Price Transparency and Market Screening

Ayça Kaya* and Santanu Roy†

July 22, 2020

Abstract

We consider repeated trading by sellers with persistent private information in dynamic lemons markets. We compare the outcomes of a *transparent market* where past trading prices are public to those of an *opaque* market, where they are private. We characterize the upper bound of trading surplus in an opaque market and construct a class of equilibria in a transparent market that improves upon this bound. We conclude that price transparency is beneficial in a repeated trading environment. The advantage of price transparency is indirect and operates through the strategic tools it provides the sellers of high quality to sustain high payoffs.

Keywords: Repeated sales, adverse selection, lemons market, price transparency.

JEL codes: D82, C73 , D61

*Department of Economics, University of Miami, Coral Gables, FL 33146. E-mail: a.kaya@miami.edu.
†Department of Economics, Southern Methodist University, Dallas, Texas 75275. E-mail: sroy@smu.edu

1 Introduction

Transparency about various aspects of past trades can play an important role in market outcomes and is an important aspect of the design of such markets for all stakeholders. Price information related to past transactions is somewhat more difficult to observe than the volume traded. An important economic question is therefore the impact of requiring or facilitating price transparency. We study this question in the context of a market where sellers with persistent private information repeatedly participate.

Such environments are common. For instance in many service industries, labor markets and many B-2-B transactions, sellers typically have the capacity to supply goods repeatedly over time and have private information about quality and other vertical attributes of their products that buyers care about. Many of these attributes are determined by factors that are relatively inflexible in the short run such as production technology, product development and the established supply chain. In consequence, quality attributes of products supplied over time as well as the private information of sellers about these attributes can be somewhat persistent.

We study the impact of the observability of realized trading prices on market outcomes. The role of such transparency is not only a theoretical curiosity: for instance in labor markets, the merits of banning potential employers (i.e. buyers) from asking potential employees (i.e. sellers) about their past wages (i.e. past transaction prices) is a current and hotly debated policy issue.[1] Further, in some US states, past transaction prices are legally considered trade secrets, and the parties cannot be required to disclose them.[2] Our study contributes to the understanding of issues surrounding such legislation.

Our model features a seller who can produce one unit per period and who receives take-it-or-leave it price offers from a sequence of short-lived buyers. The quality of the seller's output which determines both its use value and its production cost, is persistent and is the seller's private information. We consider two specifications: an *opaque market* where the buyers observe the volume of the seller's past trades but no price information versus a *transparent market* where buyers, in addition, observe the past transaction prices. We also compare the outcomes to a no-information benchmark, where the buyers observe neither the past trading volumes nor past transaction prices.

We find that a transparent market is able to quickly and fully learn the seller's quality, while full learning never takes place in an opaque market. Price observability facilitates full learning not directly by increasing the information flow but indirectly by providing strategic tools to the high quality seller to extract rents. Importantly, full learning in the transparent market is also associated with efficiency gains relative to the opaque markets.

[1]For instance, in 2019 the US House of representatives passed a "Paycheck Fairness Act," which among other things, would prohibit potential employers from asking about salary history, while at the same time protecting the rights of the workers to voluntarily reveal it. The passage of the bill in the Republican-led Senate is deemed unlikely.

[2]See for example, `https://fas.org/sgp/crs/secrecy/R43714.pdf`

To understand these results, note that when the history of trades is observable, regardless of the observability of prices, high quality can trade. The mechanism that allows this is familiar: because of the cost differential, lowered frequency of trade becomes a credible indication of high quality. The high quality's path must feature sufficiently infrequent trading so that the low quality seller prefers to reveal himself rather than mimic this path. Thus, regardless of price observability, high quality cannot trade efficiently. Price observability affects the trading rate of low quality: that an opaque market cannot fully screen the seller implies that the low quality must follow the high quality's inefficient path with some probability. In contrast, a transparent market can fully screen the seller, and thus, high quality can trade slowly while the low quality trades efficiently. The key property of a transparent market that allows full screening is that in such a market, the seller can lose a lot by trading at discount prices. This gives him the commitment power to resist undercutting, and in turn allows him to sustain positive payoffs. The threat of losing this positive payoff allows the high quality seller to trade slowly even when the market perceives a very high quality and would be willing to pay high prices.[3]

Our conclusions are based on two sets of results: characterization of the limits of market screening and the gains from trade in an opaque market, and the construction of a class of equilibria in the transparent market that achieve full screening. We further show that, as the seller becomes patient, the outcomes of these full screening equilibria approximate the maximum surplus achievable in any fully separating allocation.

The role of price transparency has been previously studied in single-sale environments where trade may occur in only one period. Kaya and Liu (2015) study this problem in a Coasian environment. (Hörner and Vieille, 2009a,b) analyze this question for a dynamic lemons market where, as in our model, a privately informed long lived seller faces a short lived buyer in every period. Fuchs et al. (2016) analyze it in a similar model but with multiple buyers in every period. In these models, price transparency necessarily refers to the observability of rejected offers. By studying a repeated sale environment, we are able to address the arguably more relevant question regarding the observability of past transaction prices.

In contrast to our findings, this literature concludes that with price transparency trading slows down, and the informed side is worse off.[4] The discrepancy is understood by noting that in their setting, the interaction ends as soon as a single trade is completed, thus, the key advantage of price observability that emerges in our setting—i.e. that it allows the high quality seller to sustain high payoffs,—is moot. Further, observability of rejected prices versus transaction prices impacts the market participants' incentives differently: When the observed rejection of a decent offer is interpreted as good news by the market, even the

[3]In contrast, in a single-sale environment, such threats are ineffective since the seller leaves the market once a transaction takes place.

[4]Also see Bergemann and Hörner (2018) for a similar conclusion in an environment where multiple bidders repeatedly participate in first price auctions for identical items.

low quality seller finds it worthwhile to do so, and the market's ability to screen is severely limited. Such signaling motives are absent when only the accepted offers are observed.[5]

The role of price transparency has been studied in other contexts. For instance, Chaves (2020) asks how the transparency of price offers in ongoing negotiations impact the entry decisions of other potential trading partners. Cullen and Pakzad-Hurson (2020) study the role of wage transparency when a single employer is bargaining with multiple agents. Bochet and Siegenthaler (forthcoming) provides experimental evidence on the interaction of price transparency and buyer competition.

Other forms of transparency impact efficiency of markets with adverse selection. Kim (2017) considers the availability of information about the seller's time on the market, Hwang and Li (2017) considers the observability of dynamically arriving outside options of long-run players, Asriyan et al. (2017) and Huangfu and Liu (2019) consider the transparency of parallel bargaining over assets with correlated values, Kaya and Roy (2020) considers variation in the length of trading records but does not address the affect of price transparency.

In what follows, Section 2 presents the model, Section 3 and Section 4 analyze opaque and transparent markets, respectively. Section 5 interprets the results. Section 6 collects formal proofs.

2 Model

A long-lived seller can produce one unit of output every period. Time is discrete and horizon is infinite, so that the interaction takes place over time periods $t = 1, 2, \cdots$. Each period, the seller meets a potential trading partner (a buyer) with unit demand who makes a take-it-or-leave-it price offer. Seller either accepts the buyer's offer and trades one unit at that price or rejects it. Regardless, buyer leaves the game, and the seller moves to the next period.

Seller's type $s \in \{L, H\}$ determines both the use value of his output and the cost of production. If in a given period trade takes place at price P, the type-s seller's payoff in that period is $P - c_s$ and her trading partner's payoff is $v_s - P$. Regardless of seller's type, gains from trade is strictly positive:

$$v_s - c_s > 0, s \in \{L, H\}.$$

The instantaneous payoff for any party who does not trade is normalized to 0. The seller maximizes the expected discounted sum of his future payoffs using discount factor $\delta \in [0, 1]$.

Seller's type is his private information. All buyers hold a common prior that assigns probability μ_0 to type $s = H$. Let μ^* be defined by $\mu^* v_H + (1 - \mu^*) v_L = c_H$. We maintain the following assumptions.

[5] Even though we do not consider this case, it is straightforward to construct a low trade equilibrium in a repeated sale environment when failed offers are observed.

Assumption 1 *Prior expected valuation is lower than c_H, i.e. $\mu^* > \mu_0$.*

Assumption 2 *Seller is sufficiently patient:*

$$\frac{v_L - c_L}{c_H - c_L} < \delta^2.$$

Public histories A public history at time t in an **opaque market** contains information about whether trade took place at each $t' < t$, and thus is an element of 2^{t-1}. A public history in a **transparent market**, in addition to the same trading information, includes the realized trading prices in periods where trade took place.

Equilibrium We consider perfect Bayesian equilibria. Let h be an arbitrary public history, at the beginning of a period, before the buyer offer is realized. Let $\mathcal{H}$ represent the set of public histories. A perfect Bayesian equilibrium consists of a strategy profile and a belief system. A buyer strategy maps $\mathcal{H}$ into an offer distribution, and a seller strategy assigns $\mathcal{H} \times \{L, H\}$ into an acceptance rule. A belief system is a map $\mu : \mathcal{H} \to [0,1]$ representing the probability that the public belief assigns to high quality. A strategy profile and a belief system forms a perfect Bayesian equilibrium if beliefs are derived using Bayes rule from public histories and strategies of others whenever possible, and the strategies maximize each player's payoff based on their beliefs and the strategies of others.

Fixing an equilibrium, throughout we let $V_s(h)$, $s \in \{L, H\}$, represent the type-s seller's continuation payoff at history h. We express all payoffs and trading volumes in average per-period terms.

2.1 No-information benchmark

If buyers observe no information about trading history, the seller acts myopically, as his continuation payoff cannot depend on his actions. Thus, the outcome is the period-by-period repetition of the static market outcome: the low quality trades efficiently and the high quality never trades. Consequently, the per-period gains from trade is $(1-\mu_0)(v_L-c_L)$.

3 Opaque markets

This section characterizes the limits of market screening in an opaque market and highlights the impact of the inability to screen on gains from trade. We first establish that the high quality seller cannot sustain positive payoffs.

Lemma 1 *In any equilibrium of the opaque market, at any history h, $V_H(h) = 0$.*

The mechanism that drives Lemma 1 is familiar: seller's continuation payoff is independent of the trading price, thus he is reluctant to turn away a buyer who only slightly undercuts what is anticipated. Thus, the equilibrium prices are driven all the way down to cost.

Observe that in an opaque market, the seller necessarily uses a reservation price strategy. Let $P_s(h)$ represent type-s seller's reservation price at history h. Lemma 1 implies that at any h, $P_H(h) = c_H$. It is worth noting that the so-called skimming property, which in this context would imply $P_L(h) < c_H$ at each h, does not necessarily hold in a repeated sale model. This complicates the formal arguments, which are relegated to Section 6. Nevertheless, the main insights can be gleaned by focusing on equilibria where it holds. Note that, the skimming property would imply that (i) the high quality cannot trade unless $\mu(h) \geq \mu^*$, and (ii) a period of no trade necessarily (weakly) increases belief. In this section, we base our intuitive discussions on these properties.

Naturally, if the belief reaches 1, each arriving buyer offers c_H, no further learning occurs, and $V_L(h) = c_H - c_L$. On the other extreme, if the belief reaches 0, the low quality seller's payoff cannot exceed $v_L - c_L$, since otherwise some buyers must make losses. Then, by Assumption 2, if at some history, acceptance were to reveal low quality and rejection were to reveal high quality, the low type's reservation price would exceed v_L, which is incompatible with separation.

This inability of the market to fully screen stems from the large gap between the payoffs that the low quality seller commands if he were to reveal his type versus if he were to pass for a high quality seller. The next two lemmas establish that there is a similarly large discrete gap in the low quality seller's continuation payoffs if the belief is strictly below versus strictly above μ^*.

Lemma 2 *If $\mu(h) > \mu^*$, then $V_L(h) \geq \delta(c_H - c_L)$.*

When $\mu(h) > \mu^*$, the buyer would never make a losing offer. Thus, he either trades with probability 1, in which case the belief remains high, or he targets the low type, in which case, rejection reveals high quality leading to the lower bound in Lemma 2.

Lemma 3 *If $\mu(h) < \mu^*$, then $V_L(h) \leq v_L - c_L$.*

If $\mu(h) < \mu^*$, the first time trade takes place in the continuation path, acceptance necessarily reveals low quality.[6] This is because an offer of c_H at such beliefs would make a loss to the buyer. The upper bound in Lemma 3 is the maximum the low quality seller can get by such trade that reveals his type.

Let (h, A) and (h, R) represent the histories obtained by augmenting history h with, respectively, a period of trade and a period of no trade. Suppose at some history h with $\mu(h) < \mu^*$, we have $\mu(h, A) < \mu^* < \mu(h, R)$. Then, by Assumption 2, and Lemmas 2 and

[6]Strictly speaking, this is not necessarily true when skimming property fails. The formal proof of Lemma 3 deals with that case.

3, $P_L(h) > v_L$, which is incompatible with low quality trading with higher probability at this history. This observation, combined with the fact that $\mu_0 < \mu^*$, leads to the first main conclusion of this section.

Proposition 1 *In an opaque market, in any equilibrium, at any equilibrium path history h,*

$$\mu(h) \leq \mu^*.$$

It is worth emphasizing that Lemma 2, and not Lemma 3, is the manifestation of the short-coming of an opaque market relative to a transparent market, leading to Proposition 1. As we show in Section 4, in a transparent market, even when the buyers are very optimistic, there are equilibria that deliver the low quality seller a low payoff of mimicking, and thus finer screening is possible.

Next, we show that the pooling required to keep belief below μ^* swamps all the gains generated by the trade of high quality, so that the total gains from trade never exceeds that in the no-information benchmark.[7]

Proposition 2 *The expected gains from trade in an opaque market are no larger than*

$$(1 - \mu_0)(v_L - c_L).$$

This result is understood by noting that whenever a buyer purchases high quality, his payoff is 0. Thus, all gains from trade of high quality is channeled towards the cross-subsidization of the low quality's (inefficient) trading at high prices.

Off-path beliefs The upper bound in Proposition 2 is attained when the low quality seller extracts all surplus upon revealing himself, i.e. receives a payoff of $v_L - c_L$. This allows high quality to trade at a relatively high frequency without inducing mimicry. In particular, if Q is the average trading frequency of high quality, the low quality seller's payoff from mimicking is $Q(c_H - c_L)$, and thus

$$Q \leq Q^* \equiv \frac{v_L - c_L}{c_H - c_L}. \tag{1}$$

If the low quality seller does not expect to receive much of the surplus after revealing his type, high quality's trading must be further slowed down. A commonly adapted restriction on off-path beliefs is that once the belief becomes degenerate it must stay that way. Under that restriction, the low quality seller can extract *no* surplus upon revealing himself. In that case, the frequency of high quality's trading cannot exceed $(1-\delta)(v_L - c_L)/(c_H - c_L)$, which approaches 0 as $\delta \to 1$.

[7] For the single sale model, Hörner and Vieille (2009b) establish an analogous upper bound on social surplus realized by an opaque market. The characterization of social surplus for an opaque market in our repeated trading model is complicated by multiplicity of equilibrium outcomes (that differ in the low quality type's payoff after revelation), the inability to rely on the skimming property and the need to consider continuation equilibria even after trade takes place.

4 Transparent markets

Recall that the root cause of the need for (inefficient) pooling in an opaque market is the high quality seller's inability to sustain positive payoffs (Lemma 1), and in turn his inability to resist prices that exceed his cost, which ties down the outcome once the belief is above μ^* (Lemma 2). In a transparent market, the observability of trading prices gives the seller the power to commit to rejecting lower offers, which allows him to sustain a positive payoff. In turn, when the seller's equilibrium payoff is positive, the threat of switching to a low-payoff equilibrium provides the seller sufficient incentives to reject high prices even when buyers are willing to offer them.

In this section we demonstrate that there exists a full-screening equilibrium where after one period of trading the seller's type is revealed, and thus, unlike in an opaque market, it is possible for the low quality seller to trade efficiently. Note that, as in the opaque market, once low quality is revealed, the seller's payoff cannot exceed $v_L - c_L$. This bounds the frequency at which the high quality can trade to be below Q^* defined in (1).

The next proposition, establishes the existence of full screening equilibria in which the low quality trades efficiently. Further, it shows that full screening equilibria can be constructed so that high quality's trading approximates its upper bound Q^* as the seller becomes patient.

Proposition 3 *In a transparent market, there exists a full screening equilibrium in which the low quality trades efficiently whenever δ satisfies*

$$1 - \delta < \frac{v_L - c_L}{v_H - c_H}. \tag{2}$$

Further, for any $\varepsilon > 0$, there exists $\overline{\delta}(\varepsilon)$ such that whenever $\delta \geq \overline{\delta}(\varepsilon)$, there exists an equilibrium in which low quality trades efficiently and the high quality's average trading frequency is no less than $Q^* - \varepsilon$.

To prove Proposition 3 we construct an equilibrium in which, the low quality trades at price v_L with probability 1 in each period, starting in the first. The high quality trades at average frequency $\overline{Q}$ and at a constant price $\overline{P}$ such that

$$\overline{Q}(\overline{P} - c_L) = v_L - c_L, \tag{3}$$

$$\overline{Q}(\overline{P} - c_H) \geq (1 - \delta)(v_H - c_H). \tag{4}$$

(2) guarantees that there exists $\overline{Q}$ and $\overline{P} < v_H$ that simultaneously satisfy (3) and (4). For high quality's average trading rate to be slowed down to $\overline{Q} < 1$, there has to be some histories at which he does not trade, even though market correctly perceives the quality of his product. If the seller deviates and trades in a period where he is not supposed to or if he trades at a price lower than $\overline{P}$, the play switches to a "punishment equilibrium," in which the high quality seller's payoff is 0.

Constraint (3) guarantees that the low quality seller is just willing to reveal himself instead of mimicking the high quality's path. Constraint (4) is precisely the condition that allows for high quality's trading to slow down even when the market is very optimistic. In the equilibria we construct, the left-hand-side of (4) is the lowest continuation payoff that the high quality receives unless punishment was previously triggered. Thus (4) is sufficient for him to reject any price below v_H if acceptance triggers punishment. In addition, (4) also guarantees that the high quality seller is willing to reject any offer below $\overline{P}$.[8]

The proof of Proposition 3 constructs buyer offer strategies that deliver the frequency $\overline{Q}$ of trading. These offer strategies are necessarily pure: at any history where the high quality is supposed to trade with positive probability, the buyer's payoff is strictly positive since $\overline{P} < v_H$ and the buyer's belief is 1. Thus, offering $\overline{P}$ is uniquely optimal. We construct pure offer strategies so that the trading path cycles between k consecutive periods of no trade, followed by m consecutive periods of trade so that the average frequency of trade is

$$Q_{m,k}(\delta) \equiv \frac{\delta^k + \cdots + \delta^{k+m}}{1 + \delta + \cdots + \delta^{k+m}},$$

for some m, k. Focusing on this class places another constraint on $\overline{Q}$:

$$\overline{Q} \in \{Q_{m,k}(\delta) \mid m, k \in \mathbb{N}\}. \tag{5}$$

It is easy to show that this set is dense in $[0,1]$, so that whenever there exists $(\overline{Q}, \overline{P})$ that satisfy (3) and (4), it is possible to find one that also satisfies (5).[9]

Finally, to understand the convergence result stated in Proposition 3, first note that (3) induces a trade-off between $\overline{P}$ and $\overline{Q}$. For the high quality to trade more frequently, the trading price must be lower. Any such adjustment that keeps the left-hand-side of (3) constant necessarily reduces the left-hand-side of (4). In this way, (4) constrains the trading frequency of the high quality. As δ grows, (4) is relaxed, and smaller $\overline{P}$ and thus larger $\overline{Q}$ become feasible, approaching Q^*.

Off-path beliefs On the path of any equilibrium covered by Proposition 3, the belief about the seller's quality becomes degenerate immediately after the first round of trading. At the same time, off-path punishments require beliefs to switch away from such degenerate support. We note that in these equilibria, the low quality seller is indifferent between immediately revealing himself versus forever pooling with the high quality. Further, if he pools with the high quality with sufficiently small probability, the structure of the equilibrium remains unchanged, since the price offers here are determined by the high quality seller's willingness to accept rather than the market's perception of quality. Thus, switching away from degenerate beliefs is not essential for the construction.

[8]Strictly speaking, rejecting a price less than $\overline{P}$ moves the game to an off-equilibrium history, as it occurs when trade is supposed to take place. We construct off-equilibrium strategies so that this is still sufficient.

[9]Note that whenever there exists a solution to (3) and (4), there exists an open set of such solutions. This is because reducing $\overline{Q}$ and increasing $\overline{P}$ in a manner to satisfy (3) relaxes (4).

5 Conclusion

Both transparent and opaque markets admit multiple equilibria. In an opaque market, in any equilibrium that attains the highest possible gains from trade, the low quality seller captures all surplus.[10] By comparing such equilibria to those characterized in Proposition 3, we conclude that price transparency benefits the sellers of high quality as well as buyers, without reducing the low quality seller's payoff. Thus, price transparency leads to a Pareto improvement. This is in sharp contrast to the conclusions derived in single-sale environments where price transparency is often detrimental to trade frequency and surplus.

Another novel insight that the repeated sale model reveals is with respect to the impact of seller patience. Insights based on single-sale models suggest that trading slows down as the long-run players get more patient.[11]. In contrast, our analysis reveals that, in a transparent market, the high quality can trade more frequently as the seller becomes patient. These observations suggest that many conclusions based on the insights of single sale models are unlikely to generalize to environments where sellers trade repeatedly. Therefore, further study of the "repeated trading model" is likely to generate a rich set of novel insights and policy implications.

6 Proofs

Proof of Lemma 1 . Let $\bar{V}_H = sup\{V_H(h)|h \in \mathcal{H}\}$. Fix $\varepsilon_1 > 0$ small enough so that $\delta\bar{V}_H < \bar{V}_H - \varepsilon_1$ and let h^* be such that $V_H(h^*) > \bar{V}_H - \varepsilon_1$. High quality must trade with positive probability at h^* because otherwise, $V_H(h^*) = \delta V_H(h^*, R) \leq \delta\bar{V}_H < \bar{V}_H - \varepsilon_1$. Let P^* be the supremum of the support of the buyer's price offer at h^*. Then,

$$(P^* - c_H)(1 - \delta) + \delta V_H(h^*, A) \geq V_H(h^*) > \delta\bar{V}_H.$$

Consider an offer $P^* - \varepsilon_2$ at h^*. When ε_2 is sufficiently small, high quality seller must accept this with probability 1, because for such ε_2,

$$(P^* - \varepsilon_2 - c_H)(1 - \delta) + \delta V_H(h^*, A) > \delta\bar{V}_H \geq \delta V_H(h^*, R).$$

Thus, the buyer has a profitable deviation. This establishes that $V_H(h) = 0$. ■

Lemma 4 records useful preliminary observations.

Lemma 4 *Fix any on or off-path history h.*

[10] Even though not included in this paper, it is easy to construct these equilibria for the opaque market. If the low quality receives less surplus, the trading of high quality must be further slowed, while the probability that the low quality must pool with the high is fixed.

[11] See for instance Janssen and Roy (2002), Deneckere and Liang (2006) as well as Hörner and Vieille (2009a).

1. *If $P_L(h) < v_L$, low quality trades with probability 1.*

2. *Trade takes place with probability 1 at price c_H, if $\mu(h) > \mu^*$ and $\frac{c_H - P_L(h)}{v_H - c_H} < \frac{\mu(h)}{1-\mu(h)}$.*

3. *$P_H(h) = c_H \geq P_L(h)$.*

Proof of Lemma 4.

1. If $P_L(h) < v_L$, the buyer's payoff is positive, so losing offers are not optimal. The claim follows if the offer is c_H. If the offer is $P_L(h)$, it must be accepted with probability 1, because otherwise the buyer would benefit from slightly increasing it.

2. If $\mu(h) > \mu^*$, the buyer's payoff is positive, so losing offers are not optimal. By standard arguments, the support of the buyer's offer is contained in $\{c_H, P_L(h^*)\}$. The inequality compares the buyer's payoff from either offer.

3. That $P_H(h) = c_H$ immediately follows by Lemma 1. If $P_L(h) > c_H$, low quality trades with probability 0 while high quality trades with probability 1, implying that $P_L(h) \leq c_L$, a contradiction.

■

Proof of Lemma 2. Assume $\mu(h) > \mu^*$. Buyer's payoff is strictly positive, thus he makes no losing offers. If trade takes place at $P < c_H$ with positive probability, then $\mu(h, R) = 1$, thus $V_L(h) \geq \delta(c_H - c_L)$. Let

$$\underline{V}_L = \inf\{V_L(h) | \mu(h) > \mu^* \text{and trade takes place only at price } c_H\}.$$

Let h^* be a history with $\mu(h^*) > \mu^*$ and at which trade takes place only at c_H, which also satisfies $V_L(h^*) < \underline{V}_L + (1-\delta)^2(c_H - c_L)$. Here, $V_L(h^*) = (1-\delta)(c_H - c_L) + \delta V_L(h^*, A)$, because, the buyer never makes a losing offer, and thus offers c_H with probability 1 and $P_L(h^*) \leq c_H$, thus accepting c_H is an optimal action for low quality seller.

Let γ_s be the probability with which type-s seller accepts c_H. Then, by increasing the offer by a small amount the buyer can increase his payoff by approximately

$$\mu(h^*)(1-\gamma_H)(v_H - c_H) + (1-\mu(h^*))(1-\gamma_L)(v_L - c_H),$$

which is non-positive if and only if $\mu(h^*, R) \leq \mu^*$. This in turn implies that $\mu(h^*, A) > \mu^*$. If at (h^*, A), trade takes place only at c_H, $V_L(h^*, A) \geq \underline{V}_L$. Otherwise, $V_L(h^*, A) \geq \delta(c_H - c_L)$. The claim follows in both cases. ■

Proof of Lemma 3. Consider a (possibly off-equilibrium) continuation path, after equilibrium path history h, along which the low type always rejects his reservation price when offered. Note that $V_L(h)$ can be calculated along this path. Let h_1 be the first continuation history along this path where equilibrium probability of trade is positive and $P_L(h_1) < c_H$. Such h_1 exists because otherwise along this path the low quality never trades and thus $V_L(h, R) = 0$, which in turn implies that $P_L(h) \leq c_L$, a contradiction by Lemma 4. Further, note that $h_1 = (h, R \cdots, R)$. Along the path, at each interim history h', either the equilibrium probability of trade is 0, in which case the belief is not updated, or trade is supposed to take place at price c_H. In the latter case, for the buyer's payoff to be non-negative the expected valuation conditional on acceptance must be no less than c_H. That is, $\mu(h', A) \geq \mu^*$. Thus, if $\mu(h') < \mu^*$, we have $\mu(h') > \mu(h', R)$. Since $\mu(h) < \mu^*$, we conclude that $\mu(h_1) \leq \mu(h) < \mu^*$. Since also $P_L(h_1) < c_H$, at h_1, the buyer never offers c_H. Thus high quality does not trade. Consequently, $\mu(h_1, A) = 0$, and thus $V_L(h_1, A) \leq v_L - c_L$. Further, $P_L(h_1) \leq v_L$ because otherwise the buyer's payoff is negative. Thus, $V_L(h) \leq V_L(h_1) \leq (1-\delta)(v_L - c_L) + \delta(v_L - c_L) = v_L - c_L$, where the first inequality follows because there is no trade between h and h_1 along this path. ■

Proof of Proposition 1. The proof is through the following claims: if trade takes place with positive probability at h, then

1. if $\mu(h) < \mu^*$ then either $0 = \mu(h, A) < \mu(h, R) \leq \mu^*$ or $\mu(h, R) < \mu(h, A) = \mu^*$.

2. if $\mu(h) = \mu^*$, then $\mu(h, A) = \mu(h, R) = \mu^*$.

For the first claim, if $P_L(h) < c_H$, then high type does not trade at h because otherwise the seller makes a loss. Then, $\mu(h, A) = 0$. If $\mu(h, R) > \mu^*$, then $P_L(h) > v_L$, and the buyer makes a loss. Thus, $\mu(h, R) \leq \mu^*$. If $P_L(h) = c_H$, then trade takes place at price c_H. Then, $\mu(h, A) \geq \mu^*$, because otherwise the buyer makes a loss. If $\mu(h, A) > \mu^*$, then $\mu(h, R) < \mu^*$, and thus $P_L(h) < v_L$. By Lemma 4, low quality trades for sure, implying $\mu(h, A) < \mu^*$, a contradiction. Thus, $\mu(h, A) = \mu^*$.

For the second claim, if $\mu(h, A) < \mu^* < \mu(h, R)$, then $P_L(h) > v_L$, thus trade takes place only at c_H. Then, $\mu(h, R) \leq \mu^*$, because otherwise the buyer has a profitable deviation to increase offer slightly above c_H, a contradiction. If $\mu(h, A) > \mu^* > \mu(h, R)$, then $P_L(h) < c_L$, low quality trades with probability 1, and thus $\mu(h, A) \leq \mu^*$, a contradiction. ■

Proof of Proposition 2. Fix an equilibrium. Let $h^t_\emptyset$ represent the t-length history that features no trading. Also let $\hat{Q}^s(h^t)$ represent the expected discounted sum of trades conditional on type $s \in \{L, H\}$ during the continuation equilibrium following h^t. Finally, let $\gamma_s(h^t)$ be the probability with which the seller type $s \in \{L, H\}$ visits history h^t.

Claim 1 *For any t, if $(h^{t-1}_\emptyset, A)$ is on path, then $\mu(h^{t-1}_\emptyset, A) \in \{0, \mu^*\}$.*

Proof of Claim 1. If $\mu(h^{t-1}_\emptyset) = \mu^*$, the result follows because belief remains constant thereafter, as established in the proof of Proposition 1. If $\mu(h^{t-1}_\emptyset) < \mu^*$, the result follows by the first claim in the proof of Proposition 1. ■

Define $T_{\mu^*} = \{t|\mu(h_\emptyset^{t-1}, A) = \mu^*\}$ and $T_0 = \{t|\mu(h_\emptyset^{t-1}, A) = 0\}$. Let $\overline{Q}_s$ be the average trading frequency of seller type $s \in \{L, H\}$. Then,

$$\overline{Q}_H \equiv \sum_{t \in T_{\mu^*}} \gamma_H(h_\emptyset^{t-1}, A) \left[(1-\delta)\delta^{t-1} + \delta^t \hat{Q}^H(h_\emptyset^t, A)\right],$$

$$\overline{Q}_L \equiv \sum_{t \in T_{\mu^*}} \gamma_L(h_\emptyset^{t-1}, A) \left[(1-\delta)\delta^{t-1} + \delta^t \hat{Q}^L(h_\emptyset^t, A)\right] + \sum_{t \in T_0} \gamma_L(h_\emptyset^{t-1}, A) \left[(1-\delta)\delta^{t-1} + \delta^t \hat{Q}^L(h_\emptyset^t, A)\right].$$

Note that, whenever $(h_\emptyset^{t-1}, A)$ is on path,

$$\gamma_H(h_\emptyset^{t-1}, A) = \begin{cases} 0 & \text{if } t \in T_0 \\ \gamma_L(h_\emptyset^{t-1}, A) \frac{1-\mu_0}{\mu_0} \frac{\mu^*}{1-\mu^*} & \text{if } t \in T_{\mu^*} \end{cases}.$$

Further,

$$\gamma_H(h_\emptyset^\infty) + \sum_{t \in T_{\mu^*} \cup T_0} \gamma_H(h_\emptyset^{t-1}, A) = \gamma_L(h_\emptyset^\infty) + \sum_{t \in T_{\mu^*} \cup T_0} \gamma_L(h_\emptyset^{t-1}, A) = 1$$

and

$$\gamma_H(h_\emptyset^\infty) \leq \gamma_L(h_\emptyset^\infty) \frac{1-\mu_0}{\mu_0} \frac{\mu^*}{1-\mu^*}.$$

The last inequality follows because $\gamma_s(h_\emptyset^t)$ is a monotone decreasing sequence in $[0, 1]$, and thus is convergent with limit $\gamma_s(h_\emptyset^\infty)$ and at each t, $\mu(h_\emptyset^t) \leq \mu^*$. Further, since no learning takes place once belief reaches μ^*, for each $t \in T_{\mu^*}$, $\hat{Q}^L(h_\emptyset^{t-1}, A) = \hat{Q}^H(h_\emptyset^{t-1}, A)$. It follows that

$$\begin{aligned}
\overline{Q}_L &\leq \frac{\mu_0}{1-\mu_0} \frac{1-\mu^*}{\mu^*} \overline{Q}_H + \sum_{t \in T_0} \gamma_L(h_\emptyset^{t-1}, A) \\
&= \frac{\mu_0}{1-\mu_0} \frac{1-\mu^*}{\mu^*} \overline{Q}_H + \left(1 - \sum_{t \in T_{\mu^*}} \gamma_L(h_\emptyset^{t-1}, A) - \gamma_L(h_\emptyset^\infty)\right) \\
&\leq \frac{\mu_0}{1-\mu_0} \frac{1-\mu^*}{\mu^*} \overline{Q}_H + \left(1 - \frac{\mu_0}{1-\mu_0} \frac{1-\mu^*}{\mu^*} \left(\underbrace{\sum_{t \in T_{\mu^*}} \gamma_H(h_\emptyset^{t-1}, A) + \gamma_H(h_\emptyset^\infty)}_{=1}\right)\right) \\
&= \frac{\mu_0}{1-\mu_0} \frac{1-\mu^*}{\mu^*} \overline{Q}_H + 1 - \frac{\mu_0}{1-\mu_0} \frac{1-\mu^*}{\mu^*}.
\end{aligned}$$

Since the low quality seller can always mimic the high quality, by Lemma 3, $v_L - c_L \geq V_L(h_0) \geq Q_H(c_H - c_L)$. Thus, the total gains from trade is less than or equal to

$$(v_L - c_L) \left[\frac{1-\mu_0}{\mu_0} \frac{\mu^*}{1-\mu^*} Q^* + \left(1 - \frac{1-\mu_0}{\mu_0} \frac{\mu^*}{1-\mu^*}\right)\right] + (c_H - c_L) Q^*,$$

where $Q^* = (v_L - c_L)/(c_H - c_L)$. The claim follows by substituting $\mu^*/(1-\mu^*) = (c_H - v_L)/(v_H - c_H)$. ■

Proof of Proposition 3. We first show that when (2) is satisfied, there exists $\overline{Q}, \overline{P}$ that solves (3) and (4). First, note that if Q_1, P_1 and Q_2, P_2 with $P_1 > P_2$ both satisfy (3) and Q_2, P_2 satisfies (4), then Q_1, P_1 also satisfy (4). Thus, a solution exists if and only if $(\overline{P}, \overline{Q}) = \left(v_H, \frac{v_L - c_L}{v_H - c_L}\right)$ satisfies (3) and (4). This is equivalent to

$$\frac{v_L - c_L}{v_H - c_L}(v_H - c_H) \geq (1-\delta)(v_H - c_H),$$

establishing the claim.

In this equilibrium, at certain histories "punishment is triggered." At these histories, belief updates to 0, and starting the next period the following equilibrium is played. This equilibrium also describes the low quality's equilibrium path.

0-belief equilibrium: When $\mu(h) = 0$, there exists an equilibrium where the buyers use the following offer strategies

- If trading price (since the punishment was triggered) was never less than v_L, or if since trading at a lower price, there were at least K consecutive periods of no trade, offer v_L, where K is large enough so that $1 - \delta + \delta^K < \delta$. This is possible as long as $\delta \geq 1/2$. That $\delta \geq 1/2$ is implied by (4), noting that $\overline{Q} < \delta$ by (3) and Assumption 2.

- Otherwise, offer c_L.

The low quality seller accepts offers weakly exceeding v_L, and rejects others. High quality seller accepts offers weakly exceeding c_H, and rejects others. The optimality of the strategies are trivially verified.

High quality's average trading frequency: Take any solution $\overline{Q}, \overline{P} < v_H$ of (3) and (4). For any $\varepsilon > 0$, one can choose m, k such that $Q_{m,k}(\delta) \in (\overline{Q} - \varepsilon, \overline{Q}]$. Define $P_{m,k}(\delta)$ by

$$Q_{m,k}(\delta)(P_{m,k}(\delta) - c_L) = v_L - c_L. \tag{6}$$

Choose ε small enough so that $P_{m,k}(\delta) < v_H$. Then,

$$Q_{m,k}(\delta)(P_{m,k}(\delta) - c_H) \geq (1-\delta)(v_H - c_H). \tag{7}$$

Equilibrium strategies At $t = 1$, buyers offer v_L. Low type seller accepts offers weakly exceeding v_L and rejects others. High type seller rejects all offers less than v_H. After acceptance of $p = v_L$, the belief updates to $\mu = 0$ and the low type's payoff is $v_L - c_L$. After acceptance of any offer $p \neq v_L$ at $t = 1$ punishment is triggered.

For $t \geq 2$, define $\tau_A(h)$ to be the length of the latest streak of periods in which trade took place. If the previous period featured no trade, $\tau_A(h) = 0$. Define $\tau_R(h)$ to be the length of the most recent streak of consecutive periods where no trade took place. Different from $\tau_A(\cdot)$, the streak need not be unbroken. Partition histories with $t \geq 2$ where punishment has not been previously triggered and the first period featured no trade into:

group (i) $\tau_A(h) > 0$ and last period's trading price is less than $P_{m,k}(\delta)$;

group (ii-$m'k'$) not in group (i), and $\tau_A(h) = m', \tau_R(h) = k'$.

Any history in group (i) or any history in group (ii-$m'k'$) with $m' > m$ or $m' > 0$ and $k' < k$ triggers punishment. In group (ii-$m'k'$),

- if $m' = 0$ and $k' < k$, or $m' = m$ and $k' \geq k$, buyers offer v_L, high quality seller rejects all offers less than v_H; low quality seller's strategy is characterized by a reservation price $P_L^{k'} \geq v_L$, which he rejects when offered.
- if $m > m' > 0$ and $k' \geq k$, buyers offer $P_{m,k}(\delta)$; high quality seller rejects all offers strictly less than $P_{m,k}(\delta)$ accepts all others; low quality seller accepts all offers strictly exceeding v_L, rejects all others.

Off-path beliefs: In all off-path histories in group (ii-$m'k'$) that don't trigger punishment, the buyer's belief is 1.

Verification: Buyers' strategies are optimal because they prescribe offering the lowest price accepted by the high quality seller at each history where that price is less than v_H. High quality seller's strategy at each history is optimal by (7). For the low quality, consider $m' = 0$ and $k' < k$, or $m' = m$ and $k' \geq k$. An acceptance triggers punishment regardless of price, while a rejection leads to a payoff no less than $Q_{k,m}(P_{k,m} - c_L)/\delta = (v_L - c_L)/\delta$, establishing the claim. Next consider $m > m' > 0$ and $k' \geq k$. Acceptance of $P < P_{m,k}$ triggers punishment regardless of P. Acceptance of $P \geq P_{m,k}$ leads to a payoff no less than $Q_{k,m}(P_{k,m} - c_L)/\delta = (v_L - c_L)/\delta$. Rejection leads to a payoff exactly equal to $Q_{k,m}(P_{k,m} - c_L)/\delta = (v_L - c_L)/\delta$, regardless of k'. The optimality of the low quality seller's strategy follows.

Next, let $(Q^*(\delta), P^*(\delta))$ be the solution of

$$Q^*(\delta)(P^*(\delta) - c_L) = v_L - c_L, \tag{8}$$

$$Q^*(\delta)(P^*(\delta) - c_H) = (1 - \delta)(v_H - c_H). \tag{9}$$

$Q^*(\delta)$ is continuous, monotone increasing with $Q^*(1) = Q^*$, and $P^*(\delta)$ is continuous, monotone decreasing with $P^*(1) = c_H$. Fix $\varepsilon > 0$. Choose $\overline{\delta}(\varepsilon)$ such that for all $\delta > \overline{\delta}(\varepsilon)$, $Q^*(\delta) > Q^*(1) - \varepsilon/2$. For each such δ choose $m(\delta), k(\delta)$ such that

$$Q^*(\delta) \geq Q_{m(\delta),k(\delta)}(\delta) > Q^*(\delta) - \varepsilon/2, \tag{10}$$

so that $Q_{m(\delta),k(\delta)} > Q^* - \varepsilon$. Such $m(\delta), k(\delta)$ exist because $\{Q_{m,k}(\delta)|m, k \in \mathbb{N}\}$ is dense in $[0, 1]$. Further, by (10) and (8), such $(\delta, m(\delta), k(\delta))$ satisfy (7). Then, by the above construction, there exists an equilibrium in which low quality trades efficiently, and high quality's average trading frequency is $Q_{m(\delta),k(\delta)}(\delta)$. ■

References

ASRIYAN, V., W. FUCHS AND B. GREEN, "Information Spillovers in Asset Markets with Correlated Values," *American Economic Review* 107 (July 2017), 2007–40.

BERGEMANN, D. AND J. HÖRNER, "Should First-Price Auctions Be Transparent?," *American Economic Journal: Microeconomics* 10 (August 2018), 177–218.

BOCHET, O. AND S. SIEGENTHALER, "Competition and Price Transparency in the Market for Lemons: Experimental Evidence," *AEJ: Microeconomics* (forthcoming).

CHAVES, I. N., "Privacy in Bargaining: The Case of Endogenous Entry," *Working Paper* (2020).

CULLEN, Z. B. AND B. PAKZAD-HURSON, "Equilibrium Effects of Pay Transparency in a Simple Labor Market," *Working Paper* (2020).

DENECKERE, R. AND M.-Y. LIANG, "Bargaining with Interdependent Values," *Econometrica* 74 (2006), 1309–1364.

FUCHS, W., A. OERY AND A. SKRZYPACZ, "Transparency and distressed sales under asymmetric information," *Theoretical Economics* 11 (2016), 1103–1144.

HÖRNER, J. AND N. VIEILLE, "Public vs. Private Offers in the Market for Lemons," *Econometrica* 77 (2009a), 29–69.

———, "Public vs. Private Offers: The Two-Type Case to Supplement 'Public vs. Private Offers in the Market for Lemons'," *https://www.econometricsociety.org/sites/default/files/6917_extensions_0.pdf* (2009b).

HUANGFU, B. AND H. LIU, "Information Spillover in Multi-good Adverse Selection," *Working Paper* (2019).

HWANG, I. AND F. LI, "Transparency of outside options in bargaining," *Journal of Economic Theory* 167 (2017), 116–147.

JANSSEN, M. C. W. AND S. ROY, "Dynamic Trading in a Durable Good Market with Asymmetric Information," *International Economic Review* 43 (2002), 257–282.

KAYA, A. AND Q. LIU, "Transparency and price formation," *Theoretical Economics* 10 (2015), 341–383.

KAYA, A. AND S. ROY, "Market Screening with Limited Records," *Working Paper* (2020).

KIM, K., "Information about sellers' past behavior in the market for lemons," *J. Econ. Theory* 169 (2017), 365–399.