Price Transparency and Market Screening

Ayça Kaya* and Santanu Roy†

September 17, 2026

Abstract

In a lemons market with repeated trading by sellers with persistent private information, price transparency affects gains from trade and the market's ability to learn the seller's type. With price transparency, a class of complete-learning equilibria approaches the upper bound on gains from trade imposed by private information when the seller is patient. Without price transparency, prior work shows that complete learning is impossible and gains from trade cannot reach this bound. Price transparency changes incentives by inducing sellers to reject profitable but lower-than-expected offers to preserve future rents, allowing high quality to trade slowly without distorting low quality's trade.

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 such as bids, offers and the prices at which goods were traded, 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 a dynamic lemons market in which a long-lived seller with persistent private information repeatedly trades with a sequence of short-lived buyers, where the past trading prices are observable to subsequent buyers. By contrasting our novel results with the existing results pertaining to markets where such price information is not available, we shed light on the impact of price transparency on two aspects of the market outcomes: first, whether the market can conclusively learn the seller's type; and second, on the limits on the gains from trade stemming from the asymmetric information.

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] Finally, observability of past transaction prices is a central feature of regulation introduced in

[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`

over-the-counter (OTC) financial markets.[3] For example, the Trade Reporting and Compliance Engine (TRACE), introduced by the U.S. Securities and Exchange Commission (SEC) for corporate bond markets in the early 2000s, requires dealers to publicly disseminate transaction prices and volumes in near real time. Similarly, the European Union's Markets in Financial Instruments Directive II (MiFID II), which came into force in 2018, mandates trade reporting and requires a minimum share of transactions to be executed on regulated electronic trading platforms. Our study contributes to the understanding of issues surrounding these kinds of legislation or regulation.

Specifically, our results establish a theoretical upper bound on gains from trade that can be obtained when past transaction prices are observable and show that equilibria approximating this upper bound on welfare exist as the seller becomes patient. In Section 4 we compare these findings to the known results in the literature on limits of gains from trade featuring different levels of transparency to derive policy implications and highlight underlying economic mechanisms.

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 can be high (H) or low (L), determines both its use value ($v_H > v_L$) and its production cost ($c_H > c_L$). The quality is persistent over time and is the seller's private information. For either type of the seller, the use value exceeds the cost of production, and therefore trade is always efficient.

We assume that the high quality's cost exceeds the low quality's use value, and therefore this is a lemon's market (Akerlof (1970)): when the market is pessimistic about the quality, and in the absence of learning by the market, the high quality seller cannot trade. When the seller participates in trading repeatedly, the market can infer the seller's private information from his past trading behavior. Naturally, credible learning requires that the high quality seller trades less frequently than efficient. A very simple observation, for instance, is that if the market immediately and perfectly learns the seller's type the high quality seller's expected discounted frequency of trading cannot exceed Q^* satisfying

$$v_L - c_L = Q^*(c_H - c_L).$$

[3] In the OTC finance literature, the observability of past transaction prices is commonly referred to as "post-trade price transparency", in contrast to "pre-trade price transparency", which refers to the public observability of current dealer quotes.

This is because the right hand side is the smallest payoff the low quality seller can obtain by mimicking while the left hand side is the highest payoff he can obtain by revealing that he is low type. Our first main result establishes that, in any equilibrium in this environment, regardless of whether the market fully learns the buyer's type or if it does, whether this happens immediately, the same upper bound applies to the high quality seller's trading.[4]

Next, we demonstrate (by construction) that when past transactions and prices are publicly observed, the market can quickly and fully learn the seller's quality. Moreover, the outcomes of such equilibria can approximate the said upper bound on the gains from trade as the seller becomes arbitrarily patient.

In the equilibria we construct, the high quality seller trades at a constant price, say P, at a per-period frequency Q, while the low quality seller trades efficiently. There are two incentive problems that restrict the equilibrium values of P and Q. First is the no-mimicking constraint: the high quality seller's path should not be too attractive for the low quality seller to mimic. When the seller is not too impatient, this requires that the high quality seller trades sufficiently slowly. Then, the second restriction emerges: the high quality seller should be willing to slow his trading down even after the market has learned its quality, in which case the buyers would be willing to offer prices that exceed his cost. Thus, it is necessary that the high quality seller is willing to forego instantaneous profits. High quality seller is willing to do this only if this helps him protect future rents, which are positive only when $P > c_H$. [5]

To understand the asymptotic welfare result, note that the first restriction (no mimicking by the low quality seller) introduces a trade-off between P and Q: if the high quality's trading price is very high, the trading frequency must be low and vice versa in order not to render this path attractive for the low quality seller to mimic. Since efficiency depends on the frequency of trade, and not on trading price, the smaller the P, the more frequently can the high quality trade and hence the larger the gains from trade; with the limiting value $P = c_H$ delivering the upper bound Q^* on this trade. As the seller becomes more patient, the one-time payoff of $v_H - c_H$ becomes less valuable relative to future gains even at a fixed future frequency of trading. That is why, in the patience limit, prices arbitrarily close to c_H–and therefore per-period rents arbitrarily close to 0–can incentivize the high

[4]We maintain an additional restriction that the high quality seller never trades below cost.

[5]To foreshadow the comparison with the cases where transaction prices are not public, notice that if trading prices were not observable, at histories where trading is on the path of the equilibrium, the high quality seller would strictly prefer to accept any offer exceeding his cost. And thus, slow trading by the high quality could not be sustained if beliefs exceed the lemons threshold.

quality seller to wait, leading to our final main result that shows that the upper bound on gains from trade discussed above can be asymptotically approached.

In what follows, we first discuss the related literature. Section 2 presents the model, and Section 3 presents our main results. Section 4 discusses markets with other forms of transparency and identifies a common thread that reconciles our results with those in the existing literature. All proofs are relegated to an appendix.

Related literature The role of price transparency in markets with adverse selection 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 since a single trade ends the interaction. 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. [6] 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 positive rents,—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 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 transaction prices (i.e., the accepted offers) are observed.[7]

In a variation of the single sale framework, Asriyan et al. (2017) analyze a model of parallel bargaining by informed sellers of correlated assets where each asset can be traded only once; however, knowing that an asset has been traded can be informative about the value of other assets. In their model, transaction price conveys no additional information to

[6] 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.

[7] 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.

uninformed buyers (beyond whether trade has occured); transparency of trading oucomes can increase or decrease welfare and trading volume, depending on the degree of asset correlation and the structure of demand.[8]

The effect of transparency of past trades is somewhat more nuanced in markets where a seller may split the sale over multiple periods. Fuchs et al. (2026) study the sale of one unit of a perfectly divisible asset to short-lived buyers; the seller chooses when to sell as well as whether to split the sale over time. When trade takes place only at discrete dates, sellers split their trade over time, creating a second dimension of private information when buyers cannot observe past trades; the latter can severely limit the market's ability to screen. They show that the effect of price (and volume) transparency on welfare is ambiguous. The strategic considerations in their framework are significantly different from that in our paper because in our environment the seller can sell a unit in every period and thus there are no intertemporal capacity constraints.

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. Kaya and Roy (2022) considers variation in the length of trading records but does not address the effect of price transparency. Kaya and Roy (2024) studies the impact of observability of past trading volumes (but not the prices) on market efficiency and how this varies with the degree of buyer competition and market's initial perception. In that paper, buyer competition plays a similar role to price transparency here, by allowing the high quality to extract rents. Section 4 of this paper contains a discussion of how price transparency as analyzed in our paper can improve welfare relative to observability of trade volumes.

The role of price transparency has been studied in other contexts. The literature on OTC financial markets has studied the effect of (initially) uninformed customers being able to observe past transaction prices. Thus, Vairo and Dworczak (2025) study separating equilibria in a dynamic market with two competing dealers and show that price transparency weakly reduces efficiency when customers observe both contemporaneous price quotes, but improves efficiency when customers observe only one. An earlier paper by Kakhbod and Song (2020) shows that when customers observe the entire history of price offers, including rejected offers, richer punishment strategies can be used to sustain less informative (pool-

[8] See, also, Huangfu and Liu (2019)

ing) equilibria.[9] The dynamic trading setup in OTC financial markets is very different from our paper where an exogenously fixed quantity of the good or service is offered for sale every period regardless of previous trades. Further, we study screening rather than signaling; in particular, we have one informed seller that trades sequentially with short-lived, uninformed monopsonists. In our setting, price transparency improves both learning and allocative efficiency.[10]

2 Model

A long-lived seller can produce one unit of output every period. Time is discrete and horizon is infinite. Time periods are indexed $t = 1, 2, \cdots$. 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, the buyer leaves the game, and the seller moves to the next period.

Seller's type $\theta \in \{L, H\}$ determines both the use value (v_θ) of his output and the cost (c_θ) of production. Seller's type is her private information. Regardless of the seller's type, gains from trade is strictly positive: $v_\theta - c_\theta > 0, \theta \in \{L, H\}$.

All buyers hold a common prior that assigns probability μ_0 to type $\theta = H$. We assume that $c_H > v_L$ and that

$$\mu_0 < \mu^* := \frac{c_H - v_L}{v_H - v_L},$$

so that the expected use value $\mu_0 v_H + (1-\mu_0) v_L$ is strictly less than the cost c_H of producing high quality.

If, in a given period, trade takes place at price P, the type-θ seller's payoff in that period is $P - c_\theta$ and her trading partner's payoff is $v_\theta - P$. The instantaneous payoff for any party who does not trade is 0. The seller maximizes the expected discounted sum of his future payoffs using discount factor $\delta \in [0, 1]$.

We express all payoffs and trading volumes in average per-period terms. For instance, if along a path of play seller type θ trades with a sequence of probabilities α_t, $t = 1, \cdots$,

[9] See also Kakhbod and Song (2024).

[10] In a different context, 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 (2021) provides experimental evidence on the interaction of price transparency and buyer competition.

then per-period gains from trade he generates is

$$Q_\theta = (1-\delta)\sum_{t=1}^{\infty} \delta^{t-1}\alpha_t.$$

Histories Our main model is a **market with price transparency**. In this model, a public history at time t contains information about the calendar time, whether trade took place in each $t' < t$, and realized trading prices in periods $t' < t$ where trade took place.[11] Let $\mathcal{H}$ represent the set of such public histories and let $h_t \in \mathcal{H}$ be an arbitrary t-period history. We also let $h_\emptyset \in \mathcal{H}$ represent the null history. Lastly, (h_t, P) represents a $t+1$-period continuation history of h_t where in period $t+1$, trade takes place at price P and $(h_t, \emptyset)$ represents a continuation history where period $t+1$ features no trade.

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. 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 $\alpha^{h,\theta} : \mathbb{R}_+ \to [0,1]$, where $\alpha^{h,\theta}(P)$ is the probability with which type θ seller accepts an offer of P at history $h \in \mathcal{H}$. 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 restrict attention to equilibria that satisfy the following condition:

Condition 1 An equilibrium exhibits no trade below cost by the high quality if at no equilibrium path history, the seller of type H trades at a price $P < c_H$ with positive probability.

Henceforth, "equilibrium" refers to perfect Bayesian equilibria satisfying Condition 1.

[11] Seller's private history also includes information about the prices that were offered but rejected. In all equilibria we consider, strategies depend only on public histories.

3 Gains from trade

In Section 3.1, we establish a theoretical upper bound on the gains from trade in any equilibrium. Then, in Section 3.2 we construct a class of equilibria that exhibit complete learning. Finally in Section 3.3, we show that as the seller becomes more patient, the equilibria in this class will approximate the theoretical maximum surplus established in Section 3.1

3.1 Bounding the gains from trade

As is well-understood, in markets with adverse selection high-quality sellers cannot trade frequently because of the strong incentive of the low-quality sellers to mimic them. In particular, the trading path of the high-quality seller must not be too attractive for the seller of a low-quality object. This requires that the high-quality seller trades slowly and/or at low prices. Part of the equation that determines how attractive mimicking is for the low-quality seller is the alternative he faces, i.e., the payoff he receives if he does not mimic the high-quality seller, but instead reveals himself. Intuitively, if this alternative payoff is high, the non-mimicking constraint that governs the high quality seller's trading is relaxed, and, in turn, the high-quality seller can trade relatively frequently. In the next lemma, as a first step towards our bound on overall gains from trade, we establish an upper bound on the payoff that the low-quality seller can receive in any equilibrium. Then, by the preceding discussion, this bound informs how severely adverse selection limits gains from trade.

Lemma 1 *In any equilibrium, at any on-path history h, if $\mu(h) < \mu^*$, then $V_L(h) \leq v_L - c_L$.*

The bound $v_L - c_L$ on the low-quality seller's payoff would be immediate in a static environment when the buyer's belief is less than μ^*: since the high-quality seller is not willing to trade below a price c_H, and since the expected quality of the object for sale does not justify offering a price of c_H or above, the highest price that could be offered would be v_L, and therefore the highest payoff that the low-quality seller can receive would be $v_L - c_L$. In the dynamic environment that we study, the argument is less straightforward, since the low-quality seller does not necessarily accept all prices exceeding cost, as continuation value after such acceptance may be lower than after rejection. This may allow a buyer to safely offer a price exceeding c_H as long as it is accepted with sufficiently low probability by the low-quality seller so that the average quality conditional on trade remains above that price. The insight that the proof of Lemma 1 carefully leverages is that the low-quality seller must be willing to reject any equilibrium path offers exceeding c_H at histories where the belief is low, and therefore his payoff can be calculated along a path where he rejects such offers.

The next proposition uses the fact that the low quality seller's equilibrium payoff should be no less than what he would get if he mimicked the trading path of the high quality seller. Letting Q_H represent the average discounted quantity of trade by the high quality seller along the equilibrium path, and using the fact that the high quality seller's expected trading price is bounded from below by c_H, the low quality seller's payoff from mimicking is no less than $Q_H(c_H - c_L)$. Then using Lemma 1, necessarily $Q_H \leq Q^*$, where

$$Q^* := \frac{v_L - c_L}{c_H - c_L}, \tag{1}$$

so that the low quality seller's equilibrium payoff is no less than what he could obtain if she mimicked high quality's trading path. The proof of Proposition 1 makes this argument formal.

Proposition 1 *In any equilibrium exhibiting no trade below cost, the expected discounted per-period gains from trade is no larger than*

$$W^*(\mu_0) := \mu_0 Q^*(v_H - c_H) + (1 - \mu_0)(v_L - c_L).$$

The expression in Proposition 1 uses the fact that the high quality seller's trading frequency cannot exceed Q^* and the low quality seller's trading frequency naturally cannot exceed 1.

3.2 Immediate complete learning equilibria

In this section we characterize a class of equilibria which approximate the upper bound on gains from trade established in Proposition 1 as the seller becomes arbitrarily patient.

This class of equilibria feature immediate and complete learning by the market (i.e. full separation in the first period) and have a very simple "trigger-strategy" structure. Low quality seller trades every period at price v_L while high quality seller trades at a fixed price P with expected per-period frequency Q_H satisfying $v_L - c_L \leq Q_H(P - c_L)$. Specifically, each of these equilibria are identified by a pair (T, P). On the path of this equilibrium the

H-type seller does not trade in the first T periods. Thereafter, she trades at price P with probability 1 in each period. The L-type seller trades at price v_L every period. Thus, in this equilibrium $Q_H = \delta^T$ and $Q_L = 1$.

If trade occurs unexpectedly, e.g. at a period $t < T$ after no trade at $t = 1$ or if trade occurs at a lower than expected price, the play reverts to a punishment equilibrium in which both types of the seller receives a payoff of 0. The existence of such a punishment equilibrium is established in Lemma 2 in the Appendix.

In our construction, the pair (P, T) satisfies

$$v_L - c_L = \delta^T (P - c_L), \tag{2}$$

so that the L-type seller is indifferent between mimicking H-type seller's equilibrium path and revealing his type in the first period; and

$$(1 - \delta)(v_H - c_H) < \delta^{T-1}(P - c_H), \tag{3}$$

so that the H type seller is willing to wait the additional $T - 1$ periods before trade starts after his type is revealed starting period 2. Proposition 2 shows that it is possible to construct such an equilibrium as long as (2) and (3) are satisfied.

Proposition 2 *Assume that $\delta > 1/2$ and there exists (P, T) with $P < v_H$, satisfying (2) and (3). Then, there is an immediate complete learning equilibrium in which $Q_H = \delta^T, Q_L = 1$ so that the expected gains from trade is*

$$\mu_0 \delta^T (v_H - c_H) + (1 - \mu_0)(v_L - c_L).$$

3.3 Attaining the upper bound

In the previous two sections we established an upper bound on equilibrium welfare and constructed a class of equilibria. Next we show that as the seller becomes arbitrarily patient equilibria within that class can attain gains from trade that are arbirarily close to that upper bound. This result is formally stated in Theorem 1.

Theorem 1 *For every $\varepsilon > 0$, there exists $\bar{\delta} < 1$ such that when $\delta > \bar{\delta}$ there exists an equilibrium which generates gains from trade no less than $W^*(\mu_0) - \varepsilon$.*

The main thrust of the proof is to show that as δ becomes sufficiently large it is possible to find a pair (P, T) which satisfies (2) and (3); and such that δ^T is arbitrarily close to Q^*. Then the claim follows by Proposition 2.

The conditions (2) and (3) on (P, T) are simply that the high quality seller should want to wait to trade at higher prices while the low quality seller should not, so that waiting is a credible signal of high quality. For a fixed (P, T) as δ grows, waiting becomes attractive for both types of the seller. However, low quality seller's alternative of trading at v_L every period is independent of δ while the high quality seller's alternative of trading one time at price v_H becomes relatively less attractive as δ grows.

If the high quality's trading price P decreases, the no mimicking condition (2) becomes easier to satisfy, and thus the waiting period T can be shortened. In fact, when P is equal to its lower bound c_H, (2) constrains the trading frequency to be its upper bound Q^* regardless of δ.[12] On the other hand, when P is small, the per-period rent $P - c_H$ that the high quality seller expects to get after the waiting is over becomes small. But, as δ grows, even such small per-period rent becomes attractive relative to the high quality's alternative. Because of this, as $\delta \to 1$, prices arbitrarily close to c_H can be part of an equilibrium. For such P, the associated δ^T that satisfies (2) approaches Q^*.

4 Discussion: transparency, market screening and efficiency

Consider an opaque market, i.e., a market where subsequent buyers do not observe any particulars of the long-run seller's past, including his entry time, past transaction volumes or prices. In this case, the seller accepts any offer that exceeds his cost, since this decision has no impact on future outcomes. Given this, each buyer solves an isolated problem of picking a take-it-or-leave-it price offer. In a lemons market (i.e. when $\mu_0 < \mu^*$), and buyers arrive one by one, each buyer offers a price $P = c_L$.[13] In this outcome, the low quality trades efficiently, while the high quality never trades.

[12] This is subject to integer constraints in this discrete-time model. However, the loss due to integer constraints vanish as $\delta \to 1$.

[13] Note that, since the seller never leaves the market, regardless of trading outcomes, the buyers do not update their belief based on the equilibrium strategies. Further, It is easy to show that even if the buyers could offer screening contracts, for initial beliefs $\mu_0 < \mu^{**} := \frac{c_H - c_L}{v_H - c_L}$, the optimal menu would exclude the high quality and trade only with the low. Since we consider $\mu_0 < \mu^* := \frac{v_L - c_L}{v_H - c_L} < \mu^{**}$, for the parameter values we consider, an opaque market with or without screening contract, excludes high quality.

Such opaqueness renders the market essentially static, and provides no opportunities for the market to learn. Transparency of past behavior allows the market to screen the long run seller, by creating different possible trading paths to which different types of the seller self-select. Even though transparency introduces screening and allows the high quality to trade, it does not completely resolve the inefficiency introduced by the lemons problem, because credible learning rules out efficient trading by the high quality seller. Further, it may reduce trading of the low quality. If the latter is the case, then introducing transparency generates a trade-off between the high quality's and the low quality's trading, and the question of whether it increases and decreases overall gains from trade becomes non-trivial.

The existing results on the impact of transparency on gains from trade are mixed, and depends both on the type of transparency introduced and the particulars of the environment. For instance Kaya and Roy (2024) shows that, in repeated trading, if solely the past trading quantities, and not the prices are observed, outcomes depend on the market structure: such transparency increases gains from trade if in each period the seller meets multiple buyers (intra-temporal competition) while it reduces gains from trade if the seller meets a single buyer (intra-temporal monopsony).

In this paper, we fill a gap by showing that making the past trading prices transparent increases gains from trade even in the case of intra-period monopsony. In doing so, we uncover a common thread underlying these existing results: whether the market can screen the seller without lowering the low quality's trading depends crucially on whether the high quality seller can sustain rents in equilibrium. This is the case in the equilibria we construct in Section 3.2 for a market with trading-price transparency as well as a market with intra-period competition and transparency of just past trading volumes. Thus, those cases are when transparency increases gains from trade relative to an opaque market.

To understand the intuition recall that market screening requires the high quality to trade less than efficiently; i.e. with probability less than 1, at certain histories. There are two possible ways this may happen: either the buyer(s) don't find it profitable to offer prices that exceed high quality seller's cost, or the high quality seller is willing to reject prices that exceed his cost. The first is possible only when the buyers are not too optimistic, which in equilibrium can occur only if the low quality seller also follows (mimics) the high quality's inefficient path with sufficiently high probability. This represents a reduction of gains from trading of low quality relative to the opaque market. If on the other hand, the high quality seller is willing to reject high prices, the low quality seller need not pool, and can still trade efficiently in equilibrium. In turn, the high quality seller is willing to turn down positive surplus in the current period only if this protects future rents. Thus, in environments where future rents are not possible, inefficient pooling by the low quality seller cannot be avoided.

Remark: More optimistic markets Our paper focuses on low initial beliefs so that in an opaque market, low quality never trades. When, as in this model, the seller meets a unique buyer every period, for beliefs that are slightly above μ^*, high quality still does not trade in an opaque market. The belief cutoff above which high quality can trade is $\mu^{**} = \frac{c_H - c_L}{v_H - c_L} > \mu^*$.[14] This is due to the monopsony distortion: each buyer can extract all the surplus from the low quality seller, making it particularly rewarding to target him. Thus, buyers require a higher level of optimism (higher beliefs) before they are convinced to trade with the high quality seller. When past trading volumes but not prices can be observed, the low quality seller can capture information rents, but high quality cannot. This weakens the monopsony distortion, making targeting the high quality more attractive. In fact, Kaya and Roy (2024) establishes that for this intermediate belief range, with past trade volume observability there is an essentially unique equilibrium the outcome of which approaches full efficiency as the seller becomes patient, so that the information rents become important. In this case, price transparency, by allowing the high quality seller to also sustain rents, leads to equilibria where the high quality trades less than efficiently.[15] These observations reveal that policy discussions about transparency of market outcomes must take into account the subtle impacts of market structure and market's initial optimism.

[14]To see this note that a monopsonist buyer holding belief μ, receives a payoff of $\mu v_H + (1-\mu)v_L - c_H$ when targeting the high quality seller (and thus trading with probability 1, and he receives a payoff of $(1-\mu)(v_L - c_L)$ when targeting only the low quality seller. The cutoff belief μ^{**} equates these payoffs.

[15]Note that the upper bound in gains from trade established in Proposition 1 does not apply to higher beliefs, thus there may exist equilibria with higher gains from trade for this belief range. However, the equilibrium construction in Section 3.2 is valid independent of initial belief. Together with the fact that without price transparency, the market approaches efficiency (as shown in Kaya and Roy (2024)), we conclude that for this range of beliefs additional price transparency may reduce efficiency by creating additional undesirable equilibrium outcomes.

A Appendix: Omitted proofs

A.1 Proofs for Section 3.1

Proof of Lemma 1. Let

$$\bar{V}_L := \sup\{V_L(h) \mid h \text{ is on the equilibrium path and } \mu(h) < \mu^*\}.$$

This supremum exists because the set over which its taken is non-empty ($\mu(h_\emptyset) < \mu^*$) and $V_L(h)$ is bounded above by $v_H - c_L$.

Take $\varepsilon > 0$ such that $\varepsilon < (1-\delta)\bar{V}_L$, and take h such that $V_L(h) > \bar{V}_L - \varepsilon$.

Let Π_h be the support of the buyer's non-losing (serious) offers at history h, and let F be the cdf of the buyer's offer distribution, including non serious offers so that the probability of trade is $F(\Pi_h)$. First note that $\Pi_h \neq \emptyset$ (i.e. the buyer makes a serious offer with positive probability.) Because otherwise, $\mu(h,\emptyset) = \mu(h)$ and $V_L(h) = \delta V_L(h,\emptyset) < \delta\bar{V}_L$. This is a contradiction because $V_L(h) > \bar{V}_L - \varepsilon > \delta\bar{V}_L$.

Take $P \in \Pi_h$. Let $\alpha_\theta(P \mid h)$ be the probability that the buyer type $\theta = L, H$ accepts this offer. Then, the seller's payoff from offering P is

$$(1-\mu(h))\alpha_L(P \mid h)(v_L - P) + \mu(h)\alpha_H(P \mid h)(v_H - P)$$

$$=$$

$$((1-\mu_h)\alpha_L(P \mid h) + \mu_h\alpha_H(P \mid h))\left(\mu(h,P)v_H + (1-\mu(h,P))v_L - P\right).$$

This follows because $\mu(h,P) = \frac{\mu(h)\alpha_H(P|h)}{(1-\mu(h))\alpha_L(P|h)+\mu(h)\alpha_H(P|h)}$.

If $P > v_L$, then L-type seller rejects it with positive probability; i.e. $\alpha_L(P \mid h) < 1$, because otherwise $\mu(h,P) \leq \mu(h) < \mu^*$, and thus the buyer's payoff from offering P is negative. Thus, $(h,\emptyset)$ is on the equilibrium path, and since L-type seller is rejecting with positive probability

$$\delta V_L(h,\emptyset) \geq (1-\delta)(P - c_L) + \delta V(h,P).$$

Next, for future reference, note that by the law of iterated expectations,

$$\mu(h) = \int_{P\in\Pi_h} \mu(h,P)dF(P) + (1 - F(\Pi_h))\mu(h,\emptyset). \tag{4}$$

There are two possibilities for $(h,\emptyset)$: (i) if $\mu(h,\emptyset) < \mu^*$, then $V_L(h,\emptyset) \leq \bar{V}_L$; (ii) if $\mu(h,\emptyset) \geq$

μ^*, then by (4), there exists $P' \in \Pi_h$ such that $\mu(h, P') < \mu^*$. That conditional on acceptance of P' the belief is less than μ^* implies that $P' < c_H$, because otherwise, the buyer makes negative payoff. Further, since H type does not accept such offers $\mu(h, P') = 0$, which implies that $P' \leq v_L$ and that $V_L(h, P') \leq \bar{V}_L$. Finally, since P' is a serious offer, L type is willing to accept it, so that

$$\delta V_L(h, \emptyset) \leq (1-\delta)(P' - c_L) + \delta V_L(h, P') \leq (1-\delta)(v_L - c_L) + \delta \bar{V}_L.$$

We conclude that in either case,

$$\delta V_L(h, \emptyset) \leq (1-\delta)(v_L - c_L) + \delta \bar{V}_L.$$

So we have shown that if $P > v_L$, then

$$(1-\delta)(P - c_L) + \delta V(h, P) \leq \delta V_L(h, \emptyset) \leq (1-\delta)(v_L - c_L) + \delta \bar{V}_L.$$

Now consider $P \in \Pi_h$ with $P < c_H$. By the same arguments above given for P',

$$(1-\delta)(v_L - c_L) + \delta \bar{V}_L \geq (1-\delta)(P - c_L) + \delta V_L(h, P) \geq \delta V_L(h, \emptyset).$$

Here, the first inequality is because P is accepted only by L type, and thus it must be no more than v_L, and its acceptance must lead to belief $0 < \mu^*$. The second inequality is because P is a serious offer and thus L must be willing to accept it.

Now let

$$\tilde{V}_L(h; P) = \max\{(1-\delta)(P - c_L) + \delta V_L(h, P), \delta V_L(h, \emptyset)\}.$$

That is, $\tilde{V}_L(h; P)$ is the L-type seller's continuation payoff at history h conditional on realized price offer P. We have shown that for any $P \in \Pi_h$,

$$\tilde{V}_L(h; P) \leq (1-\delta)(v_L - c_L) + \delta \bar{V}_L,$$

and thus

$$V_L(h) \equiv \int \tilde{V}_L(h; P) dF(P) \leq (1-\delta)(v_L - c_L) + \delta \bar{V}_L.$$

Since $V_L(h) > \bar{V}_L - \varepsilon$ we have $\bar{V}_L - \varepsilon \leq (1-\delta)(v_L - c_L) + \delta \bar{V}_L$, which implies

$$\bar{V}_L < v_L - c_L + \frac{\varepsilon}{1-\delta}.$$

The claim follows since this must hold for any sufficiently small ε. ■

Proof of Proposition 1. Fix an equilibrium exhibiting no trade below cost by the H type. Let $\gamma_\theta(h)$ be the probability with which type $\theta \in \{L, H\}$ reaches history h, $\tau(h)$ be the length of history h; $F_h(P)$ be the cdf of the price offers that the buyer makes at history h and $\alpha_\theta(h; P)$ be the probability that he accepts price P at history h, if offered. Finally, let $\eta_\theta(h) := \int \alpha_\theta(h; P) dF_h(P)$ be the probability with which type θ trades at history h conditional on reaching it. Then,

$$V_\theta(h_\emptyset) = (1-\delta) \sum_{h \in \mathcal{H}} \gamma_\theta(h) \delta^{\tau(h)} \int \alpha_\theta(h; P)(P - c_\theta) dF_h(P)$$

Define the expected discounted total probability of trade conditional on type θ to be

$$Q_\theta := (1-\delta) \sum_{h \in \mathcal{H}} \gamma_\theta(h) \delta^{\tau(h)} \eta_\theta(h)$$

Since necessarily $V_H(h_\emptyset) \geq 0$, we have

$$(1-\delta) \sum_{h \in \mathcal{H}} \gamma_H(h) \delta^{\tau(h)} \int \alpha_H(h; P) P dF_h(P) \geq Q_H c_H.$$

Further, it is necessary that

$$V_L(h_\emptyset) \geq (1-\delta) \sum_{h \in \mathcal{H}} \gamma_H(h) \delta^{\tau(h)} \int \alpha_H(h; P) P dF_h(P) - Q_H c_L,$$

so that L type has no incentive to follow H type's equilibrium path. The latter two inequalities imply

$$V_L(h_\emptyset) \geq Q_H(c_H - c_L).$$

Next, since by Lemma 1, $V_L(h_\emptyset) \leq v_L - c_L$, we necessarily have $v_L - c_L \geq Q_H(c_H - c_L)$, or equivalently $Q_H \leq Q^*$. The claim of the proposition follows by also noting that $Q_L \leq 1$. ■

A.2 Proofs for Section 3.2

Lemma 2 (Punishment equilibria) *Consider any history h with $\mu(h) = 0$. There exists a continuation equilibrium in which $V_L(h) = V_H(h) = 0$.*

Proof of Lemma 2. In this continuation equilibrium, $\mu(h') = 0$ for any continuation history h' of h, on or off-the equilibrium path. At each history the buyer offers c_L with probability 1, θ type seller accepts any offer greater than or equal to c_θ with probability 1, and rejects all others with probability 1. Given the seller's strategy, and given $\mu(h') = 0$, the buyers' strategies are optimal at every continuation history h'. The seller's strategy is optimal because at every continuation history h', $V_\theta(h') = 0$, $\theta = L, H$. ∎

Lemma 3 (L-seller's equilibrium path) *Assume that $1 > \delta > 1/2$. Consider any history h with $\mu(h) = 0$. There exists a continuation equilibrium in which $V_L(h) = v_L - c_L$ and $V_H(h) = 0$.*

Proof of Lemma 3. The following strategies and beliefs form an equilibrium: at every continuation history h', $\mu(h') = 0$.

If the history h' features no trade at $P < v_L$ since h then the buyer offers v_L. At all such histories, L type seller accepts all offers v_L or above with probability 1, rejects all others with probability 1, while the H type seller accepts all offers c_H or above with probability 1, and rejects all others with probability 1. Otherwise, the continuation equilibrium starting at h' is the one characterized in Lemma 2, so that, in particular, $V_L(h') = 0$.

Consider h' which features no trade at $P < v_L$ since h. Given that $\mu(h') = 0$ and given L seller's strategy, the buyer cannot make positive profits. Thus, offering v_L is optimal. The H-type seller's strategy is optimal because at every continuation history $V_H(h') = 0$. To see that the L-type seller's strategy is optimal, consider an offer $P < v_L$. If the L-seller accepts this offer he receives a payoff of $(1 - \delta)(P - c_L)$ while if he rejects it he receives a payoff $\delta(v_L - c_L)$. Since $\delta > 1/2$, it is optimal for the L-type seller to reject such offers. Now consider an offer $P \geq v_L$. If the L-type seller accepts this offer he receives a payoff $(1 - \delta)(P - c_L) + \delta(v_L - c_L)$ while if he rejects it he receives a payoff $\delta(v_L - c_L)$. Thus, it is optimal for the L-type seller to accept such offers. ∎

Proof of Proposition 2. We first consider histories h_1 of length 1. If the history features trade at $P \geq v_L$, then $\mu(h_1) = 0$ and the equilibrium described in Lemma 3 is played. Thus, in particular, $V_L(h_1) = v_L - c_L$. If the history featured trade at $P < v_L$, then $\mu(h_1) = 0$ and the equilibrium described in Lemma 2 is played. Thus, in particular, $V_L(h_1) = 0$. If the history features no trade, then $\mu(h_1) = 1$ and the following continuation equilibrium is played:

We classify continuation histories of h_1 featuring no trade as follows:

- $\mathcal{H}_1$: History length is $t \leq T - 1$, and the history features no trade.
- $\mathcal{H}_2$: History length is $t > T - 1$, and the history features no trade in the first T periods, and features no trade at prices strictly less than P.
- $\mathcal{H}_3$: All other histories.

Beliefs:

$$\mu(h) = \begin{cases} 1 & \text{if } h \in \mathcal{H}_1 \cup \mathcal{H}_2 \\ 0 & \text{if } h \in \mathcal{H}_3 \end{cases}.$$

In histories $h \in \mathcal{H}_3$, the play follows the equilibrium described in Lemma 2. Next we describe strategies at $h \in \mathcal{H}_1 \cup \mathcal{H}_2$.

Buyer strategies:

- If $h = h_\emptyset$, offer v_L with probability 1.
- If $h \in \mathcal{H}_1$, offer v_L with probability 1.
- if $h \in \mathcal{H}_2$, offer P with probability 1.

Seller's strategies:

- If $h \in \mathcal{H}_1$, reject all prices less than or equal to $\bar{P}_t^\theta$ with probability 1, where t is the length of the history and

  $$(1 - \delta)(\bar{P}_t^\theta - c_\theta) = \delta^{T-t}(P - c_\theta).$$

  Accept all others with probability 1. We note that $\bar{P}_t^H > v_H$ for each $t \geq 1$ by (3).
- If $h = h_\emptyset$, then
  - If $\theta = L$, accept all offers that weakly exceed v_L, and reject all others with probability 1.

- If $\theta = H$, accept all offers that strictly exceed $\bar{P}_0^H$, and reject all others with probability 1.

- If $h \in \mathcal{H}_2$, then accept all offers greater than or equal to P and reject all others with probability 1.

Verification

- Belief consistency: There are three types of on-path histories:
  - History features trade every period at price v_L, which has positive probability only conditional on L-type seller. At these histories belief is 0, which is consistent.
  - $h \in \mathcal{H}_1$. Given the describes strategies, only H type seller reaches these histories on the equilibrium path. This $\mu(h) = 1$ is consistent.
  - History length is $t > T$, and the history features no trade in the first T periods, and features trade at price P in every period thereafter. These are a subset of $\mathcal{H}_2$, and thus the belief is 1. Further, these have positive probability only conditional on type-H seller, and thus the belief is consistent.

- Optimality of buyer strategies:
  - at $h_\emptyset$: By assumption, H-type seller accepts no price that is less than or equal to c_H. Thus, offering $P > v_L$ leads to a negative payoff since $\mu_0 < \mu^*$. Offering $P \leq v_L$ leads to payoff 0. Thus, in particular, offering v_L is optimal.
  - If $h \in \mathcal{H}_1$, since belief is 1, offering $P' > \bar{P}_t^H$ leads to a negative payoff, and offering anything less leads to payoff 0. Thus, in particular, it is optimal to offer v_L.
  - If $h \in \mathcal{H}_2$, since belief is 1, offering $P' \geq P$ leads to a payoff of $v_H - P'$ and offering anything lower leads to a payoff o 0. Thus, it is optimal to offer P.

- Optimality of seller strategies:
  - If $\theta = L$ and $h = h_\emptyset$:

    For $P' < v_L$, $V_L(h_\emptyset, P) = (1-\delta)(P'-c_L)$ while $V_L(h_\emptyset, \emptyset) = \delta(v_L - c_L)$. The latter is greater, and thus rejecting such offers is optimal. For $P' \geq v_L$, $V_L(h_\emptyset, P) =$

$(1-\delta)(P'-c_L)+\delta(v_L-c_L)$ while $V_L(h_\emptyset,\emptyset)=\delta(v_L-c_L)$. The former is greater, and thus accepting such offers is optimal.

- If $h \in \mathcal{H}_1$ or $\theta = H$ and $h = h_\emptyset$. Given buyer strategies, for any P', and $\theta = L, H$, $V_\theta(h,P') = (1-\delta)(P'-c_\theta)$, while $V_\theta(h,\emptyset) = \delta^{T-t}(P-c_\theta)$, establishing the optimality of the strategies.
- If $h \in \mathcal{H}_2$: If $P' < P$, then $V_\theta(h,P') = (1-\delta)(P'-c_\theta)$, while $V_\theta(h,\emptyset) = \delta(P-c_\theta)$. Since $\delta > 1/2$, the latter payoff is larger, and thus it is optimal to reject such offers. If $P' \geq P$, $V_\theta(h,P') = (1-\delta)(P'-c_\theta) + \delta(P-c_\theta)$, while $V_\theta(h,\emptyset) = \delta(P-c_\theta)$. Since $P' \geq P > c_H$, the former payoff is larger and thus it is optimal to accept such offers.

■

B Proofs for Section 3.3

Proof of Theorem 1. We show that for any $\varepsilon > 0$, there exists $\bar{\delta} > 0$ such that when $\delta > \bar{\delta}$, we can find $P < v_H$ and T satisfying (2) and (3) such that $\delta^T \geq Q^* - \varepsilon$. The claim of the theorem follows by appropriate adjustment of ε and via Proposition 2.

We re-express (2) and (3) as

$$P = \frac{v_L - c_L}{\delta^T} + c_L \text{ and } P > c_H + \frac{1-\delta}{\delta^{T-1}}(v_H - c_H)$$

Using the first equality to replace P in the second inequality we obtain

$$\frac{v_L - c_L}{\delta^T} + c_L > c_H + \frac{1-\delta}{\delta^{T-1}}(v_H - c_H).$$

Re-arranging and using $Q^* = (v_L - c_L)/(c_H - c_L)$ yields

$$Q^* - \delta(1-\delta)\frac{v_H - c_H}{c_H - c_L} > \delta^T. \tag{5}$$

Define

$$T(\delta) := \min\{T \mid (5) holds\} \text{ and } P(\delta) := \frac{v_L - c_L}{\delta^{T(\delta)}} + c_L.$$

As long as $\delta(1-\delta) < \frac{v_L - c_L}{v_H - c_H}$ such $T(\delta)$ exists because the right-hand-side of (5) approaches 0 as T goes to infinity while the left-hand-side is positive. By construction, $T(\delta), P(\delta)$

satisfy (3) and (2).

Let $1/2 < \bar{\delta} < 1$ be large enough such that $\bar{\delta}(1-\bar{\delta}) < \frac{v_L - c_L}{v_H - c_H}$ and $1 - \bar{\delta} < \varepsilon/2$ and $\delta(1-\bar{\delta})\frac{v_H - c_H}{c_H - c_L} < \varepsilon/2$. Such $\bar{\delta}$ can be found because all inequalities are strictly satisfied at $\bar{\delta} = 1$, and the left-hand-side of each inequality is continuous in $\bar{\delta}$. Also note that each left-hand-side is decreasing in δ for $\delta > 1/2$. Thus, these inequalities are satisfied for all $\delta > \bar{\delta}$. Next, we show that for any $\delta > \bar{\delta}$, $\delta^{T(\delta)} > Q^* - \varepsilon$.

First we claim that

$$Q^* - \delta(1-\delta)\frac{v_H - c_H}{c_H - c_L} - \delta^{T(\delta)} < \varepsilon/2.$$

This is because otherwise $Q^* - \delta(1-\delta)\frac{v_H - c_H}{c_H - c_L} - \delta^{T(\delta)} > \varepsilon/2$, and since $\delta^{T(\delta)-1} - \delta^{T(\delta)} < \varepsilon/2$, we have $Q^* - \delta(1-\delta)\frac{v_H - c_H}{c_H - c_L} > \delta^{T(\delta)-1}$, contradicting to the fact that $T(\delta)$ is the minimum T satisfying this inequality.

Re-writing the above inequality we have

$$Q^* - \delta^{T(\delta)} < \varepsilon/2 + \delta(1-\delta)\frac{v_H - c_H}{c_H - c_L} < \varepsilon,$$

where the last inequality follows because $\delta(1-\delta)\frac{v_H - c_H}{c_H - c_L} < \varepsilon/2$, establishing the claim.

Further, we note that by choosing ε small enough, we have for all $\delta > \bar{\delta}$, $\delta^{T(\delta)}$ arbitrarily close to $Q^* = (v_L - c_L)/(c_H - c_L)$, and thus by (2), $P(\delta)$ is arbitrarily close to c_H, which is strictly less than v_H. ■

References

Akerlof, G. A. (1970). The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics 84*(3), 488–500.

Asriyan, V., W. Fuchs, and B. Green (2017). Information spillovers in asset markets with correlated values. *American Economic Review 107*(7), 2007–40.

Bergemann, D. and J. Hörner (2018, August). Should first-price auctions be transparent? *American Economic Journal: Microeconomics 10*(3), 177–218.

Bochet, O. and S. Siegenthaler (2021). Competition and price transparency in the market for lemons: Experimental evidence. *AEJ: Microeconomics*.

Chaves, I. N. (2020). Privacy in bargaining: The case of endogenous entry. *Working Paper*.

Cullen, Z. B. and B. Pakzad-Hurson (2020). Equilibrium effects of pay transparency in a simple labor market. *Working Paper*.

Fuchs, W., P. Gottardi, and H. Moreiro (2026). Time trumps quantity. *working paper*.

Fuchs, W., A. Oery, and A. Skrzypacz (2016). Transparency and distressed sales under asymmetric information. *Theoretical Economics 11*(3), 1103–1144.

Hörner, J. and N. Vieille (2009a). Public vs. private offers in the market for lemons. *Econometrica 77*(1), 29–69.

Hörner, J. and N. Vieille (2009b). 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*.

Huangfu, B. and H. Liu (2019). Information spillover in multi-good adverse selection. *Working Paper*.

Hwang, I. and F. Li (2017). Transparency of outside options in bargaining. *Journal of Economic Theory 167*(C), 116–147.

Kakhbod, A. and F. Song (2020). Dynamic price discovery: Transparency vs. information design. *Games and Economic Behavior 122*, 203–32.

Kakhbod, A. and F. Song (2024). Public vs. private offers with informed and forward-looking dealers. *working paper*.

Kaya, A. and Q. Liu (2015). Transparency and price formation. *Theoretical Economics 10*(2), 341–383.

Kaya, A. and S. Roy (2022). Market screening with limited records. *Games and Economic Behavior 132*, 106–132.

Kaya, A. and S. Roy (2024). Repeated trading: Transparency and market structure. *American Economic Review 114*, 2388–2435.

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

Vairo, M. and P. Dworczak (2025). Optimal transparency in over-the-counter markets. *working paper*.

■