R_A\Big) = \lambda_1 = \tfrac{1}{2},$$
$$\mathbb{P}_\lambda\Big(R_\infty(\mathbf{v}^*, \mathbf{c}) = R_B\Big) = \lambda_2 = \tfrac{1}{2}.$$

Equal probabilities here reflect this particular construction, not a general implication of uniform arrival rates. The rectangles overlap: they are alternative learning outcomes for the same profile, not cells of a common partition.

Both rectangles lie entirely in $v_1 < v_2$, where the prior density is constant. Their conditional distributions are therefore product-uniform. For a uniform valuation on $[\ell, u]$, the monopoly price is $\max\{\ell, u/2\}$. In each rectangle, this price equals the lower endpoint for both characteristics. Every subsequent offer is accepted and no further information is generated: both rectangles are absorbing.

Figure 2 in Section 3.3 compares these outcomes with information design. The first arrival changes the prices offered before a rejection, and hence the information retained when full-disclosure learning stops. Information design can instead ε-learn both coordinates for arbitrarily small $\varepsilon > 0$. The singleton in Panel (c) depicts this precision limit, not exact learning at a finite date under a fixed design.

OA.2. Profiling consumers with competing sellers

We extend the one-characteristic model with short-lived consumers to Hotelling competition with linear transportation costs.

In each period, two short-lived sellers enter the market: Seller A is located at zero and Seller B at one. The platform sends a message observed by both sellers, who then set prices simultaneously. The consumer purchases from A or B or makes no purchase. The consumer and both sellers then leave the market and are replaced in the next period.

The consumers' common location y^* is drawn from F on $[0, 1]$ before period 1. Their valuations are $v_A^* = w - y^*\tau$ for product A and $v_B^* = w - (1 - y^*)\tau$ for product B, where $\tau > 0$. We assume the baseline benefit w is sufficiently large to ensure full market coverage at the equilibrium prices constructed below (specifically, $w \geq \tau\big(1 + 2/\min_{[0,1]} f\big)$). In addition to Assumption 2 about F in the main text, we assume that f is symmetric around $1/2$.

Let p_A, p_B be prices charged by sellers A and B, respectively. We extend Assumption 1 from the main text to the competitive setup:

Assumption 5 (Short-lived consumers)**.** *A consumer at location y^* purchases from Seller A if and only if $v_A^* - p_A \geq \max\{0, v_B^* - p_B\}$ and from Seller B otherwise (as long as $v_B^* - p_B \geq 0$).*

We consider perfect Bayesian equilibria as defined in the body of the paper. For the result

below, we restrict sellers to pure pricing strategies. We use the same definition of ε-learning from the body of the paper, substituting "valuations" with "locations."

Let y_A^0 be the location of the indifferent consumer when Seller B charges zero and Seller A best responds under the prior F. Define y_B^0 symmetrically. We have that

$$F(y_A^0) - \left(\frac{1}{2} - y_A^0\right) f(y_A^0) = 0, \qquad 1 - F(y_B^0) + \left(\frac{1}{2} - y_B^0\right) f(y_B^0) = 0. \tag{22}$$

Log-concavity assures that these solutions are unique, and symmetry implies $y_B^0 = 1 - y_A^0$.

Proposition 7 (Competing sellers)**.**

(i) ***No learning below*** y_A^0 ***or above*** y_B^0***.*** *For every public information design and every induced equilibrium in pure pricing strategies, any two locations both in* $[0, y_A^0]$ *or both in* $(y_B^0, 1]$ *induce the same distribution over platform histories.*

(ii) ***Learning between*** y_A^0 ***and*** y_B^0***.*** *For every* $\varepsilon > 0$*, the platform can* ε*-learn all locations in* (y_A^0, y_B^0)*.*

The result mirrors the learning result with one seller. The platform cannot distinguish consumers who have a strong preference for one of the products, but it can learn detailed information about the location of consumers who do not. For a uniform prior, the two thresholds are $y_A^0 = 1/4$ and $y_B^0 = 3/4$. The point y_A^0 identifies the indifferent consumer when Seller A faces the highest possible competition: Seller B charges zero. Pooling histories associated with locations in $[0, y_A^0]$ with histories associated with higher locations cannot induce a lower cutoff. A symmetric argument applies to y_B^0.

Learning uses the same idea as in Theorem 1: concealing past purchases changes sellers' beliefs and induces prices that reveal more information. With competing sellers, it is the difference between their prices that determines the consumer's choice. Consider learning below $1/2$. Rather than revealing that the consumer's location is below $1/2$, the platform sends the same message after these histories and some histories corresponding to higher locations. This pooling can raise Seller B's price relative to Seller A's. Some locations that

previously led to purchases from B now lead to purchases from A, allowing the platform to distinguish them. The construction varies which histories receive the same message to generate arbitrarily fine distinctions between y_A^0 and y_B^0.

Proof. **Part 1: No learning below y_A^0 or above y_B^0.** Suppose that, before a period, all locations in $[0, y_A^0]$ induce the same distribution over platform histories. This is true before period 1. Conditional on any message that these locations can receive, the sellers' posterior density on $[0, y_A^0]$ is therefore proportional to f.

Fix Seller B's price $p_B \geq 0$. If Seller A's price induces a cutoff $\bar{y} \in (0, y_A^0)$, her profit, up to a positive normalization, is

$$p_A F(\bar{y}), \qquad \bar{y} = \min\left\{\frac{1}{2} + \frac{p_B - p_A}{2\tau}, \frac{w - p_A}{\tau}\right\}.$$

Using $p_B \geq 0$ and $w \geq \tau$, both one-sided derivatives of $p_A F(\bar{y})$ with respect to p_A are bounded above by

$$F(\bar{y}) - \frac{p_A}{2\tau} f(\bar{y}) \leq F(\bar{y}) - \left(\frac{1}{2} - \bar{y}\right) f(\bar{y}) < 0.$$

The last inequality follows from (22) and the fact that F/f is increasing. Reducing the price is therefore profitable. A price inducing a cutoff at or below zero is also not optimal, since Seller A can earn positive profit by serving some of these consumers. Thus every equilibrium price pair at such a message induces a cutoff at least y_A^0, and all locations in $[0, y_A^0]$ continue to generate the same purchase decision. Induction proves the claim. The argument for $(y_B^0, 1]$ is symmetric.

Part 2: Learning between y_A^0 and y_B^0

Information design. In period 1, the platform sends an uninformative message. The equilibrium prices are $p_A = p_B = \tau / f(1/2)$. A purchase from Seller A reveals $y^* \leq 1/2$, and a purchase from Seller B reveals $y^* > 1/2$.

Suppose past purchases have revealed whether $y^* \leq b$ on every on-path history, for some $b \in [y_A^0, y_B^0]$. For each $a \in \{0, 1/2, 1\}$ with $a \neq b$, use one period as follows. If $a < b$, send

message (a, b) after all histories consistent with $a < y^* \leq b$, including $y^* = 0$ when $a = 0$. If $a > b$, send it after all histories consistent with $b < y^* \leq a$. On the remaining histories, reveal whether $y^* \leq \min\{a, b\}$ or $y^* > \max\{a, b\}$. These messages are feasible because earlier purchases have revealed on which side of b and $1/2$ the location lies. The common message (a, b) conceals any finer distinctions that earlier purchases revealed within that interval.

After each message, sellers have the prior restricted to the interval specified by the message. They charge the prices derived in (23) below, with the endpoints in increasing order. Write $\bar{y}(a, b)$ for the indifferent location after message (a, b). We verify below that the prescribed prices constitute an equilibrium and that $\bar{y}(a, b)$ lies between a and b and belongs to $[y_A^0, y_B^0]$.

After message (a, b), a purchase from Seller A reveals $y^* \leq \bar{y}(a, b)$, and a purchase from Seller B reveals $y^* > \bar{y}(a, b)$. On histories receiving another message, earlier purchases already reveal on which side of $\bar{y}(a, b)$ the location lies. Thus, by the end of this period, the platform knows whether $y^* \leq \bar{y}(a, b)$ on every on-path history.

Starting with $b = 1/2$, apply this procedure to every cutoff generated in the preceding round. Use the same schedule on every history. Each round takes finitely many periods, and the platform retains all information from earlier purchases.

We next verify sellers' prices and then show that the cutoffs generated by these messages become arbitrarily close to every location in $[y_A^0, y_B^0]$.

Pricing with posterior $F_{[a,b]}$. Suppose both sellers have posterior $F_{[a,b]}$, with $a < b$. For prices p_A and p_B, let $\bar{y} = 1/2 + (p_B - p_A)/(2\tau)$ be the indifferent location. When $\bar{y} \in (a, b)$, the sellers' profits are proportional to $p_A[F(\bar{y}) - F(a)]$ and $p_B[F(b) - F(\bar{y})]$. Differentiating each seller's profit with respect to her own price gives the first-order conditions

$$p_A = 2\tau \frac{F(\bar{y}) - F(a)}{f(\bar{y})}, \qquad p_B = 2\tau \frac{F(b) - F(\bar{y})}{f(\bar{y})}. \tag{23}$$

Substituting these conditions back in $\bar{y} = 1/2 + (p_B - p_A)/(2\tau)$ gives

$$2F(\bar{y}) + \left(\bar{y} - \frac{1}{2}\right) f(\bar{y}) = F(a) + F(b). \tag{24}$$

The first-order conditions are sufficient for profit maximization. When $\bar{y} \in (a, b)$, the derivatives of the sellers' profits with respect to their own prices have the signs of

$$\frac{F(\bar{y}) - F(a)}{f(\bar{y})} - \frac{p_A}{2\tau}, \qquad \frac{F(b) - F(\bar{y})}{f(\bar{y})} - \frac{p_B}{2\tau}.$$

Log-concavity implies that $(F(\bar{y}) - F(a))/f(\bar{y})$ is increasing in $\bar{y}$ and $(F(b) - F(\bar{y}))/f(\bar{y})$ is decreasing in $\bar{y}$. Raising p_A lowers $\bar{y}$, so the first expression decreases with p_A. Raising p_B raises $\bar{y}$, so the second expression decreases with p_B. Each seller's profit therefore increases until her first-order condition holds and decreases thereafter.

Prices inducing $\bar{y} \notin [a, b]$ either yield no sales or can be raised without reducing demand, so they cannot improve on the proposed prices.

If a deviating price causes some consumers not to purchase from either seller, the seller earns no more than calculated above. At the proposed prices, all consumers purchase. Hence these deviations are also unprofitable.

It remains to verify that (24) has a solution between the endpoints of every interval used by the design. These intervals have endpoints $a \in \{0, 1/2, 1\}$ and $b \in [y_A^0, y_B^0]$. When $a > b$, the posterior is $F_{[b,a]}$, and (24) still applies.

Symmetry and log-concavity imply that f is non-decreasing below $1/2$. Moreover, since F/f is increasing, (22) implies that, for $\bar{y} \in [y_A^0, 1/2]$,

$$\left(\frac{1}{2} - \bar{y}\right) f(\bar{y}) \leq F(\bar{y}), \qquad F(\bar{y}) f'(\bar{y}) \leq f(\bar{y})^2.$$

Together with $f'(\bar{y}) \geq 0$, these inequalities give

$$\frac{d}{d\bar{y}}\left[2F(\bar{y}) + \left(\bar{y} - \frac{1}{2}\right) f(\bar{y})\right] = 3f(\bar{y}) + \left(\bar{y} - \frac{1}{2}\right) f'(\bar{y}) \geq 2f(\bar{y}). \tag{25}$$

Symmetry gives the same inequality on $[1/2, y_B^0]$. Log-concavity implies that f is absolutely continuous on $[y_A^0, y_B^0]$. Since $f > 0$, integrating (25), which holds almost everywhere, shows that the left side of (24) is strictly increasing in $\bar{y}$ on $[y_A^0, y_B^0]$. By (22), the left side equals $F(y_A^0)$ at $\bar{y} = y_A^0$ and $1 + F(y_B^0)$ at $\bar{y} = y_B^0$. For every choice of a and b specified above, the right side, $F(a) + F(b)$, lies in $[F(y_A^0), 1 + F(y_B^0)]$. Continuity and strict monotonicity therefore imply that (24) has a unique solution in $[y_A^0, y_B^0]$. A direct check shows that this solution also lies between a and b.

For $a \neq b$, the prices in (23), evaluated at $\bar{y}(a, b)$ and with the endpoints in increasing order, therefore constitute an equilibrium. The same best-response argument applies to endpoint solutions, where one seller charges zero. Set $\bar{y}(1/2, 1/2) = 1/2$; this case adds no new cutoff and requires no message.

Learning. We show by induction that, after n rounds, every $y \in [y_A^0, y_B^0]$ has a generated cutoff b' satisfying $|F(b') - F(y)| \leq 2^{-n}$. Thus the cutoffs generated by the design approximate every location in this interval. The induction uses two properties. First, replacing an endpoint with a nearby known cutoff reduces the approximation error by at least half when distances are measured by F. Second, every target cutoff can be generated from one of the specified posteriors.

First, using an endpoint b' in place of b produces cutoffs satisfying, for every $a \in \{0, 1/2, 1\}$ and $b, b' \in [y_A^0, y_B^0]$,

$$|F(\bar{y}(a, b)) - F(\bar{y}(a, b'))| \leq \frac{1}{2} |F(b) - F(b')| . \tag{26}$$

Indeed, changing b to b' changes the right side of (24) by $F(b') - F(b)$. Integrating (25) between the two solutions gives (26).

Second, for every target cutoff $y \in [y_A^0, y_B^0]$, there exist $a \in \{0, 1/2, 1\}$ and $b \in [y_A^0, y_B^0]$

such that $y = \bar{y}(a, b)$. To show this, note that

$$\frac{1}{2} - F(y_A^0) = \int_{y_A^0}^{1/2} f(y)\, dy \geq \left(\frac{1}{2} - y_A^0\right) f(y_A^0) = F(y_A^0),$$

where the inequality follows because f is non-decreasing below $1/2$. Hence $F(y_A^0) \leq 1/4$, and symmetry implies $F(y_B^0) - F(y_A^0) \geq 1/2$. As b varies over $[y_A^0, y_B^0]$, the right side of (24) ranges over

$$\begin{aligned} &[F(y_A^0), F(y_B^0)] && \text{when } a = 0, \\ &[1/2 + F(y_A^0), 1/2 + F(y_B^0)] && \text{when } a = 1/2, \\ &[1 + F(y_A^0), 1 + F(y_B^0)] && \text{when } a = 1. \end{aligned}$$

These intervals overlap and together cover $[F(y_A^0), 1 + F(y_B^0)]$, which is the range of the left side of (24) for $\bar{y} \in [y_A^0, y_B^0]$. Thus every cutoff in this interval can be obtained from one of the specified posteriors.

We can now prove the claim by induction. Before the first round, the initial cutoff $1/2$ satisfies the bound for $n = 0$.

Suppose the claim holds after n rounds. Fix any $y \in [y_A^0, y_B^0]$. By the second property, choose a, b such that $y = \bar{y}(a, b)$. By the induction hypothesis, there is a generated cutoff b' satisfying $|F(b') - F(b)| \leq 2^{-n}$. By the end of the next round, the design has generated $\bar{y}(a, b')$. If $a = b' = 1/2$, this is the initial cutoff, which has already been distinguished. By (26),

$$|F(\bar{y}(a, b')) - F(y)| \leq \frac{1}{2}|F(b') - F(b)| \leq 2^{-(n+1)}.$$

This proves the induction step. Since F is continuous and strictly increasing, the generated cutoffs are dense in $[y_A^0, y_B^0]$.

Fix $\varepsilon > 0$ and $y^* \in (y_A^0, y_B^0)$. After finitely many rounds, there are generated cutoffs below and above y^*, both within ε of y^*. The platform's posterior is then concentrated between these cutoffs, and remains so thereafter. Hence the platform ε-learns every location in (y_A^0, y_B^0).

Finally, following a price different from the prescribed price, or a purchase inconsistent with the information from earlier transactions, the platform sends a separate message in every later period. After this message, sellers hold the prior F and charge $p_A = p_B = \tau/f(1/2)$. These prices are optimal under that belief. The message has zero probability on the equilibrium path, so assigning the prior is consistent with the belief requirement. Consumers follow the assumed purchasing rule at every history. This completes the design and equilibrium.

□

OB. Long-lived consumers

OB.1. Stationarity and the support of public beliefs

Scope of stationarity. We impose stationarity only for valuations in the support of the current public belief. This appendix explains why this qualification is needed to accommodate the standard full-disclosure equilibrium while maintaining sequential rationality.

Let μ_t denote the current seller's posterior over the consumer's valuation, which we call the public belief. For a consumer entering at history e, stationarity requires a fixed acceptance rule $\bar{\beta}_e(p, v)$ such that

$$\beta_e(h^{t-1}, m_t, p, v) = \bar{\beta}_e(p, v) \qquad \text{whenever } v \in \operatorname{supp} \mu_t.$$

This requirement applies to every current offer, including deviations by sellers. Consumers entering with the same public belief use the same stationary rule. If the consumer's valuation lies outside the current posterior support, his behavior may depend on history but must remain sequentially optimal.

Connection to bargaining models. The qualification is consistent with the treatment of remaining types in standard bargaining constructions. Gul et al. (1986, p. 162) define consumer stationarity only through the market state following a price below all previous prices. Their equilibrium construction specifies the reservation-price rule for consumers remaining in the market (p. 182), and their bargaining interpretation identifies residual demand with the seller's posterior over a single consumer's valuation (Section 6.2). In other words, Gul et al. require stationarity only along histories of decreasing prices. At other histories, this is not imposed. The related reservation-price approach is reviewed by Ausubel et al. (2002, Section 3). In the recursive bargaining model of Sobel and Takahashi (1983, p. 414), rejection updates the seller's existing posterior: an offer rejected by all supported types leaves its upper bound unchanged. Our definition makes the support qualification explicit in this tradition.

A full-disclosure example. Consider the initial consumer, with valuation drawn from $U[0,1]$ and discount factor $\delta_B = 1/2$, facing successive short-lived sellers. All past offers and purchase decisions are disclosed. The issue arises before any purchase, so replacement consumers play no role. In the familiar equilibrium, if the current posterior is $U[0,b]$, a consumer with valuation $v \le b$ has reservation price $r(v)$, and the seller posts $P(b)$, where

$$r(v) = \frac{2}{3}v \quad (v \le b), \qquad P(b) = \frac{b}{3}.$$

Indeed, the marginal consumer satisfies $r(v) = (1-\delta_B)v + \delta_B P(v)$, and the seller maximizes $p\big(1 - 3p/(2b)\big)$ over prices with positive demand.

Now consider a consumer with $v = 9/10$ and the following history:

1. The first seller charges $P(1) = 1/3$. The consumer deviates and rejects. Since the prescribed rule accepts this offer exactly when $v \ge 1/2$, the next seller's posterior is $U[0,1/2]$.
2. The next seller deviates and charges $p = 11/20$. Applying the original rule to this excluded

valuation would prescribe acceptance, because $11/20 < r(9/10) = 3/5$. Acceptance yields $9/10 - 11/20 = 7/20$.

3. Every valuation in the current posterior support rejects $11/20$. Rejection therefore leaves the posterior at $U[0, 1/2]$, and the following seller charges $P(1/2) = 1/6$. Rejecting and accepting that next offer yields

$$\delta_B \left(\frac{9}{10} - \frac{1}{6} \right) = \frac{11}{30} > \frac{7}{20}.$$

Thus, extending the original stationary rule to this excluded valuation violates sequential rationality.

The role of the qualification. After the first rejection, the consumer's valuation $9/10$ lies outside $\operatorname{supp} \mu_t = [0, 1/2]$. He can therefore respond to the actual continuation price. At this history his reservation price is

$$(1 - \delta_B)\frac{9}{10} + \delta_B P(1/2) = \frac{8}{15} < \frac{11}{20},$$

so he rejects the offer, as sequential rationality requires. Since $p = 11/20 > 1/3$, Gul et al. would not impose stationarity either, as the price sequence is not decreasing. The supported valuations retain their original stationary rule. Allowing this continuation behavior for excluded valuations preserves the familiar equilibrium price path and learning outcome. The example explains why the reservation-price rule must be applied on its stated support, while optimal behavior must still be specified at every information set.

OB.2. Full disclosure when consumers are long-lived

Under full disclosure, every seller observes the full history. The proof distinguishes posteriors whose support reaches zero, which we call the no-gap case, from posteriors with a strictly

positive lower endpoint, which we call the gap case. Let $Q, P, \sigma : V \to V$ satisfy

$$Q(x) = (1-\delta_B)x+\delta_B P(x), \qquad \sigma(b) \in \arg\max_{x\in[0,b]} Q(x)[F(b)-F(x)], \qquad P(b) = Q(\sigma(b)). \tag{27}$$

When the posterior is truncated from above at b, $Q(x)$ is the price that induces marginal consumer x, $\sigma(b)$ is the seller's optimal cutoff, and $P(b)$ is the associated offer. The proof of Proposition 1 below establishes existence by applying the weak-Markov construction of Fudenberg et al. (1985) with a zero seller discount factor. Fix such a solution, including a selection σ whenever the argmax is not a singleton. Set $b_0^* = 1$ and define

$$b_t^* = \sigma(b_{t-1}^*), \qquad t \geq 1. \tag{28}$$

Lemma 8 (Cutoff sequence properties)**.** *For any stationary solution of* (27)*, the sequence defined by* (28) *is strictly decreasing and converges to zero.*

Proof. Because $\sigma(0) = 0$, equation (27) implies $P(0) = Q(0) = \delta_B P(0)$. Hence, $P(0) = Q(0) = 0$. Fix $b > 0$. The objective defining $\sigma(b)$ is zero at both $x = 0$ and $x = b$, whereas it is strictly positive for every $x \in (0, b)$: indeed, $Q(x) \geq (1 - \delta_B)x > 0$ and $F(b) - F(x) > 0$. Thus, $0 < \sigma(b) < b$, so $\{b_t^*\}_{t\geq 0}$ is strictly decreasing and converges to some $\ell \geq 0$.

Suppose, toward a contradiction, that $\ell > 0$. Because Q is bounded, the value attained by the selected cutoff satisfies

$$Q(b_t^*)\Big[F(b_{t-1}^*) - F(b_t^*)\Big] \longrightarrow 0.$$

Fix any $x \in (0, \ell)$. For all sufficiently large t, this x is feasible in the problem defining $\sigma(b_{t-1}^*)$, and its objective value converges to

$$Q(x)\Big[F(\ell) - F(x)\Big] > 0.$$

This contradicts the optimality of $b_t^* = \sigma(b_{t-1}^*)$. Therefore, $\ell = 0$. □

Proof of Proposition 1. Step 1: The no-gap case. Before the first purchase, no replacement has entered, so the initial consumer remains active. In a stationary equilibrium, this consumer's pure strategy can be represented by a nondecreasing reservation-price function $r(v)$: type v accepts p if and only if $p \leq r(v)$ whenever v lies in the current posterior support. Equivalently, every price induces a valuation threshold x, with types $v \geq x$ accepting and types $v < x$ rejecting. This standard screening property is established by Fudenberg et al. (1985); see also Ausubel et al. (2002, Lemma 1).

The initial seller's posterior is $F_{[0,1]}$. Along a sequence of rejections, each rejected price induces a threshold x and reveals that $v < x$. Under full disclosure, the seller's posterior is therefore truncated from above to $F_{[0,x)}$. Hence, before the first acceptance, the posterior remains of the form $F_{[0,b)}$, and the game remains in the no-gap case.

Let $P(b)$ denote the equilibrium price at posterior $F_{[0,b)}$. Suppose a current price p induces an interior threshold x. Continuity of payoffs implies that type x is indifferent. The cutoff dynamics established above imply that $\sigma(x) < x$. If type x rejects, the next posterior is $F_{[0,x)}$ and the next seller posts $P(x)$; because the next cutoff is $\sigma(x)$, type x accepts that offer. Therefore,

$$x - p = \delta_B[x - P(x)].$$

Hence, the price that induces cutoff x is

$$Q(x) := (1 - \delta_B)x + \delta_B P(x). \tag{29}$$

Since the seller is short-lived, her cutoff and price solve

$$\sigma(b) \in \arg\max_{x \in [0,b]} Q(x)\Big[F(b) - F(x)\Big], \tag{30}$$

$$P(b) = Q(\sigma(b)). \tag{31}$$

The omitted factor $1/F(b)$ is positive and does not affect the maximizer. Equation (29) is the stationary consumer-indifference equation, equation (30) determines the seller's optimal cutoff, and equation (31) gives the associated price.

We apply the weak-Markov construction in Fudenberg et al. (1985, Propositions 1–2 and Appendices 1–2), allowing the seller and consumer to have different discount factors and setting the seller's discount factor to zero. This removes the seller's continuation-value term from the maximization problem and leaves the current-profit objective in (30). The consumer-indifference condition is unchanged. The construction for a positive lower bound and the compactness argument as that bound tends to zero continue to apply. We select a pure on-path maximizer; randomization may still be used to complete strategies after off-path histories.

We will also use the standard monotonicity property that σ and P can be selected to be nondecreasing; see Ausubel et al. (2002, Section 3.1.1). For completeness, we provide the direct revealed-preference argument. Fix $0 < b_1 < b_2$, and write $a_i = \sigma(b_i)$ and $q_i = Q(a_i) = P(b_i)$. If $a_2 < a_1$, optimality at the two states gives

$$q_1[F(b_1) - F(a_1)] \geq q_2[F(b_1) - F(a_2)],$$
$$q_2[F(b_2) - F(a_2)] \geq q_1[F(b_2) - F(a_1)].$$

Adding these inequalities implies $q_2 \geq q_1$. At state b_1, however, cutoff a_2 would then generate a weakly higher price and strictly greater demand than cutoff a_1, a contradiction. Thus, $a_2 \geq a_1$. If $a_2 > a_1$, optimality at b_2 requires $q_2 \geq q_1$; if $a_2 = a_1$, then $q_2 = q_1$. Hence both σ and P are nondecreasing.

Step 2: Successive rejections. We can now characterize the equilibrium path that ensues if every price offer is rejected. Equation (28) defines the cutoff sequence recursively. After $t-1$

rejections, the posterior is $F_{[0,b^*_{t-1})}$. Equations (29)–(31) imply that seller t posts

$$p^*_t = P(b^*_{t-1}) = (1-\delta_B)b^*_t + \delta_B P(b^*_t) = (1-\delta_B)b^*_t + \delta_B p^*_{t+1}.$$

Iterating this identity n times yields

$$p^*_t = (1-\delta_B)\sum_{k=0}^{n-1} \delta_B^k b^*_{t+k} + \delta_B^n p^*_{t+n}.$$

Because equilibrium prices lie in $[0,1]$, the last term converges to zero. Letting $n \to \infty$ proves (6).

Step 3: The gap case. Suppose that the first acceptance occurs in period t. Rejection in period $t-1$ revealed $v < b^*_{t-1}$, while acceptance in period t reveals $v \geq b^*_t$. At the beginning of period $t+1$, the posterior is therefore $F_{[b^*_t, b^*_{t-1})}$, a positive-gap posterior. We first specify equilibrium-path play with every seller charging b^*_t (Step 3a). We then verify that this price remains optimal after an off-path rejection that lowers the posterior's upper endpoint while preserving its lower endpoint b^*_t (Step 3b). This verification matters because a rejecting consumer remains active. After an off-path acceptance, the consumer exits, so we complete the resulting subgame with a stationary equilibrium (Step 3c).

Step 3a: Equilibrium-path behavior. Prescribe that every seller following the first acceptance posts b^*_t. For each consumer who enters this continuation, define the stationary reservation-price function

$$r_t(v) := \begin{cases} v, & v < b^*_t, \\ (1-\delta_B)v + \delta_B b^*_t, & v \geq b^*_t \end{cases}$$

and prescribe acceptance if and only if $p \leq r_t(v)$. Given the seller's continuation verified in Step 3b, accepting p gives type v the payoff $v-p$, whereas rejecting gives $\delta_B \max\{v - b^*_t, 0\}$. Thus, the reservation rule is sequentially optimal. In particular, every type in $[b^*_t, b^*_{t-1})$ accepts

b_t^*. A price $p > b_t^*$ induces the cutoff

$$x = \frac{p - \delta_B b_t^*}{1 - \delta_B}, \qquad \text{or equivalently} \qquad p = (1 - \delta_B)x + \delta_B b_t^*.$$

We next verify the seller's incentives at posterior $F_{[b_t^*, b_{t-1}^*)}$. For $x \in [b_t^*, b_{t-1}^*]$, define

$$\begin{aligned}\pi_t(x) &:= \Big[(1 - \delta_B)x + \delta_B b_t^*\Big]\Big[F(b_{t-1}^*) - F(x)\Big] \\ &= \Big\{Q(x) + \delta_B[b_t^* - P(x)]\Big\}\Big[F(b_{t-1}^*) - F(x)\Big],\end{aligned}$$

as the seller's unnormalized profit from posting $p = (1 - \delta_B)x + \delta_B b_t^*$, which induces cutoff x. The second equality follows from $Q(x) = (1 - \delta_B)x + \delta_B P(x)$. Any price $p < b_t^*$ is accepted by every type and yields strictly less profit than $p = b_t^*$; hence, a downward deviation is not optimal. It remains to consider an upward deviation $p > b_t^*$, which induces $x \in (b_t^*, b_{t-1}^*]$. Two inequalities rule out such a deviation. First, because $b_t^* = \sigma(b_{t-1}^*)$, optimality in the preceding no-gap subgame gives

$$Q(x)\Big[F(b_{t-1}^*) - F(x)\Big] \le Q(b_t^*)\Big[F(b_{t-1}^*) - F(b_t^*)\Big]. \tag{32}$$

Second, since $\sigma(b_t^*) < b_t^*$ and P is nondecreasing,

$$P(b_t^*) = Q(\sigma(b_t^*)) = (1 - \delta_B)\sigma(b_t^*) + \delta_B P(\sigma(b_t^*)) \le (1 - \delta_B)\sigma(b_t^*) + \delta_B P(b_t^*),$$

so $P(b_t^*) \le \sigma(b_t^*) < b_t^*$. Because P is nondecreasing and $P(b_t^*) \le b_t^*$,

$$[b_t^* - P(x)]\Big[F(b_{t-1}^*) - F(x)\Big] \le [b_t^* - P(b_t^*)]\Big[F(b_{t-1}^*) - F(b_t^*)\Big]. \tag{33}$$

If $P(x) > b_t^*$, the left-hand side is nonpositive. Otherwise, both nonnegative factors on the left are weakly smaller than their counterparts on the right. Expanding the preceding expression

for $\pi_t(x)$, we can now conclude that

$$\begin{aligned}\pi_t(x) &= Q(x)\Big[F(b^*_{t-1}) - F(x)\Big] + \delta_B[b^*_t - P(x)]\Big[F(b^*_{t-1}) - F(x)\Big] \\ &\leq \Big\{Q(b^*_t) + \delta_B[b^*_t - P(b^*_t)]\Big\}\Big[F(b^*_{t-1}) - F(b^*_t)\Big] \\ &= b^*_t\Big[F(b^*_{t-1}) - F(b^*_t)\Big] = \pi_t(b^*_t),\end{aligned}$$

where the inequality follows by multiplying (33) by δ_B and adding it to (32); the final two equalities follow from $Q(b^*_t) = (1-\delta_B)b^*_t + \delta_B P(b^*_t)$ and the definition of $\pi_t(b^*_t)$, respectively. Thus, $\pi_t(x) \leq \pi_t(b^*_t)$ for every $x \in (b^*_t, b^*_{t-1}]$. Hence, no deviation above b^*_t is optimal.

At the equilibrium price b^*_t, every type in $[b^*_t, b^*_{t-1})$ accepts. Acceptance does not change the posterior, and the consumer exits and is replaced by an identical consumer. Hence, the same positive-gap posterior recurs in every subsequent period. This is precisely the post-acceptance behavior stated in the proposition.

Step 3b: Continuation following an off-path rejection. Suppose that a seller facing posterior $F_{[b^*_t, b^*_{t-1})}$ posts an off-path price above b^*_t that induces cutoff $b \in (b^*_t, b^*_{t-1})$, and the consumer rejects. Under full disclosure, the next seller's posterior is $F_{[b^*_t, b)}$. More generally, repeated rejections can generate any posterior of this form with $b < b^*_{t-1}$. Because the same consumer remains active after a rejection, we must verify that the next seller still posts b^*_t.

For $x \in [b^*_t, b]$, let

$$\pi_{t,b}(x) := \Big[(1-\delta_B)x + \delta_B b^*_t\Big][F(b) - F(x)]$$

denote the next seller's unnormalized profit from inducing cutoff x. Log-concavity of f implies that the truncated tail $F(b) - F(x)$ is log-concave in x. The other factor is positive and log-concave, so $\pi_{t,b}$ is log-concave. Step 3a shows that π_t, which corresponds to the upper

endpoint b^*_{t-1}, is maximized at b^*_t and hence has a nonpositive right derivative there. Since

$$\begin{aligned}\pi'_{t,b}(b^*_t+) &= (1-\delta_B)[F(b)-F(b^*_t)] - b^*_t f(b^*_t)\\ &\le (1-\delta_B)[F(b^*_{t-1})-F(b^*_t)] - b^*_t f(b^*_t)\\ &= \pi'_t(b^*_t+) \le 0,\end{aligned}$$

log-concavity implies that $\pi_{t,b}$ is nonincreasing on $[b^*_t, b]$. Thus, the next seller optimally posts b^*_t. The same argument applies after every subsequent rejection. This validates the continuation payoff $\delta_B \max\{v - b^*_t, 0\}$ used in Step 3a and establishes sequential rationality of the prescribed reservation rule.

Step 3c: Other off-path histories. An off-path acceptance may create a positive-gap posterior whose lower endpoint differs from b^*_t. The accepting consumer exits, and the current seller is short-lived, so play in the resulting subgame affects neither player's current payoff.

We complete every such positive-gap subgame using the weak-Markov construction of Fudenberg et al. (1985, Proposition 1), specialized to a zero seller discount factor as in Step 1. No-gap continuations are completed using the stationary equilibrium constructed in Step 1.

Together, Steps 3a–3c define a stationary equilibrium continuation after the first acceptance. Along its equilibrium path, every seller posts b^*_t and every entering consumer with $v \ge b^*_t$ accepts. Together with Lemma 8, this proves part 1.

Step 4: Learning. If the first t offers are rejected, part 1 implies that $v^* \in [0, b^*_t)$, and full disclosure gives the platform posterior $F_{[0,b^*_t)}$. If the first acceptance occurs in period $t' \le t$, then the previous offers were rejected and the period-t' offer was accepted. Hence, $v^* \in [b^*_{t'}, b^*_{t'-1})$, with the upper endpoint included when $t' = 1$, and the platform's posterior is the prior conditional on that interval. By part 1, every subsequent seller posts $b^*_{t'}$ and every replacement consumer accepts. These purchases reveal no additional information, so the posterior remains unchanged. □

Corollary 1 (Rejection-path offers as consumers become patient)**.** *For any family, indexed by $\delta_B \in (0,1)$, of the stationary equilibria constructed in Proposition 1, and every fixed $t \geq 1$,*

$$\lim_{\delta_B \uparrow 1} p_t^*(\delta_B) = 0.$$

Proof. Consider any such family of stationary equilibria. We first prove that $p_1^*(\delta_B) \to 0$. Fix δ_B and suppress its argument. Let

$$m_t := F(b_{t-1}^*) - F(b_t^*)$$

denote the probability mass of valuations for which a purchase occurs for the first time in period t. Two bounds drive the argument when prices remain bounded away from zero. The probability mass purchasing for the first time must be at least of order $1-\delta_B$, but successive masses must shrink by a factor of that same order. Applying the second bound twice yields a contradiction as $\delta_B \uparrow 1$. Fix $\varepsilon \in (0,1]$. Whenever $p_t^* \geq \varepsilon$, equation (6) implies $b_t^* \geq p_t^* \geq \varepsilon$. Seller t can deviate to the price $Q(\varepsilon/2)$ and induce cutoff $\varepsilon/2$. Because $Q(\varepsilon/2) \geq (1-\delta_B)\varepsilon/2$, optimality of p_t^* gives

$$p_t^* m_t \geq Q(\varepsilon/2)\Big[F(b_{t-1}^*) - F(\varepsilon/2)\Big] \geq c_\varepsilon(1-\delta_B),$$

where

$$c_\varepsilon := \frac{\varepsilon}{2}\Big[F(\varepsilon) - F(\varepsilon/2)\Big] > 0.$$

Since $p_t^* \leq 1$, it follows that

$$m_t \geq c_\varepsilon(1-\delta_B) \qquad \text{whenever } p_t^* \geq \varepsilon. \tag{34}$$

Seller t can also post $p_{t+1}^* = Q(b_{t+1}^*)$, thereby inducing cutoff b_{t+1}^* and selling to the mass

$m_t + m_{t+1}$. Seller optimality therefore implies

$$p_t^* m_t \geq p_{t+1}^*(m_t + m_{t+1}).$$

Moreover, equation (6) gives $p_{t+1}^* \leq b_{t+1}^* < b_t^*$. The price recursion established in the proof of Proposition 1 yields

$$0 < p_t^* - p_{t+1}^* = (1 - \delta_B)(b_t^* - p_{t+1}^*) \leq 1 - \delta_B.$$

Consequently,

$$m_{t+1} \leq \frac{1 - \delta_B}{\varepsilon} m_t \qquad \text{whenever } p_{t+1}^* \geq \varepsilon. \tag{35}$$

Suppose, toward a contradiction, that $p_1^*(\delta_B)$ does not converge to zero. Then there are $\eta \in (0, 1]$ and a sequence $\{\delta_{B,n}\}$ converging to one along which $p_1^*(\delta_{B,n}) \geq \eta$. Set $\varepsilon = \eta/2$. For all sufficiently large n,

$$1 - \delta_{B,n} < \min\left\{\frac{\eta}{4}, c_\varepsilon \varepsilon^2\right\}.$$

Fix such an n, set $\delta_B = \delta_{B,n}$, and suppress this argument. Because prices decrease by at most $1 - \delta_B$ per period, $p_1^*, p_2^*, p_3^* \geq \varepsilon$. Applying (35) twice and using $m_1 \leq 1$ gives

$$m_3 \leq \left(\frac{1 - \delta_B}{\varepsilon}\right)^2 m_1 \leq \left(\frac{1 - \delta_B}{\varepsilon}\right)^2.$$

By (34), however,

$$m_3 \geq c_\varepsilon(1 - \delta_B),$$

which contradicts $1 - \delta_B < c_\varepsilon \varepsilon^2$. Hence $p_1^*(\delta_B) \to 0$. Finally, the monotonicity of P and the strict decrease of $\{b_t^*\}$ imply $0 \leq p_t^*(\delta_B) \leq p_1^*(\delta_B)$ for every $t \geq 1$. Thus, $p_t^*(\delta_B) \to 0$ for every fixed t. □

The following example perturbs the uniform prior and shows that patience can raise or lower the cutoff in each period relative to the benchmark with short-lived consumers.

Example 2 (Cutoffs can move in either direction)**.** *Consider* $F_\xi(v) = v + \xi(v^2 - v)$, $\xi \in (-1, 1)$. *Figure 5 plots the percentage change in the period-t cutoff relative to its myopic value for the first ten periods. For* $\xi = 0.5$*, the displayed increases in patience raise these cutoffs; for* $\xi = -0.5$*, they lower them.*

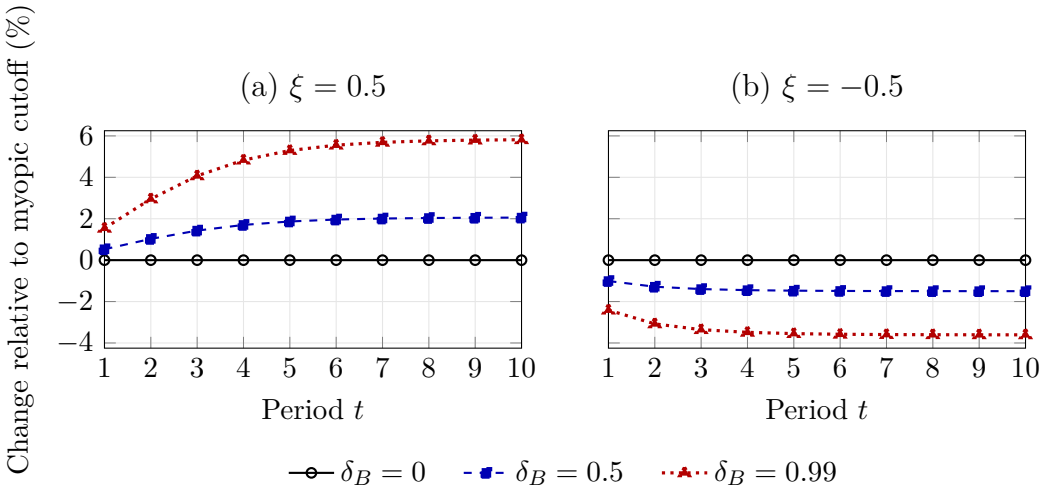


Figure 5: **Percentage change in the full-disclosure cutoffs.** Panels (a) and (b) use $\xi = 0.5$ and $\xi = -0.5$, respectively. Each curve plots $\Delta_t(\delta_B, \xi) := 100\left[\frac{b_t^*|_{\delta_B,\xi}}{b_t^*|_{0,\xi}} - 1\right]$ for the first ten periods.

Details for Example 2. For $|\xi| < 1$, the density is strictly positive and

$$\frac{d^2}{dv^2} \log f_\xi(v) = -\frac{4\xi^2}{[1 + \xi(2v - 1)]^2} \leq 0,$$

so the maintained log-concavity assumption is satisfied. Normalize the inducing-price schedule by setting $q_{\delta_B,\xi} := Q_{\delta_B,\xi}/(1 - \delta_B)$. Equation (27) becomes

$$q_{\delta_B,\xi}(x) = x + \delta_B q_{\delta_B,\xi}\Big(\sigma_{\delta_B,\xi}(x)\Big), \qquad \sigma_{\delta_B,\xi}(b) \in \arg\max_{x \in [0,b]} q_{\delta_B,\xi}(x)[F_\xi(b) - F_\xi(x)]. \tag{36}$$

Numerical construction of Figure 5. For $\xi \in \{-0.5, 0.5\}$ and $\delta_B \in \{0, 0.5, 0.99\}$, we approximate $\sigma_{\delta_B,\xi}(b)/b$ and $q_{\delta_B,\xi}(b)/b$ by degree-40 Chebyshev polynomials on 401 Chebyshev–Lobatto nodes. Starting from $\sigma(b) = b/2$, we alternate between computing the normalized price schedule for a given cutoff rule and updating that rule. For a given σ, the price schedule is

$$q_{\delta_B,\sigma}(x) = \sum_{k=0}^{\infty} \delta_B^k \sigma^k(x).$$

Here $\sigma^0(x) = x$ and $\sigma^{k+1}(x) = \sigma(\sigma^k(x))$. We then update $\sigma(b)$ using the unique interior solution to

$$q'_{\delta_B,\sigma}(x)[F_\xi(b) - F_\xi(x)] - q_{\delta_B,\sigma}(x) f_\xi(x) = 0.$$

The continuation series is truncated when its largest new term is below 10^{-16}, and the policy iteration stops when the sup-norm change is below 3×10^{-14}. If $|\sigma(x)| \leq \eta|x|$ for some $\eta < 1$, the error after truncating the series at K is bounded by $x(\delta_B\eta)^{K+1}/(1 - \delta_B\eta)$. We obtain the underlying cutoff paths recursively from $b_0^* = 1$ and $b_t^* = \sigma(b_{t-1}^*)$ and compute the plotted percentage changes from them. Repeating the calculation with 301 nodes and degree 30, and with 601 nodes and degree 55, leaves all plotted values unchanged at the reported precision; the largest first-order-condition residual along the paths is below 3.2×10^{-14}. At every collocation state, we also verify the seller's objective over the full feasible interval. The computed normalized inducing-price schedule is strictly log-concave in all six cases, so the objective is strictly log-concave and its first-order-condition root is the unique global maximizer. An independent grid search finds no profitable deviation, with comparison error below 3.3×10^{-15}. □

OB.3. Learning for arbitrarily patient consumers

Use Appendix B.1's notation and assumptions, writing $\delta = \delta_B$ and $\alpha = 1 - \delta$.

Theorem 4 (Learning with long-lived consumers)**.** *For every $\varepsilon > 0$ and every $\delta_B \in (0, 1)$ the platform can ε-learn all valuations in finite time.*

Proof. The construction has two stages. A finite sequence of declining offers either identifies an interval between consecutive cutoffs or places the valuation in a bottom interval of length less than ε. Following an acceptance, the construction from Appendix B.1 refines the identified interval.

Initial offers. For each integer $K \geq 1$, choose $1 = \tau_0 > \cdots > \tau_K > 0$ and $P_1 > \cdots > P_K > P_{K+1} = a$, where $a = \widehat{p}(\tau_K)$, $P_k = \alpha\tau_k + \delta P_{k+1}$, and, for $k = 1, \ldots, K$,

$$\alpha[F(\tau_{k-1}) - F(\tau_k)] = P_k f(\tau_k). \tag{37}$$

Construct backward from $\tau_K = x \in (0,1)$ and $P_{K+1} = \widehat{p}(x)$, using the price equation and $\tau_{k-1} = F^{-1}(F(\tau_k) + P_k f(\tau_k)/\alpha)$. Since $P_k > \alpha\tau_k$, (37) makes static profit's derivative positive at τ_k, giving $\tau_k < \widehat{p}(\tau_{k-1}) \leq \widehat{p}(1) < 1$. The recursion is continuous and feasible for small x. At its maximal feasible endpoint $x^* \leq \widehat{p}(1) < 1$, limits exist recursively; intermediate inverse arguments stay below $F(\widehat{p}(1)) < 1$. The final argument must reach one; otherwise, continuity extends feasibility; hence $\tau_0 = 1$. Backward induction gives $P_{k+1} < P_k < \tau_k$.

To obtain the desired precision, let $\underline{f} = \min f > 0$ and $\bar{f} = \max f$. The static first-order condition gives $\widehat{p}(u) \leq qu$, where $q = \bar{f}/(\bar{f} + \underline{f}) < 1$. Choose a finite K such that $q^K < \varepsilon$. Then $\tau_K < q^K < \varepsilon$.

Design. Send D_1 to the first seller (replacing Appendix B.1's rule 1).

In the postulated equilibrium, sellers offer P_k after D_k, for $1 \leq k \leq K+1$. The initial consumer and consumers entering at D_{K+1} accept $p \leq R(v)$, where

$$R(v) = \max\Big(\{r(v)\} \cup \{P_j : 1 \leq j \leq K,\ \tau_j \leq v\}\Big).$$

Here $r(v) = v - \delta(v-a)^+$. This is proven below once we complete the description of the design.

If the initial consumer or a consumer who entered when the message was D_{K+1} rejects a price $p \in B_k = (P_{k+1}, P_k] \cap [r(\tau_k), P_k]$, $k < K$, send the next seller D_{k+1} with probability $\psi_k(p) = [p - r(\tau_k)]/[\delta(P_{k+1} - a)]$ and D_{K+1} otherwise; after rejections of other prices send D_{K+1}. Since $P_k - r(\tau_k) = \delta(P_{k+1} - a)$, these probabilities are valid and $\psi_k(P_k) = 1$. The bands are disjoint and exclude a.

Note that the rule $R(\cdot)$ implies that the consumer buys iff $v \geq \tau_k$ in B_k, and iff $p \leq r(v)$ outside the bands. To see why additional acceptances occur only within the bands, suppose $r(v) < p \leq P_j$ and $\tau_j \leq v$. If $P_{k+1} < p \leq P_k$, then $j \leq k < K$, so $v \geq \tau_k$ and $p > r(\tau_k)$. Thus $p \in B_k$.

Consequently, on path (of the postulated equilibrium), acceptance of P_k at D_k, $k \leq K$, identifies $I_k = [\tau_k, \tau_{k-1})$ (including 1 for $k = 1$).

When the initial consumer, or a consumer who entered at D_{K+1}, accepts P_k at D_k, the next consumer starts at the second test of the construction in Appendix B.1, applied to I_k. Replace b_0 by $c_k = P_{k+1}$, b_1 by τ_k, and the upper endpoint 1 by τ_{k-1}. Accordingly, use $r_k(v) = v - \delta(v - c_k)^+$ and replace $1 - F(v)$ by $F(\tau_{k-1}) - F(v)$ in the profit expressions. The test prices, lottery probabilities, and purchasing rules (13)–(15) are otherwise unchanged.

For each interval I_k, use separate messages $m_{p,k,t}$, $m_{r,k,t+1}$, and $m_{b,k,t}$, beginning with $m_{p,k,2}$. Follow Appendix B.1's instructions for sending these messages, with two replacements:

- After rejection by a consumer who entered at $m_{p,k,t}$, replace $\varnothing$ by D_{k+1}.
- Whenever Appendix B.1 prescribes $\varnothing$ after an acceptance, send D_{K+1} instead.

The rejection instructions continue to apply while that consumer remains active. The index t records the test number; deviations that change its calendar date do not change its message.

By (37), $r_k(v)[F(\tau_{k-1}) - F(v)]$ has derivative zero at τ_k and decreases thereafter by log-concavity. When consecutive cutoffs coincide at $u \geq \tau_k$, Appendix B.1's prescribed price earns the seller $\alpha\delta(u - c_k) > 0$ more than selling to everyone at $r_k(u)$. Continuity permits a finite grid with all learned intervals shorter than ε and each test more profitable than selling to everyone at the preceding cutoff's reservation price (9).

Incentives. *Sellers.* At D_k, $k \leq K$, the seller believes $v \sim F_{[0,\tau_{k-1})}$ and that the consumer uses the threshold rule R. On this interval, $R(v) \leq \alpha v + \delta P_{k+1}$. At $p = \alpha y + \delta P_{k+1}$, $0 \leq y \leq \tau_{k-1}$, consumers must have $v \geq y$, so profit is at most $(\alpha y + \delta P_{k+1})[F(\tau_{k-1}) - F(y)]/F(\tau_{k-1})$. By the initial-offer condition (37) and log-concavity, this bound is maximized at $y = \tau_k$. At P_k, the purchasing consumers are exactly those with $v \geq \tau_k$, so the bound is attained. Prices below δP_{k+1} earn less than the bound at $y = 0$; prices above $\alpha\tau_{k-1} + \delta P_{k+1}$ sell nothing. At D_{K+1},

belief $F_{[0,\tau_K)}$ gives $R = r \leq v$; static profit bounds actual profit and is attained at $a = \widehat{p}(\tau_K)$.

Consumers. Outside the bands, rejection brings a forever; buying is optimal iff $v - p \geq \delta(v-a)^+$, or $p \leq r(v)$. Check the bands in decreasing order of k. For $p \in B_k$, write $\psi = \psi_k(p)$. Rejection brings P_{k+1} with probability ψ and a otherwise. The lottery definition gives $p = \alpha\tau_k + \delta[\psi P_{k+1} + (1-\psi)a]$. If $v \geq \tau_k$, the already-verified continuation buys next period at either price. Buying now therefore gains $v - p - \delta[\psi(v - P_{k+1}) + (1-\psi)(v-a)] = \alpha(v-\tau_k) \geq 0$. If $v < \tau_k$, the already-verified continuation is at least as good as buying next period or never buying. Waiting thus gives at least $\delta[\psi(v-P_{k+1}) + (1-\psi)(v-a)]$, so buying gains at most $\alpha(v-\tau_k) < 0$. Thus the cutoff τ_k, and hence R, is optimal. For a consumer who entered at $m_{p,k,t}$, rejection of c_k sends D_{k+1} again, since $c_k < r_k(\tau_k)$. Thus, this consumer continues to face c_k until purchase. The remaining consumer incentive checks in Appendix B.1 apply with the same substitutions.

Beliefs and learning. When strategies are followed, each seller's belief is F restricted to her message's interval. If deviations make a message arrive earlier or later, the seller retains the same belief about v. Initial and repeated rejections of P_{k-1} produce D_k; accepting P_k at D_k starts Appendix B.1 within I_k. Within I_k, another rejection of the same test price sends $m_{r,k,t+1}$ again. A consumer who entered at $m_{b,k,t}$ receives that same message after either decision. An initial P_K rejection reveals $[0, \tau_K)$ and sends D_{K+1}. There, both acceptance and rejection of a send D_{K+1}; the next seller does not observe the decision, so her belief remains $F_{[0,\tau_K)}$. Observing a seller's price leaves beliefs about valuation and history unchanged: she knows only the date and message (no-signaling). Reservation rules remain fixed. Distinct entry messages give distinct beliefs; consumers entering at D_1 or D_{K+1} use R, with no entry at intermediate D_k (Definition 2).

All K rejections identify $[0, \tau_K)$, shorter than ε; the first acceptance identifies I_k, which Appendix B.1 refines. With N_k refinement tests, $T(\varepsilon, \delta_B) = K + \max_k N_k + 1$ suffices. □

OB.4. Consistency of Beliefs

OB.4.a. Common perturbations of consumers and sellers

We construct a common perturbation of consumers' and sellers' behavior. At every fixed date and permitted message, the resulting Bayesian beliefs converge to the public valuation beliefs supporting Theorem 2 and to sellers' beliefs about the past history that generated their message. The perturbations allow mistakes and serve to justify beliefs. Sequential rationality is verified only for the limiting strategies, as in the proof of Theorem 2. The next subsection also requires approximate sequential rationality for the perturbed strategies, using Myerson and Reny's notion of conditional ϵ-equilibrium.

Notation. Fix the information design and strategies in the theorem's construction. Recall the grid $0 < b_0 < \cdots < b_N < 1$, let $b_{N+1} = 1$, and write

$$r(v) = \begin{cases} v, & v < b_0, \\ (1-\delta_B)v + \delta_B b_0, & v \geq b_0, \end{cases} \qquad p_n = (1-\delta_B)b_n + \delta_B r(b_{n-1}).$$

The prescribed seller price $p^*(m)$ equals p_n at $m_{p,n}$, $r(b_{n-1})$ at $m_{r,n+1}$, b_n at $m_{b,n}$, and b_0 at $\varnothing$. For $n = 2, \ldots, N$, put

$$J_n = [b_{n-1}, b_n), \qquad W_n = [r(b_n), p_n], \qquad \lambda_n(p) = \frac{p - r(b_n)}{p_n - r(b_n)} \quad (p \in W_n).$$

A consumer entering at $m_{p,n}$ accepts a price in W_n exactly when $v \geq b_n$; outside W_n, he accepts exactly when $p \leq r(v)$. Whenever this consumer rejects $p \in W_n$, the design sends $m_{r,n+1}$ with probability $\lambda_n(p)$ and otherwise sends $\varnothing$. Other rejections send $\varnothing$. These rules depend on his *entry message* and continue to apply while he remains active.

A common perturbation. Choose a feasible price $P > 1$ and an atomless distribution Q with a strictly positive density on the price set.[25] For $0 < \eta < 1/3$, perturb every seller's strategy to

$$\sigma_t^\eta(dp \mid m) = (1 - \eta - \eta^2)\delta_{p^*(m)}(dp) + \eta\delta_P(dp) + \eta^2 Q(dp), \tag{38}$$

where δ_p is the point mass at p. At every consumer information set, replace the prescribed acceptance probability β by

$$\beta^\eta = (1 - \rho_\eta)\beta + \frac{\rho_\eta}{2}, \qquad \rho_\eta = e^{-1/\eta}. \tag{39}$$

Use fresh, independent random draws at every move, keeping the prior and information design fixed. These rules define a single probability law $\mathbb{P}^\eta$ over play. Under the prescribed strategy, consumers reject P at every valuation. Offers of P have probability η, whereas draws from Q have probability η^2. Thus, in the leading histories identified below, waiting is generated by rejections of P, which do not select among valuations. Consumer mistakes vanish faster than every power of η.

Let h^{t-1} record past prices, decisions, and messages. At each date t and message m permitted at that date, define the seller's joint posterior

$$\Pi_t^\eta(\cdot \mid m) = \mathbb{P}^\eta\Big((v, h^{t-1}) \in \cdot \mid m_t = m\Big).$$

Its valuation marginal $\mu_t^\eta(\cdot \mid m)$ is the public belief. Individual prices drawn from Q have probability zero, but every message has positive probability under $\mathbb{P}^\eta$ at each date at which the information design permits it to be sent. Bayes' rule therefore determines these posteriors by integrating over the histories that generate the message.

[25] With nonnegative unbounded prices, one may take $P = 2$ and $Q(dp) = e^{-p}dp$. With prices in $[0, \bar{p}]$, $\bar{p} > 1$, take $P \in (1, \bar{p})$ and let Q be uniform.

Proposition 8 (Common perturbation consistency)**.** *At every fixed date t and permitted message m, the perturbed posterior $\Pi_t^\eta(\cdot \mid m)$ converges to a limiting belief over valuations and past histories. Convergence is in total variation: the largest difference between the probabilities that the perturbed and limiting beliefs assign to any event tends to zero as $\eta \downarrow 0$. The limiting valuation marginal is*

$$\mu_t(\cdot \mid m) = \begin{cases} F, & m = m_{p,1},\ t = 1, \\ F_{[b_{n-1},1]}, & m = m_{p,n},\ t = n,\ 2 \le n \le N, \\ F_{[0,b_1)}, & m = \varnothing,\ t \ge 2, \\ F_{J_n}, & m = m_{r,n+1},\ t \ge n+1,\ 2 \le n \le N, \\ F_{[b_n,b_{n+1})}, & m = m_{b,n},\ t \ge n+3,\ 1 \le n < N, \\ F_{[b_N,1]}, & m = m_{b,N},\ t \ge N+1. \end{cases} \tag{40}$$

Here F_I denotes the prior conditional on I. These limiting beliefs support the strategies and outcomes of the construction in Theorem 2.

Proof. Messages with positive equilibrium probability. On path, rejection of p_1 leads to $\varnothing$ forever. Rejection of p_n, $n \ge 2$, leads to $m_{r,n+1}$; acceptance of $r(b_{n-1})$ then leads to $m_{b,n-1}$ forever. Acceptance of p_N leads to $m_{b,N}$ forever. Each message in (40) therefore has positive equilibrium probability at every date listed there, except $m_{r,n+1}$ at $t \ge n+2$. At any fixed date, the joint distribution of the valuation, past history, and current message converges to its equilibrium distribution, because the probability of any earlier tremble tends to zero. Conditioning proves the claim for all positive-probability cases.

Off-path rejection messages. On the equilibrium path, message $m_{r,n+1}$ occurs only in period $n+1$. We now consider this same message occurring at a later date $t \ge n+2$, following deviations that keep the consumer who entered in period n in the market.

Fix $n \geq 2$ and such a date t, and set

$$c_n = \int_{W_n} \lambda_n(q) Q(dq) > 0.$$

First consider histories without consumer mistakes. Reaching $m_{p,n}$ requires acceptance of the exact prescribed prices $p_1, \ldots, p_{n-1}$, so $v \geq b_{n-1}$. Generating $m_{r,n+1}$ requires a rejection in W_n, so $v < b_n$. Thus $v \in J_n$, and the consumer who entered at date n must still be active. On J_n, both prescribed continuation prices, b_0 and $r(b_{n-1})$, are accepted. Keeping this consumer active therefore requires a seller tremble in every period from $n+1$ through $t-1$.

The leading histories use prescribed prices through date n, with the first $n-1$ accepted and p_n rejected. Offers P are then rejected at dates $n+1, \ldots, t-2$. Finally, a draw $q \in W_n$ from Q is rejected and the lottery selects $m_{r,n+1}$. The intermediate sequence is empty if $t = n+2$.

Including the requirement that all consumer decisions are prescribed, the joint probability measure of these histories over (v, q) is

$$(1-\eta-\eta^2)^n (1-\rho_\eta/2)^{t-1} \eta^{t-n} \mathbf{1}_{J_n}(v) F(dv)\, \mathbf{1}_{W_n}(q) \lambda_n(q) Q(dq). \tag{41}$$

Every other history producing this message without a consumer mistake requires at least one additional factor of η: either the first rejection follows a seller tremble, or an intermediate waiting price is drawn from Q instead of being P. For fixed n, t, their total probability is $O(\eta^{t-n+1})$.

Histories with consumer mistakes. The probability of any consumer decision departing from the prescribed rule before date t is at most $(t-1)\rho_\eta/2$. By (41), the conditioning message has probability at least a positive constant times η^{t-n}. Hence the conditional weight of all remaining histories is

$$O\left(\eta + \frac{\rho_\eta}{\eta^{t-n}}\right) \longrightarrow 0.$$

Conditional on the histories in (41), the valuation and final price are independent, with laws

$$v \sim F_{J_n}, \qquad q \sim \frac{\mathbf{1}_{W_n}(q)\lambda_n(q)Q(dq)}{c_n};$$

all other history coordinates are fixed. This law is independent of η, and the remaining histories have vanishing weight, proving joint convergence. The active consumer retains entry message $m_{p,n}$. □

OB.4.b. Myerson and Reny's perfect conditional ϵ-equilibrium

Online Appendix OB.4.a constructs a common perturbation of consumers and sellers that generates the beliefs supporting Theorem 2. We now extend that construction to establish Myerson and Reny's (2020) *perfect conditional ϵ-equilibrium*. The additional requirement is that behavior is approximately optimal throughout an admissible net of perturbations. We first define these terms, then explain the additions to the preceding construction and verify optimality.

A finite upper bound on prices. For this subsection, fix any finite $\bar{p} > 1$ and restrict prices to $[0, \bar{p}]$. All prescribed prices in the construction are below one, and the prescribed consumers reject every price above one, so the construction and its optimality arguments remain valid. Use the bounded-price specification from Online Appendix OB.4.a: choose $P \in (1, \bar{p})$ and let Q be uniform on $[0, \bar{p}]$.

The bound matters because consumer mistakes give even very high offers a small positive acceptance probability. Without a price bound, sellers could earn arbitrarily large expected profits.

Conditional ϵ-equilibrium. Write $g = (\sigma, \beta)$ for the prescribed seller and consumer strategies, respectively. A *conditional ϵ-equilibrium* allows no continuation deviation to gain more than ϵ conditional on any positive-probability event in a player's information. For a

seller, such an event can be a particular message at a given date. For a consumer, it can be a measurable collection of valuations, observed histories, and current prices; an exact observation is included whenever it has positive probability.

More precisely, let Z be such an event at a decision of player i, and let d_i be any strategy change from that decision onward. Writing $U_i(\cdot \mid Z; F')$ for conditional expected payoff when the valuation distribution is F', the requirement is

$$U_i(d_i, g_{-i} \mid Z; F') \leq U_i(g \mid Z; F') + \epsilon. \tag{42}$$

The deviation keeps the player's earlier behavior unchanged. All feasible continuation strategies are considered, including every price in $[0, \bar{p}]$.

Because Q is continuous, each individual price has probability zero. We therefore add small atoms on finite sets of prices and valuations. Later indices retain the earlier sets and may add new points, so each fixed price and valuation eventually receives positive probability. Together with an accuracy level, these expanding sets index a net of perturbations.

Definition 4 (Admissible nets and perfect conditional ϵ-equilibrium)**.** Call a net (g^α, F^α) *admissible* for (g, F) if:

(i) The largest difference between g^α and g in the probability of any action event, over all players and all information sets, tends to zero. The largest difference between F^α and F in the probability of any valuation event also tends to zero.

(ii) For each fixed information set and each fixed feasible action, there is an index beyond which that action has positive probability. This index may depend on the information set and action. The same requirement applies separately to each valuation in a set of prior probability one.

The information design is kept fixed. For $\epsilon > 0$, the profile g is a *perfect conditional ϵ-equilibrium* if there is an admissible net for which every g^α is a conditional ϵ-equilibrium under F^α. Such a net is called an *ϵ-test net*.

These are the definitions of Myerson and Reny (2020, Sections 4 and 6.1), applied to all successive sellers and consumers in our model. Their formal model has finitely many players; we use the same requirements for our countably many successive sellers and consumers, and verify them directly below.[26]

Extending the common perturbation. We retain the seller mixture and the consumer perturbation in (38)–(39). We add small probabilities to individual prices and valuations to obtain an admissible net, and stop mistakes after a finite date to control conditional optimality at all dates.

An index is $\alpha = (T, D, V, k)$, where $T \geq N + 3$ is an integer, $D \subset [0, \bar{p}]$ and $V \subset [0, 1]$ are nonempty finite sets, and k is a positive integer. Later indices have larger T and k and contain the earlier sets D and V. The number $1/k$ will bound the gain from any continuation deviation.

For a small $\eta \in (0, 1/4)$, use through date T

$$\sigma_t^\alpha(dp \mid m) = (1 - \eta - \eta^2 - \rho_\eta)\delta_{p^*(m)}(dp) + \eta\delta_P(dp) + \eta^2 Q(dp) + \rho_\eta D^\circ(dp), \tag{43}$$

$$\beta^\alpha = \beta^\eta = (1 - \rho_\eta)\beta + \rho_\eta/2. \tag{44}$$

Here D° gives equal probability to the prices in D, and β^α is the acceptance probability at every consumer information set through date T. After date T, use the prescribed strategies g exactly. Perturb the initial valuation distribution to

$$F^\alpha = (1 - \rho_\eta)F + \rho_\eta V^\circ, \tag{45}$$

where V° gives equal probability to the valuations in V. Random choices are independent conditional on the relevant information. The information design is unchanged.

[26] Myerson and Reny explicitly contemplate countably many dates. Here messages, prices, and purchase decisions occur in successive stages. We represent the platform's random move by its finite message distribution. Each message in its support already has positive probability, so only the initial valuation law needs perturbing.

Thus, through date T, the consumer perturbation is exactly (39). The seller perturbation transfers probability ρ_η from the prescribed price in (38) to D°. Similarly, the valuation perturbation assigns total probability ρ_η to draws from V°. Since ρ_η vanishes faster than every power of η, these additions preserve the limiting beliefs established in Proposition 8.

The cutoff T applies only to mistakes. The game continues forever, and consumers can consider deviations of any duration. At later indices, the cutoff increases, so every fixed date is eventually included in the perturbation.

Proposition 9 (Perfect conditional ϵ-equilibrium of the construction)**.** *For every finite $\bar{p} > 1$, the strategy profile in Theorem 2, restricted to prices in $[0, \bar{p}]$, is a perfect conditional ϵ-equilibrium under Definition 4 for every $\epsilon > 0$. Its ϵ-test nets can be chosen so that, at every fixed date and permitted message, the seller's joint posterior converges to the belief in Proposition 8.*

Proof. Step 1: Control sellers' beliefs at every date. Let $a = \eta + \eta^2 + \rho_\eta$ be the total seller perturbation probability. Compare perturbed play with prescribed play by using the same valuation and random draws whenever possible. The probability that the two plays ever differ is at most

$$e_T = \rho_\eta + T(a + \rho_\eta) :$$

the valuation differs with probability at most ρ_η, and departures from prescribed behavior can occur only during the first T dates.

At every date, each message with positive equilibrium probability has probability at least $c > 0$, where c is independent of the date. This follows from (40): these probabilities are drawn from a finite list of positive prior interval probabilities. Consequently, when $e_T < c/2$, the largest difference between the perturbed and equilibrium conditional probabilities of any event in (v, h^{t-1}) is at most $4e_T/c$, at all such messages and dates.

The remaining messages are off-path occurrences of $m_{r,n+1}$. Under these perturbed strategies, none can occur after date $T + 1$. Such a message requires a rejection in W_n by

the consumer who entered at $m_{p,n}$. After date T, that incumbent can face only b_0 or $r(b_{n-1})$, both strictly below W_n. These offers cannot generate another $m_{r,n+1}$.

There are therefore only finitely many remaining cases to check: message $m_{r,n+1}$ observed at date t, with $n+2 \leq t \leq T+1$. Without the additional probabilities assigned to the prices in D and valuations in V, play through date T is exactly the common perturbation in Online Appendix OB.4.a. We can therefore use Proposition 8 directly: for fixed T, the histories in (41) have probability bounded below by a positive constant times η^T, and their conditional law approaches the stated limiting belief. Assigning these additional probabilities changes the probability of any event through the message at date $T+1$ by at most $(1+T)\rho_\eta$. The effect after conditioning vanishes because

$$\frac{(1+T)\rho_\eta}{\eta^T} = \frac{(1+T)e^{-1/\eta}}{\eta^T} \longrightarrow 0.$$

Thus there is a bound $d_T(\eta) \to 0$, for fixed T, on the posterior error at every positive-probability seller message, uniformly over dates and over the finite sets D, V. Here error means the largest difference in the probabilities assigned to any event in the valuation and past history. The comparison belief is the one in Proposition 8.

Step 2: Verify conditional optimality for sellers. At each message, the prescribed price maximizes profit under the limiting joint belief and the prescribed consumer rule. A posterior error of at most $d_T(\eta)$ changes profit at any price by at most $\bar{p}\, d_T(\eta)$. Consumer mistakes change it by at most $\bar{p}\rho_\eta/2$. Finally, the seller's own price randomization costs at most $\bar{p}a$. Hence the gain from any alternative price, including a mixed price choice, is at most

$$\bar{p}\Big[2d_T(\eta) + \rho_\eta + a\Big]. \tag{46}$$

The same bound applies when conditioning on any collection of messages, by averaging over the messages in that collection.

Step 3: Verify conditional optimality for consumers. Fix any valuation, observed history, and current offer at which a consumer acts. The proof of Theorem 2 verifies optimality of the prescribed continuation at every such history, including valuations outside the public belief's support.

Set $M = 1 + \bar{p}$. Under any fixed consumer continuation strategy, compare play against prescribed sellers with play against perturbed sellers, using identical draws until the first different price. If that first difference is j periods ahead, the difference in the consumer's continuation payoff is at most $M\delta_B^j$. Its probability is at most a. Summing over $j \geq 0$ bounds the change in expected payoff by $Ma/(1-\delta_B)$. The consumer's own mistakes add at most $M\rho_\eta/(1-\delta_B)$ to this bound.

Comparing an arbitrary continuation deviation with the prescribed perturbed strategy therefore gives the bound

$$\frac{M(2a+\rho_\eta)}{1-\delta_B}. \tag{47}$$

We have measured payoffs from the current date. Discounting back to the consumer's entry date can only reduce the bound. It holds uniformly over valuations, histories, and continuation deviations, and hence also conditional on any positive-probability event in the consumer's information.

Step 4: Choose the perturbations and verify admissibility. For each (T, D, V, k), choose $\eta > 0$ small enough that $d_T(\eta) \leq 1/k$, both (46) and (47) are at most $1/k$, and

$$\eta \leq \frac{1}{k(T+1)}.$$

Such a choice exists by Step 1. Every resulting profile is a conditional $1/k$-equilibrium. The largest change in any seller action probability is at most a, in any consumer action probability at most $\rho_\eta/2$, and in any valuation probability at most ρ_η. All these bounds vanish along the net.

Eventually T exceeds any fixed decision date. Both consumer actions then have positive probability, and each fixed seller price eventually belongs to D and also has positive probability. Each fixed valuation likewise eventually belongs to V. The unchanged message distributions already give positive probability to each message in their support. Thus the net is admissible.

For any $\epsilon > 0$, the part of the net with $k \geq \lceil 1/\epsilon \rceil$ is an ϵ-test net for g. Fix a message and a date at which it is permitted. Eventually T reaches that date and the posterior error is at most $1/k$, proving the claimed belief convergence. □

Implication for the outcome. Under this extension of Myerson and Reny's definition, the prescribed outcome distribution is a limit of perfect conditional ϵ-equilibrium distributions: keep g fixed as $\epsilon \downarrow 0$ (Myerson and Reny, 2020, Section 6.2). It therefore meets their equilibrium-distribution criterion.