Repeated Trading: Transparency and Market Structure

Ayça Kaya* and Santanu Roy†

October 10, 2022

Abstract

We analyze the effect of transparency of past trading outcomes in markets with incomplete information. A long lived seller who can sell every period and has private information about product quality (that is persistent over time) receives offers from short lived uninformed buyers that enter the market over time. Transparency allows buyers to observe the quantities sold in previous periods. The qualitative effect of transparency on market efficiency depends on the extent of buyer competition as well as the prior distribution of quality. When the market is characterized by high monopsony power (single buyer) in every period, transparency reduces welfare and efficiency if the ex ante expected quality is low but improves welfare and efficiency when the expected quality is high. The effect is exactly the reverse when the market is characterized by intra-period competition among multiple buyers; transparency increases (reduces) efficiency and welfare when the expected quality is low (high). Thus, when the prior belief about quality is adverse, competition makes transparency socially desirable but when the beliefs are better, it is the lack of competition that provides the rationale for improving market transparency. Competition generates more learning (finer screening) relative to monopsony when the market is transparent.

*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

Contents

1 Introduction

In many markets, sellers don't have the opportunity to form long-term relationships with their customers. Rather, they complete a single or infrequent transactions with a varying set of buyers. Many service sectors are like this: e.g. the work of a contractor, a real estate agent or a travel agent is rarely required frequently by the same customers. Similarly, sellers of items such as furniture, ceramics and other artisanal products, as well as some issuers of securities fit into this category. Often, in such markets sellers may be privately informed about the value they provide and may lack the ability to credibly communicate it. This may cause them to miss opportunities of mutually beneficial trade. Such inefficiencies may be partially resolved if newly arriving buyers can observe past trading behavior and use that to make inferences about the quality of the good or service on offer. On the flip side, the understanding that trading behavior is used in this manner by the arriving buyers gives the sellers incentives to distort their own sales decisions, creating a distinct cause of inefficiency which would not exist if past behavior is not observable. Thus it is an open question whether making past trading behavior observable would improve market efficiency or not, and how the answer may differ based on market conditions.

We explore this question focusing on the impact of the observability of past trading volumes abstracting from other channels via which information may flow to newly arriving customers. This is not just a theoretical curiosity. With the advancing technology, credibly providing this information to future buyers has become easier. And thus our analysis may shed light on policy questions as to whether such revelation should be mandated or not and what sorts of markets are likely to benefit from it.

Our analysis reveals that the impact of trade-volume transparency on market efficiency varies depending on two factors: (i) the degree of buyer competition, via its impact on the buyers' bargaining power, and (ii) the initial market perception of quality. Figure 1 summarizes our conclusions, with a + indicating cases where transparency promotes more efficient trading and a − indicating the cases where the opposite is true.

Our results also clarify the mechanisms via which transparency impacts market outcomes, highlighting two channels: first, transparency allows the market to screen the seller based on past trading patterns, and thus to some extent alleviates the lemons problem. Of course, this comes at a cost. The extent of this cost as well as the degree of the market's ability to learn depends on whether the buyer side of the market is competitive or whether they hold significant bargaining power, leading to the variation in the impact of

	low initial perception of quality	high initial perception of quality
with temporary monopsony	$-$	$+$
with buyer competition	$+$	$-$

Figure 1: Impact of transparency on gains from trade. A $+$ sign indicates an improvement and a $-$ sign indicates a reduction.

transparency. Second, the lemons problem is typically exacerbated when the buyers hold significant bargaining power. For these cases, transparency facilitates the resolution of the lemons problem by strengthening the lower quality sellers' bargaining power.

We obtain these results within the context of a formal model featuring a long-lived seller who has the capacity to build one unit of a good each period. The (binary, $\theta \in \{L, H\}$) quality of her output is exogenously fixed, is persistent and is her private information. The quality determines both the use value (v_θ) of the object to potential buyers and its cost of production (c_θ). These values satisfy

$$v_H > c_H > v_L > c_L.$$

Thus, the gains from trade is always positive and the cost of producing a high quality unit is higher than the use value of the low quality, so that the environment is potentially a lemons market. Each period, the seller meets one or more potential buyers. The buyers make simultaneous price offers. The seller can accept one or reject all. The buyers' bargaining power depends on the number of buyers meeting the seller at a time. In a market featuring intra-period monopsony, the seller meets a single buyer each period, while if the market features intra-period competition, she meets two or more buyers. An opaque market is one where arriving buyers observe nothing about the history of transactions. In a transparent market buyers observe the seller's history of trades, but not the trading prices. We consider perfect Bayesian equilibria. Formally, we say that transparency is welfare-reducing if all equilibria of an opaque market generate as much gains from trade as a transparent market with at least one equilibrium generating strictly more than any equilibrium of a transparent market. The case where transparency is welfare-improving is analogously defined.

To understand our results, first consider a market with intra-period buyer competition. When such a market is opaque, the lemons problem precludes trade of the high quality when market's perception is low. Specifically, lemons problem manifests when the initial probability assigned to high quality is less than μ^*, which is defined by

$$\mu^*(v_H - c_H) + (1-\mu^*)(v_L - c_H) = \underbrace{(1-\mu^*)(v_L - v_L)}_{=0}. \tag{1}$$

When the belief is above this cutoff, an opaque market is necessarily efficient. When such a market is transparent a costly learning equilibrium cannot be ruled out: regardless of the initial belief the market may coordinate on an equilibrium in which slow trading is interpreted to indicate high quality. In the extreme version of this equilibrium which features complete learning, the low quality seller trades efficiently, while the high quality trades a positive amount which is bounded away from efficiency. Thus, for initial beliefs where the opaque market would suffer from the lemons problem, this equilibrium is an unambiguous improvement while for higher initial beliefs it does worse. These observations are suggestive of our first main result:

Theorem 1 *Consider a market with intra-period buyer competition. Transparency is welfare-improving when $\mu_0 < \mu^*$ and is welfare-reducing when $\mu_0 > \mu^*$.*

We are able to derive further insights into how the market's ability to learn relates to gains from trade by constructing a class of equilibria for low initial beliefs. These are partial pooling equilibria, similar in structure to the complete-learning equilibrium discussed above, and which can be ranked with respect to the degree of informativeness when indexed by the probability with which the low quality seller pools with the high quality's slower trading path. We show that in this class, the more accurate is the market's learning, the higher the overall gains from trade. This indicates that the market's ability to learn accurately when there is intra-period buyer competition is a crucial factor rendering transparency beneficial when initial perception of quality is low. This ability is also what makes transparency undesirable when the initial perception of quality is high.

Next, consider a market with intra-period monopsony. The opaque version of such a market features a unique equilibrium in which, depending on the initial perception of quality, either only the low quality trades or both qualities trade efficiently. The cutoff belief above which the opaque market achieves efficiency is μ^{**} defined by

$$\mu^{**}(v_H - c_H) + (1 - \mu^{**})(v_L - c_H) = (1 - \mu^{**})(v_L - c_L). \tag{2}$$

Comparing (2) to (1) reveals that $\mu^{**} > \mu^*$: in an opaque market, the buyers' temporary monopsony power exacerbates the lemons problem. Intuitively, compared to competitive buyers, buyers with monopsony power can extract more surplus (the right-hand-side of equations (1) and (2)) from trading with the low quality seller, and thus are more reluctant to target the high quality seller. In such a market, transparency changes the nature of this *monopsony distortion*. Further, this monopsony distortion indirectly affects the market's ability to learn. As a result, the impact of transparency in such a market sharply contrasts its impact in a market with buyer competition:

Theorem 2 *Consider a market with intra-period monopsony. Transparency is welfare-reducing when $\mu_0 < \mu^*$ and is welfare-improving when $\mu_0 \in (\mu^*, \mu^{**})$. Transparency has no impact on market outcomes when $\mu > \mu^{**}$.*

On the one hand, transparency creates the option for the low quality seller to build a reputation for quality by mimicking trading patterns that the market would associate with high quality. When the market is uncertain about quality, this option strengthens the low quality seller's bargaining position and allows her to extract positive information rents, which in turn alleviates the monopsony distortion. This is why when the market's initial belief is between μ^* and μ^{**}, so that the only force precluding efficient trade in an opaque market is this distortion rather than the lack of information, transparency is welfare-improving.

On the other hand, the buyers' temporary monopsony power interferes with the market's ability to learn: market can credibly screen only if the high quality seller is willing to slow down trade in return for higher rents in the future. When buyers have the bargaining power, unlike the low quality seller, the high quality seller extracts no rents and is thus unwilling to wait. This severely restricts the market's ability to learn. When the initial belief is high (i.e. larger than μ^{**}) so that an efficient pooling equilibrium exists, this is a blessing: it allows the market to avoid coordinating on a costly learning equilibrium, and the outcome is necessarily efficient. When the initial belief is low, the market's inability to learn hurts its gains from trade—since, for the same reasons as in the case with intra-period buyer competition—finer screening is associated with higher overall surplus.

Further, because of the buyers' monopsony power, the low quality's payoff can be very low once revealed. This gives her very strong incentives to mimic the high quality's path. Thus credible screening may require very slow trading by the high quality, further reducing the overall gains from trade.

The analysis that supports Theorems 1 and 2 includes two types of results: (i) construction of specific equilibria and (ii) establishing bounds on equilibrium gains from trade without reference to these specific equilibria. It is worth highlighting two technical difficulties that our formal analysis overcomes. First, even though all equilibria we construct satisfy the so-called "skimming property," it is not possible to establish, from first principles, that this property must be satisfied by all equilibria.[1] This makes it challenging to establish the said payoff bounds. We are able to overcome this by referring to the properties of full equilibrium histories, and equilibrium probability distributions over these histories implied by equilibrium conditions. Second, the equilibrium construction, especially for the case with intra-period monopsony is non-trivial and all equilibria necessarily feature trading cycles, where the probability of trade conditional on quality varies across histories. Because of this, directly constructing equilibrium strategies becomes insurmountably tedious. We overcome this difficulty by resorting to techniques introduced by Abreau et al. (1990) whereby we modify the definitions of self-generating sets of payoffs to be appropriate for our specific setup and through these identify sets of equilibrium payoffs.

The rest of this paper is organized as follows: Section 1.1 discusses related literature. Section 2 introduces the formal model, Section 3 discusses the opaque market outcomes. Section 4 analyzes the impact of transparency in markets with intra-period buyer competition. Section 5 does the same for markets with intra-period monopsony.

1.1 Related literature

We study a dynamic lemons market where a seller has the ability to produce and sell *a unit in each period*, sequentially meeting *short-lived* potential buyers. Our specific focus is

[1] Skimming property is satisfied by an equilibrium if at any history, the high quality seller's reservation price is strictly higher than that of the low quality seller. This property holds in general when the seller has a single indivisible object to sell. In our context, this may fails because the low quality seller may have more to gain from waiting for frequent middling offers, and thus may be willing to reject certain offers that would be accepted by the high quality seller in spite of her current lower cost of production.

on the impact of *trade-volume transparency* on market outcomes.

There is an extensive literature studying dynamic lemons markets where the seller has one indivisible unit for sale where uninformed parties make repeated offers (e.g., Evans (1989), Vincent (1989), Deneckere and Liang (2006)), so that once they trade they leave the market.[2] Our notion of transparency (i.e. of past trading volumes) is moot in those models as trade can take place only once. Nevertheless, there are studies that explore the impact of different forms of transparency on market outcomes in such markets (e.g. Hörner and Vieille (2009) and Fuchs et al. (2016) consider observability of past rejected offers. Kim (2017) considers the observability of time-on-the-market.[3] The overarching conclusion in these studies is that transparency reduces the gains from trade due to equilibrium distortions to combat the low quality seller's strong incentives to mimic. In contrast, in our context, the results are more subtle and sensitive to market structure and perceptions. It is also worth noting that the equilibria of our transparent markets share some similarities to the equilibria typically observed in these single-sale models. In particular, in both, slower trading indicates higher quality. However, construction of our equilibria has additional challenges because of the need to keep track of continuation play after trade occurs.

A paper that studies repeated trading between two long-lived players is Hart and Tirole (1988) with a focus on the role of commitment to long-term contracts and not on notions of transparency. In addition, they focus on the case of independent valuations and thus the complications we face in equilibrium construction do not arise in their case.[4] A setting where trade does not immediately end the interaction is when the good for sale is divisible and can be traded incrementally over time. Gerardi et al. (2022) studies such an environment. Even though that model also shares the feature that in equilibrium slower trade is indicative of high quality, its focus is not on transparency but on the characterization of trading patterns. Finally, since in both of these papers, the bargaining takes place between two long-lived players, our notion of transparency is not relevant. In a recent study Fuchs

[2] See also Janssen and Roy (2002) for the analysis of a dynamic lemons market with decentralized equilibria. In addition to these papers who first studied this environment, several studies explore variations in this model. For instance, Vincent (1990) studies the impact of buyer competition, Ortner (forthcoming) studies the case where the seller's production cost may change, Fuchs and Skrzypacz (2019) poses a market design question and explores optimal times to allow/disallow trade in a lemons markets.

[3] The comparison of Noldeke and Damme (1990) which studies public offers and Swinkels (1999) which studies private offers in a labor market environment also sheds light on the role of transparency in dynamic lemons markets with single sale.

[4] An important result of Hart and Tirole (1988) is that in the repeated sale model with a long horizon and no commitment to long-term contracts, the uninformed side never learns due to the so-called ratchet effect. This is reminiscent of our result on limits on learning with intra-period monopsony.

et al. (2022) analyze the sale of one unit of a divisible asset with unknown quality (with a continuum of possible types) to a market of short-lived buyers and explore the impact of trade transparency. Similar to Janssen and Roy (2002), they study this question in a market with period-by-period decentralized equilibria. A seller in their model strategically chooses when to sell as well as whether to split the sale over time. This consideration is distinct from the strategic choices of the sellers in our model who have the ability to sell a unit each period. Fuchs et al. (2022) shows that if the trade is allowed to take place continuously, the seller's type is gradually screened with each seller selling at a specific instant, similar to the outcome of an indivisible single-sale model, and transparency of past trades does not play a role. In contrast, 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. In this case, without trade-volume transparency, the market's ability to screen is severely limited, and the qualitative impact of transparency on welfare is ambiguous.

There are a few papers that study repeated trade between a long-lived player and a sequence of short-lived players. Pei (forthcoming) considers a repeated sale environment, and shows that a long-lived seller cannot build reputation for producing only high quality when the sequence of short-lived buyers can observe only a bounded number of the seller's past actions. Similar to our question, Dilme (2022), explores the impact of the availability of information on past volumes of trade on efficiency, but in a Coasian environment where the short-lived informed buyers' valuations are independent of the long-lived seller's production costs, and shows that correctly designed noisy information about past trades creates more surplus than both full transparency and perfect confidentiality. Kaya and Roy (2022) also studies a repeated sale environment, and shows that the gains from trade can be non-monotone in the length of the records of past trades. In addition to focusing on a different question, that paper is confined to the case of competitive buyers and low initial perception of quality and thus cannot capture the subtler impacts of market structure and market perception on how transparency affects market outcomes. In a companion working paper (Kaya and Roy (2020)), we study how further transparency affects market outcomes, starting with a market where past trades are observable. Focusing on the case of intra-period monopsony and low initial beliefs, that paper shows that price observabilty can improve the outcomes by allowing the high quality seller to extract rents, playing a role similar to buyer competition in this paper.[5]

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 N potential trading partners (buyers) each with unit demand who makes take-it-or-leave-it price offers. Seller either accepts one of the buyers' offer and trades one unit at that price or rejects all prices. Regardless, all buyers leave the game, and the seller moves to the next period, meeting N new buyers.

We consider both the case when $N = 1$, so that each buyer has temporary monopsony power and the case where $N > 2$ so that the market features buyer competition. We refer to the first case as a "market with intra-period monopsony" and the second case as a "market with intra-period buyer competition."

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 $s = H$.

Assumption 1 *The seller is sufficiently patient: $\delta > \bar{\delta}$ where*

$$\bar{\delta} = \max\left\{\sqrt{\frac{v_L - c_L}{c_H - c_L}}, \frac{1}{1 + \frac{v_L - c_L}{v_H - c_L}}\right\}.$$

Histories In an **opaque market,** buyers observe only the calendar time and no particulars of the seller's trading history. In a **transparent market**, a public history at time t

[5] That paper precedes the current one. Some of the results (specifically Propositions 4 and 5) were originally included in that paper, and are now moved here. That paper now contains only the results concerning the impact of price transparency.

contains information about whether trade took place at each $t' < t$, and thus is an element of 2^{t-1}. Define $\mathcal{H}^\infty$ to be the set of all complete (terminal) public histories.

Let $\mathcal{H}_t$ be the set of all t-period public histories so that $\bigcup_{t=1}^\infty \mathcal{H}_t$ represent the set of public histories, with a typical element h. A private history of the seller includes the public histories, his type $\{H, L\}$ and the sequences of past realized price offers, including the currently active offer. Let $\mathcal{H}^S$ represent the set of all private histories of the seller. Given two histories h', h of respective lengths $t' < t$, we say that h is a continuation history of h' if the two histories coincide in the first t' periods. Fix a public history h', and let $\mathcal{H}^\infty(h')$ represent the set of all complete histories that are a continuation history of h'. Finally, let $h_\emptyset$ represent the null history.

Equilibrium We consider perfect Bayesian equilibria. A perfect Bayesian equilibrium consists of a strategy profile and a belief system. A behavior strategy of a buyer arriving at $t < \infty$ is a map $\sigma_t^B : \mathcal{H}_t \to \Delta\mathbb{R}_+$, specifying a probability distribution over price offers. A behavior strategy of the seller is a map $\sigma^S : \mathcal{H}^S \to [0, 1]$, specifying an acceptance probability for the currently active offer. A behavior strategy profile $(\sigma^S, \{\sigma_t^B\}_{t=1}^\infty)$ naturally induces a probability $\gamma_s(h|h')$, $s \in \{L, H\}$ that the seller of type s reaches history h conditional on having reached history h' of which h is a continuation history. A belief system is a map $\mu : \bigcup_{t=1}^\infty \mathcal{H}_t \to [0, 1]$ 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.[6]

For a given history h, let $q_i(h) \in \{0, 1\}$ be an indicator function representing whether trade took place in period i along history h. For $h \in \mathcal{H}$, and any $(t'-1)$-length history h' that h is a continuation history of, it is convenient to define $Q(h|h') = (1-\delta)\sum_{i=t'}^\infty \delta^{i-1} q_i(h)$ to be the expected discounted average trading volume along the continuation history h starting from history h'. Then, fixing an equilibrium and implied probabilities $\gamma_s(\cdot|\cdot)$, $s \in \{L, H\}$,

$$\overline{Q}_s(h') = \sum_{h \in \mathcal{H}^\infty(h')} \gamma_s(h|h') Q(h|h')$$

is the expected discounted average trading volume of type $s \in \{L, H\}$ in the continuation

[6]Formally, fixing a public history h, $\mu(h) = \mu_0\gamma_H(h)/[\mu_0\gamma_H(h) + (1-\mu_0)\gamma_L(h)]$, whenever the denominator is not zero.

equilibrium. Note that, given an equilibrium, the expected gains from trade is given by

$$\mu_0 \overline{Q}_H(h_\emptyset)(v_H - c_H) + (1 - \mu_0)\overline{Q}_L(h_\emptyset)(v_L - c_L). \tag{3}$$

Since price offers must be measurable with respect to public histories, the continuation payoff of the seller can be expressed as a function only of these histories and his private type. Fixing an equilibrium, throughout we let $V_s(h)$, $s \in \{L, H\}$, represent the type-s seller's continuation payoff at public history h. We express all payoffs in average per-period terms.

3 Opaque markets

If buyers observe no information about the trading history, the market has no tools to screen the seller, and therefore the market's belief is never updated. Thus, the seller acts myopically, as his continuation payoff cannot depend on his actions. In turn, the outcome in an opaque market is the period-by-period repetition of the static market outcome. The next proposition formally states the outcome in such markets. The proof is omitted.

Proposition 1 *In an opaque market the low quality trades with probability 1 in each period.*

- *In a market with intra-period buyer competition, the high quality never trades if* $\mu_0 < \mu^*$ *and trades with probability 1 in each period if* $\mu > \mu^*$*, where* μ^* *is defined in* (1).
- *In a market with intra-period monopsony, the high quality never trades if* $\mu_0 < \mu^{**}$ *and trades with probability 1 in each period if* $\mu > \mu^{**}$*, where* μ^{**} *is defined in* (2).

The outcome of an opaque, or equivalently, a static market is shaped by two economic forces. First is the lemons problem: the high quality cannot trade when the market's perception of quality is low. In our model, with intra-period buyer competition the cutoff belief below which lemons problem occurs is μ^*, which is less than the corresponding belief μ^{**} for a market with intra-period monopsony. The second economic force, "monopsony distortion," explains this discrepancy. A buyer with monopsony power can extract more surplus from a low quality seller, and therefore would attempt to trade with high quality only when the probability of it is higher.

Naturally, transparency changes how these two forces manifest. We next take up these issues in the context of first markets with intra-period buyer competition and then with intra-period monopsony.

4 Transparency with intra-period buyer competition

In this section we provide the analysis that support the claims in Theorem 1 which we reproduce here:

Theorem 1 *Consider a market with intra-period buyer competition. Transparency is welfare-improving when $\mu_0 < \mu^*$ and is welfare-reducing when $\mu_0 > \mu^*$.*

We start by constructing an equilibrium which achieves full separation of seller types. The equilibrium structure is stark but familiar (e.g. Admati and Perry (1987)). When the seller first enters the market, conditional on low quality, she immediately starts trading. Conditional on high quality, she waits $k(\delta)$ periods, and thereafter, trades each period with probability 1 at price v_H, where $k(\delta)$ is given by[7]

$$k(\delta) = \min\{k \mid \delta^k(v_H - c_L) < v_L - c_L\}. \tag{4}$$

Equation (4) specifies $k(\delta)$ to be the shortest wait that can deter mimicking by the low quality seller when at the conclusion of the pause, the buyers always offer v_H. The lengthy pause of trade is achieved by the threat of "belief punishments." If trade is not observed immediately at $t = 1$, the market learns that the seller is of high quality. But if trade unexpectedly occurs at some $1 < t \leq t(\delta)$, the belief goes down, and an equilibrium that delivers the high quality seller a payoff of 0 and the low quality seller a payoff of $v_L - c_L$ is played. This threat ensures that the seller's reservation price remains high enough during the first $k(\delta)$ periods of pause to preclude profitable trade.

Proposition 2 *A transparent market with intra period buyer competition admits an equilibrium in which after the first period the belief is either $\mu = 0$ or $\mu = 1$.*

The proof of Propositon 2 formally describes the equilibrium strategies and beliefs and verifies that they form an equilibrium.

We make two observations about the separating equilibrium of Proposition 2. First, one key factor that makes complete learning possible in a market with buyer competition

[7]The trading dynamics and off-equilibrium punishments in this construction are chosen to be stark in order to facilitate its description. There are other payoff-equivalent equilibria that feature trading cycles that are similar to the partial pooling equilibria discussed below, and formalized in Appendix A.4. Further, the off-path beliefs upon observing unexpected trade need not jump to 0. There are other equilibria starting with intermediate beliefs that deliver the high quality seller a payoff of 0 and the low quality seller a payoff of $v_L - c_L$ which can serve as punishment paths. See for example the construction in Appendix A.4.1.

is that the high quality seller is able to appropriate some of the gains from trade when the market's perception of quality is sufficiently high. This gives the seller incentives to incur costly pausing of trade (even when the buyers are willing to temporarily offer high prices) in order to receive rents in the future. Next, it is interesting to note how the structure of the separating equilibrium varies as the seller becomes more patient. Naturally, deterring mimicking by the low quality seller requires that as $\delta \to 1$, the length of the trading pause $k(\delta)$ must approach infinity. Nevertheless, the discounted amount of wait remains finite. Indeed, as $\delta \to 1$, $\delta^{k(\delta)} \to (v_H - c_L)/(v_L - c_L)$. This in particular shows that the differences between the opaque and transparent regimes persist even in the infinite patience limit.

The separating equilibrium of Proposition 2 exists regardless of the initial belief. Further, in this equilibrium, the low quality trades efficiently, while the high quality trades a positive expected amount which is distorted down from its efficient level. Thus, when $\mu^* < \mu_0$ this equilibrium improves upon the outcome of an opaque market, where due to the lemons problem, the high quality never trades. In contrast, when $\mu > \mu^*$ so that the opaque market's outcome is efficient, this equilibrium inefficiently reduces the high quality's trading volume.[8]

The claim of Theorem 1 for the case when $\mu_0 > \mu^*$ follows immediately because of the existence of an informative equilibrium which distorts the high quality's trading volume down, while in an opaque version of the market, the outcome is necessarily efficient. In other words, a transparent market can coordinate on an inefficient equilibrium, while this possibility is ruled out in an opaque market.

When $\mu_0 < \mu^*$, the argument is completed by the following lemma that establishes that the gains from trade in a transparent market is never less than $(1 - \mu_0)(v_L - c_L)$ which is the gains from trade in an opaque market.

Proposition 3 *In a transparent market with intra-period buyer competition the total surplus generated in any equilibrium is no less than* $(1 - \mu_0)(v_L - c_L)$.

The first step in proving Proposition 3 is to show that the low quality buyer captures all the trading surplus she generates, and thus trade never takes place below price v_L. Further, whenever the low quality seller's reservation price is less than v_L, she trades with probability 1. Then, at each history the low quality seller either trades at a price no less

[8]This insight is familiar from static signaling games (e.g. Spence (1973)), where often an inefficient separating equilibrium exists–and is selected by most common equilibrium refinements–even when pooling equilibria provide higher payoffs to each party.

than v_L or is better off rejecting such a price. Thus, the low quality seller's equilibrium payoff cannot be less than $v_L - c_L$. This leads to the lower bound on the overall gains from trade established in Lemma 3, because neither the high quality seller not the buyers can have negative equilibrium payoffs.

Partial pooling: For low initial beliefs ($\mu_0 < \mu^*$) the lemons problem precludes the transparent market from having an efficient pooling equilibrium. The transparent market also cannot replicate the no-learning outcome of the opaque market whereby the low quality trades efficiently while the high quality is excluded. Instead, in addition to the complete learning equilibrium of Proposition 2, there are equilibria with partial learning where the high quality trades a positive amount.[9] In particular, there is a class of equilibria in which the high quality seller always trades at the same price $P_H \in (c_H, v_H)$ and the expected discounted frequency of her trading is Q_H satisfying

$$v_L - c_L = Q_H(P_H - c_L).$$

Thanks to the latter condition, the low quality seller is indifferent between revealing herself (in return for a continuation payoff of $v_L - c_L$) versus mimicking the high quality throughout. To ensure that buyers are willing to offer exactly the price P_H, the low quality mimics the high quality seller's inefficient path with just sufficient probability so that along this path the buyers' expected value of the object is P_H.[10] Therefore, a higher P_H is associated with lower trading frequency for the high quality but also lower probability of pooling—and thus a higher expected frequency of trading—by the low quality. Interestingly, in spite of this trade-off, the overall gains from trade always increases as P_H increases regardless of how the intrinsic gains from trade ($v_\theta - c_\theta$) are ranked. This is because any adjustment that reduces the frequency of trading in return for a higher price that leaves the low quality seller indifferent renders the high quality seller strictly better off. Since the buyers in these equilibria receive no rents, the overall surplus increases as a result of such adjustments.

[9] See Appendix A.4 for the formal construction.

[10] The description of these equilibria is familiar and the constraints we present here are static. This obscures the difficulty of construction due to dynamics. In particular, it may not be possible to implement the frequency Q_H with an initial pause similar to that in the fully separating equilibrium of Propsosition 2. This is because, a large quantity of future trade is more attractive to the low quality seller due to her cost advantage, and because of that, the low quality seller's reservation price may exceed the high quality's towards the end of a long initial pause, giving the buyers the chance to target high quality only. We overcome this problem by constructing cyclical trading paths.

5 Transparency with intra-period monopsony

Similar to the case with buyer competition, transparency in a market with intra-period monopsony partially alleviates the lemon's problem by allowing the high quality to trade positive amounts even when initial belief is low ($\mu_0 < \mu^*$). However, the learning is limited due to the high quality's inability to extract rents and the resulting unwillingness to slow down trade to achieve reputation for high quality. It turns out that, combined with the low quality's weakened ability to extract rents when she is revealed, market's inability to screen is detrimental to the gains from trade for this range of beliefs. In fact, for low initial beliefs transparency is welfare-reducing.

In contrast, when the initial belief is very high $\mu_0 > \mu^{**}$ so that the opaque market would achieve efficiency, the market's inability to learn is a blessing as it eliminates the possibility that the transparent market could coordinate on an inefficient learning equilibrium. Indeed, for this range of beliefs, transparent market with intra-period monopsony necessarily achieves full efficiency just like an opaque market. Thus, transparency has no impact.

A novel impact of transparency appears for intermediate beliefs $\mu_0 \in (\mu^*, \mu^{**})$. Recall that for this range of beliefs, the opaque market still features no trade by the high quality even though there is no intrinsic lemons problem, because of the monopsony distortion: the buyers find it attractive to target only the low quality. Transparency improves the low quality seller's bargaining position, as she now has the option to mimic the high quality's trading path, bounding her equilibrium payoff from below. This makes it less attractive for buyers to target the low quality alone. Consequently, for this range of beliefs, transparency is unambiguously welfare-improving.

As discussed in the Introduction, these observations lead to the formal claims in Theorem 2 which we reproduce below. The rest of this section describes the formal analysis that supports these claims.

Theorem 2 *Consider a market with intra-period monopsony. Transparency is welfare-reducing when $\mu_0 < \mu^*$ and is welfare-imroving when $\mu_0 \in (\mu^*, \mu^{**})$. Transparency has no impact on market outcomes when $\mu > \mu^{**}$.*

We start by formally establishing the limits on learning in a transparent market with intra-period monopsony (Section 5.1). Then we discuss the welfare implications for different ranges of beliefs (Section 5.2). All results in Sections 5.1 and 5.2 are derived without reference to specific equilibria. In Section 5.3 we construct a class of equilibria and discuss the difficulties associated with construction.

5.1 Learning in a transparent market with intra-period monopsony

As is the case in the equilibrium constructed in Proposition 2, with intra-period competition, the high quality seller can extract rents when the market's perception of quality becomes high. As discussed, this is a crucial feature that allows complete learning in that setting. In a market with intra-period monopsony, this feature fails, as formalized in the following lemma.

Lemma 1 *In any equilibrium of the transparent market with intra-period monopsony, at any history h, $V_H(h) = 0$. Therefore, the high quality seller accepts any offer that exceeds c_H.*

An immediate implication of Lemma 1 is that a buyer arriving with belief $\mu > \mu^*$ is guaranteed a strictly positive payoff (which he can achieve by offering a price slightly above c_H). Therefore, such a buyer would never make a losing offer. Consequently, neither overall trade, nor trade with high quality can be significantly slowed down. This implies that, if the low quality seller finds herself in a market with high average quality, her payoff is necessarily large, as demonstrated in the following lemma.

Lemma 2 *In any equilibrium of the transparent market with intra-period monopsony, if $\mu(h) > \mu^*$, then $V_L(h) \geq \delta(c_H - c_L)$.*

In contrast, when the market's belief is below μ^* trade must eventually take place at a price below c_H with positive probability, revealing the low quality seller. Because of this and since once revealed, the low quality seller cannot receive a continuation payoff exceeding $v_L - c_L$, her payoff is bounded from above.

Lemma 3 *In any equilibrium of the transparent market with intra-period monopsony, if $\mu(h) < \mu^*$, then $V_L(h) \leq v_L - c_L$.*

Lemmas 2 and 3 together limit the market's ability to screen the seller. If, along an equilibrium path, the market's belief crosses the threshold μ^* either from above or below, there must be a history at which the seller makes a choice that puts him on either side of it. Importantly, it must be optimal for the low quality seller to make the choice that puts him below the threshold. The large discrepancy between the payoffs of the low quality seller on either side contradicts the optimality of such a choice. This leads to the following formal result on the limits of screening.

Proposition 4 *Consider a transparent market with intra-period monopsony. Fix an arbitrary equilibrium.*

- *If $\mu_0 < \mu^*$, then at any equilibrium path history h, $\mu(h) \leq \mu^*$.*
- *If $\mu_0 = \mu^*$, then at any equilibrium path history h, $\mu(h) = \mu^*$.*
- *If $\mu^* < \mu_0 \leq \mu^{**}$, at any equilibrium path history h, $\mu(h) \geq \mu^*$.*

5.2 Transparency and gains from trade with intra-period monopsony

In this section we study the impact of transparency on the efficiency of trade in markets with intra-period monopsony. In the next two subsections, we separately take up the cases of low and high initial beliefs.

5.2.1 Low initial beliefs

We show that when the initial belief is low, the transparent market can never do better than an opaque market. In fact, it can do much worse.

The intuition is best understood by considering a specific set of equilibria even though the formal results do not rely on this construction. When $\mu_0 < \mu^*$, the transparent market with intra period monopsony admits partial pooling equilibria similar to those in the market with buyer competition. In these equilibria, in the first period the low quality seller reveals herself by trading at price v_L with positive probability. If she does, she continues to trade with probability 1 each period thereafter. With the remaining probability she pools with the high quality seller, who does not trade in the first period, and then trades at an expected discounted average frequency $\overline{Q}_H$ thereafter.[11] Unlike that setting however, two

[11] This path of equilibrium is familiar from single sale models. However, in a repeated sale environment the construction of such equilibria is a lot more intricate. In particular, along the pooling path, i.e. after the belief jumps to μ^*, the probability of an offer of c_H cannot be independent of the history. Because if it were, the low quality seller's reservation price would be c_L, and consequently, each buyer would strictly prefer to target only the low quality seller, rather than trading with both qualities at price c_H. Because of this, the equilibrium path must always be cyclical. See Appendix B.4 for the formal construction.

factors preclude a transparent market from improving upon the level of gains from trade in an opaque market. First, due to the buyers' monopsony power, the low quality seller's payoff from revealing herself can be very low, strengthening her incentives to mimic the high quality's trading path. Consequently, to deter mimicking, the high quality's trading frequency must be severely restricted. At the extreme, if in an equilibrium within this class, the low quality seller anticipates receiving no rents once her type is revealed, then the high quality's trading frequency $\overline{Q}_H$ must satisfy

$$(1-\delta)(v_L - c_L) = \overline{Q}_H(c_H - c_L), \tag{5}$$

resulting in overall gains from trade of $(1-\mu_0)(v_L - c_L)(1-\delta)$, as opposed to the gains of $(1-\mu_0)(v_L - c_L)$ in an opaque market.Second, by Proposition 4 the market screening will necessarily be coarse, while finer screening would have been associated with higher overall surplus for analogous reasons to those discussed in the case of markets with buyer competition.

The intuition gained from this class of equilibria applies more generally. In fact, using arguments based only on equilibrium conditions, and independent of specific equilibria, we are able to establish an upper bound on the gains from trade in a transparent market with intra-period monopsony in Proposition 5. Combined with the construction of lower-welfare equilibria, this establishes that when $\mu_0 < \mu^*$, transparency is welfare reducing, as claimed in Theorem 2.

Proposition 5 *If $\mu_0 < \mu^*$, the expected average gains from trade in a transparent market with intra-period monopsony is no larger than*

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

For the class of equilibria discussed above, the bound in Proposition 5 follows by simple accounting as follows: the low quality seller must pool with the high quality with probability $\frac{\mu_0}{1-\mu_0}\frac{1-\mu^*}{\mu^*}$ so that the average quality conditional on pooling on the slower trading path is c_H. Thus, the low quality's trading frequency cannot exceed

$$\overline{Q}_L \equiv \frac{\mu_0}{1-\mu_0}\frac{1-\mu^*}{\mu^*}\overline{Q}_H + 1 - \frac{\mu_0}{1-\mu_0}\frac{1-\mu^*}{\mu^*}.$$

Further, $\overline{Q}_H$ is bounded by the low quality seller's incentives to mimic:[12]

$$v_L - c_L \geq \overline{Q}_H(c_H - c_L).$$

Using the latter two inequalities to bound the total gains from trade $\mu_0(v_H - c_H) + (1 - \mu_0)(v_L - c_L)$ and substituting the definition of μ^* yields the bound in Proposition 5. The general proof does not refer to a specific equilibrium structure, but uses similar arguments, along with the fact that starting from $\mu_0 < \mu^*$, when trade takes place for the first time, the belief either must jump to 0 or μ^*, where the latter assertion follows from Proposition 4.

5.2.2 High initial beliefs

It is once again instructive to first discuss a specific class of equilibria for the case when $\mu_0 \in (\mu^*, \mu^{**})$. These equilibria also feature partial pooling, but unlike in the case of low initial beliefs, in this case, the high quality seller must be revealed with positive probability. In particular, if μ_0 is sufficiently low (close to μ^*), there exists an equilibrium where in the first period the buyer randomizes between two offers: v_L and c_H. The former is accepted with probability 1 by only the low quality seller, while the latter is accepted with probability 1 by both types. Thus, low quality trades with probability 1, and failure to trade reveals high quality. The first buyer's randomization is such that, upon trade in the first period, belief updates to μ^*, and along this path, trade always takes place at price c_H with average discounted frequency, say $\overline{Q}_L$.[13] To deter the low quality seller from mimicking the high quality seller by rejecting offers in the first period, the frequency $\overline{Q}_L$ must satisfy

$$\delta\overline{Q}_L(c_H - c_L) + (1 - \delta)(v_L - c_L) \geq \delta(c_H - c_L).$$

The right hand side of this inequality is what the low quality seller can receive by mimicking the high type and rejecting a price offer in the last period of screening. The left hand side is what he would get if he trades at price v_L in that period and then trades at frequency $\overline{Q}_L$ at price c_H from then on.[14] Mimicking a path where high quality is exactly identified is very lucrative for the low quality seller, and deterring such mimicking requires that the

[12] The left-hand-side reflects the highest payoff the low quality seller can receive upon revealing herself.

[13] For higher beliefs in this range, market screening may take multiple (finite number of) periods. The construction of these equilibria is quite delicate. See Appendix B.4 for details.

[14] This is an inequality, because the low quality's reservation price is less than v_L at such a history. See Appendix B.4 for the full characterization.

alternative (in this case, the pooling outcome) also generates a high payoff. This bounds the trading frequency from *below.*

Proposition 6 *If $\mu^{**} > \mu_0 > \mu^*$, the expected gains from trade is no less than*

$$\mu_0(v_H - c_H) + (1 - \mu_0)(v_L - c_L)$$

$$-$$

$$(1 - \delta)[(1 - \mu_0)(c_H - c_L) - \mu_0(v_H - c_H)],$$

which strictly exceeds $(1 - \mu_0)(v_L - c_L)$.

Further, if $\mu > \mu^{**}$, the expected gains from trade is $\mu_0(v_H - c_H) + (1 - \mu_0)(v_L - c_L)$.

It is interesting to note that when $\mu_0 \in (\mu^*, \mu^{**})$, the lower bound on equilibrium gains from trade approaches full efficiency as $\delta \to 1$. Recall that for this range of beliefs the opaque market is inefficient, not because of the lemon's problem per se, but because of the monopsony distortion. It is intuitive that as the low quality seller becomes more patient, the strengthening of her bargaining power due to transparency becomes extreme, eliminating the monopsony distortion.

5.3 Equilibrium structure with intra-period monopsony

In this section, we construct a class of equilibria for the transparent markets with intra-period monopsony. Our goal is not to characterize all equilibria. Instead, this construction, in addition to establishing existence, completes our analysis by formally demonstrating that the gains from trade can be strictly below the upper bound established in Proposition 5.[15] Also in this discussion we further highlight the difficulties with equilibrium construction.

Self-generating sets of payoffs: Direct characterization of equilibrium strategies is cumbersome because all equilibria necessarily feature trading cycles which are potentially non-stationary. In our discussion of specific equilibria below we highlight the aspects of the environment that necessitate such cycles. Because of this difficulty, our strategy of characterization appeals to dynamic programming arguments similar to the techniques developed in Abreau et al. (1990) and Fudenberg et al. (1994).

[15]We need this additional analysis since that bound exactly matches the gains from trade in an opaque market.

Informally, for each belief μ, we specify sets $\mathcal{U}_\mu$ of potential equilibrium payoffs for the low quality seller and let $\mathcal{U} \equiv \{\mathcal{U}_\mu\}_{\mu\in[0,1]}$.[16] We say that $\mathcal{U}$ is **self-generating** if for each belief μ and each payoff $U \in \mathcal{U}_\mu$ one can choose continuation beliefs μ^A, μ^R, after "trade" and "no trade," respectively, and continuation payoffs $U^A \in \mathcal{U}_{\mu^A}$ and $U^R \in \mathcal{U}_{\mu^R}$ of the low quality seller such that

1. U is an equilibrium payoff of the static game obtained by replacing the continuation game with the continuation payoffs U^A, U^R, *and*
2. the strategies in such equilibrium justifies μ^A, μ^R as continuation beliefs.[17,18]

Then, when $\mathcal{U}$ is self-generating, for each μ and $U \in \mathcal{U}_\mu$, it is straightforward to iteratively construct an equilibrium where the low quality seller's payoff is U. We provide the formal definitions and results in Appendix B.4.

In what follows, we describe equilibria in three steps: First, we discuss equilibria starting with belief exactly μ^*. Second and third steps take up lower and higher initial beliefs.

5.3.1 Equilibria when belief is μ^*

By Proposition 4, if belief reaches μ^*, it is never updated and thus the two types of the seller must always trade at the same frequency, and then necessarily at price c_H. One property of such equilibria then is that buyers never make positive profits. There are many such equilibria that differ with respect to their average frequency at which this price offer is made, and thus trade takes place. Let $\tilde{Q}$ be the average frequency of trade along such a continuation path equilibrium. First, we note that the stationary strategy where each buyer would offer c_H with probability $\tilde{Q}$ and would make a losing offer with the remaining probability cannot be part of an equilibrium: If the buyers' strategy was history-independent in this manner, the seller's continuation payoff upon trade would be the same as her continuation payoff upon failure to trade. Therefore, the low quality seller's

[16]Since the high quality seller's payoff is always zero and the buyers are short-lived, it suffices to focus on the payoffs of the low quality seller.

[17]That is, if acceptance (respectively, rejection) is on the equilibrium path, μ^A (respectively, μ^R) is obtained from the equilibrium strategies (and initial belief μ) using Bayes rule.

[18]Since our goal is not to characterize all possible equilibria, in each case, for item (i) we specify sufficient conditions for equilibrium and not necessary conditions. See Appendix B.4 for details.

reservation price would be c_L. But then, a buyer arriving at such a history would have a profitable deviation to making an offer in (c_L, v_L), guaranteeing him a positive payoff.

The above thought exercise illustrates why the equilibrium probability of trade in the current period must vary with the trading history. More generally, in any such equilibrium, the low quality seller's reservation price must be no less than v_L. It must also be no larger than c_H because otherwise the buyer would be able to target only the high quality seller and thereby guarantee a positive payoff. Thus, it is necessary that

$$(1-\delta)(v_L - c_L) \leq \delta[U^R - U^A] \leq (1-\delta)(c_H - c_L), \tag{6}$$

where U^A and U^R, respectively, represent the continuation payoffs of the low quality seller after trade (acceptance) and after no trade (rejection) at history h. An equilibrium where the low quality seller's payoff is $\tilde{U}$ must feature a probability α with which the offer c_H is made in the current period and continuation values U^A, U^R such that

$$\tilde{U} = \alpha\left[(1-\delta)(c_H - c_L) + \delta U^A\right] + (1-\alpha)\delta U^R$$

and U^A, U^R must satisfy (6) and must be equilibrium payoffs themselves. In the appendix, we show that this is possible for any

$$\tilde{U} \in \mathcal{U}_{\mu^*} = \left[(1-\delta)(v_L - c_L), c_H - c_L - (1-\delta)(v_L - c_L)\right].$$

In our construction, the values that are close to the two ends of this interval are enforced using belief punishments or rewards. In particular, an equilibrium that delivers a payoff close to the lower end of the range to the low quality seller features zero probability of trade, and the low quality seller's reservation price is kept above v_L because unexpected trade is interpreted as coming from the low quality seller only. In contrast, payoffs on the upper end feature trade with probability 1, and off-path rejection is interpreted as coming from the high quality seller only.[19] Thus, construction of equilibria for initial belief μ^* relies on those for other beliefs. In turn, those equilibria rely on the construction of equilibria with initial belief μ^*, which forms a building block for all others, as clarified below.

[19] For some parameter constellations, it is possible to construct self-generating payoff sets that are all enforced by $\alpha \in (0,1)$, so that belief punishments / rewards are not necessary. For other parameter values, these are unavoidable.

5.3.2 Equilibria when $\mu_0 < \mu^*$

The multiplicity of equilibria when $\mu = \mu^*$ allows for multiple equilibria when $\mu = 0$, using belief rewards. In particular, when $\mu = 0$, fixed price equilibria where each buyer offers a specific P with $P \in [c_L, v_L]$ with probability 1 and is accepted with probability 1 can be sustained by choosing off-path belief $\mu^R = \mu^*$ and $U^R = (P - c_L)/\delta$ ensuring that the low quality seller's reservation price is exactly P. Thus, $\mathcal{U}_0 = [0, v_L - c_L]$ are equilibrium payoffs for the low quality seller when the belief is 0.

When $\mu_0 \in (0, \mu^*)$ there is a simple type of equilibrium that spans the full range of equilibrium payoffs and gains from trade. In this equilibrium, the market screens the seller in the first period by offering a price of v_L. This price is rejected by the high quality seller with probability 1, and rejected by the low quality seller with just the right probability so that upon failure of trade the belief is μ^*. Then, $U^A \in \mathcal{U}_0, U^R \in \mathcal{U}_{\mu^*}$ are chosen such that

$$(1 - \delta)(v_L - c_L) + \delta U^A = \delta U^R.$$

Note that when $U^A < v_L$ the gains from trade is strictly less than the upper bound $(1 - \mu_0)(v_L - c_L)$ established in Proposition 5. In this class of equilibria, with probability β which satisfies

$$\frac{\mu^*}{1 - \mu^*} = \frac{\mu_0}{1 - \mu_0}(1 - \beta),$$

the low quality seller trades efficiently. With the rest of the probability she pools with the high quality on an inefficient path along which the average frequency, say Q, of trade is pinned down by the low quality seller's indifference condition in the initial period:

$$(1 - \delta)(v_L - c_L) + \delta U^A = Q(c_H - c_L).$$

The upper bound on gains from trade is attained when $U^A = v_L - c_L$. For any other U^A, the pooling path features less trade (smaller Q) and thus the gains from trade is smaller.

5.3.3 Equilibria when $\mu > \mu^*$

Unlike in the case for low initial beliefs, we construct a single equilibrium for almost all initial beliefs in this range.[20] These equilibria feature finitely many periods of screening

[20] We suspect these may be the unique equilibria here. For a countably many initial belief that constitute cutoffs in the screening process, there are multiple equilibria as described below.

during which the belief either jumps to 1 or declines. The declining belief path converges to μ^*. The formal construction in the appendix still relies on self-generation arguments. Here we describe the convergence path which is monotone and does not feature cycles.

First consider an initial belief μ_0 that is higher than but close to μ^*. In this case, there is an equilibrium as follows: in the first period, the buyer randomizes between c_H and the low quality's reservation price, the former is accepted with probability 1 by both types, the latter is accepted with probability 1 by the low type only. The probability α with which c_H is offered is such that

$$\frac{c_H - v_L}{v_H - c_H} \equiv \frac{\mu^*}{1-\mu^*} = \frac{\mu_0}{1-\mu_0}\alpha,$$

so that upon trade in the first period the belief updates to μ^*. Also of course, upon failure to trade, the belief is 1. Further, the buyer's indifference requires that

$$\frac{\mu_0}{1-\mu_0} = \frac{c_H - P_L}{v_H - c_H}, \tag{7}$$

where P_L is the low quality seller's reservation price, which is thus uniquely pinned down. Low quality seller's reservation price, by definition, must satisfy

$$(1-\delta)(P_L - c_L) + \delta U^A = \delta \underbrace{(c_H - c_L)}_{U^R}, \tag{8}$$

where $U^R = c_H - c_L$, because upon rejection the belief updates to 1, and the unique equilibrium is for each buyer to offer c_H. Since it must be that $U^A \in \mathcal{U}_{\mu^*}$ and $\mathcal{U}_{\mu^*}$ is bounded from above, there is a lower bound on P_L that can satisfy (8), and thus there is an upper bound, say μ^1, on μ_0 that can satisfy (7). Thus, this construction is possible only when initial belief μ_0 is no larger than μ^1, so defined. Note that μ^1 is uniquely pinned down and is the highest belief from which the beliefs can reach μ^* in one step. The payoff of the low quality seller in this equilibrium is[21]

$$U_0 \equiv \delta(c_H - c_L) + (1-\delta)\alpha(c_H - P_L) = \delta(c_H - c_L) + (1-\delta)(c_H - v_L),$$

where the last equality follows by substituting α from the previous equation. Note that this payoff is independent of the initial belief μ_0 as long as it is in the interval (μ^*, μ^1).

[21] This is because, if the offer is P_L the low quality seller is indifferent between accepting and rejecting, and thus, in that case her payoff is $\delta(c_H - c_L)$ while her payoff is higher by $c_H - P_L$ if the offer is c_H.

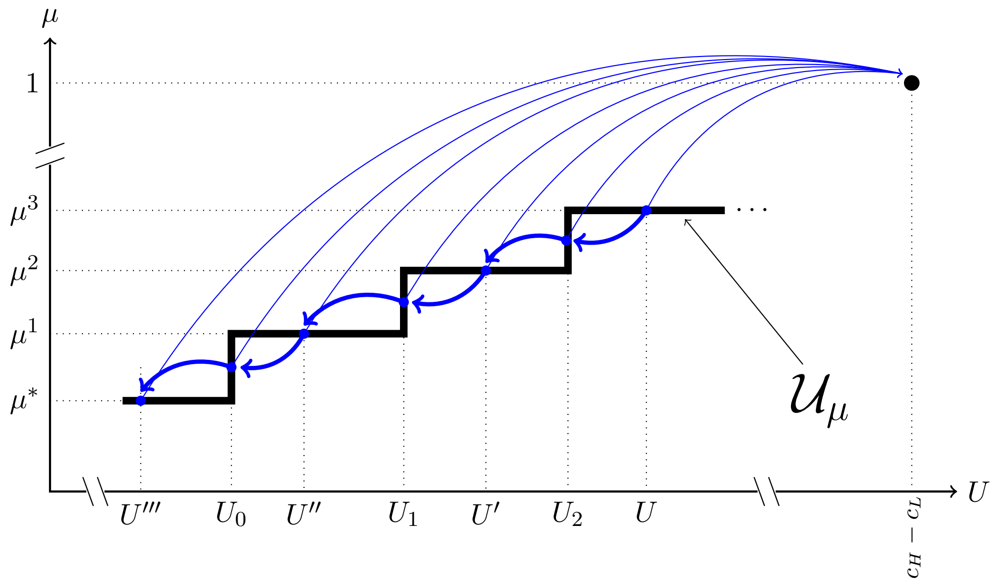

Figure 2: The heavy black step function maps a subset of beliefs $[\mu^*, \mu^{**})$ on the vertical axis to equilibrium payoffs the low quality seller at those beliefs on the horizontal axis. The (blue) heavy arrows show how the (belief, payoff) pairs change after trade. The (blue) light arrows indicate that after each failure of trade, the belief reaches 1 and the consequently, the low quality seller's payoff reaches $c_H - c_L$.

When the belief is exactly μ^1, there are a continuum of equilibria that are similar in structure to the one described so far. These equilibria differ by the probability α with which the buyer offers the high price c_H. In all these equilibria, the low quality's reservation price simultaneously satisfies (7) when $\mu_0 = \mu^1$ and (8) when $U^A = U_0$. The highest payoff, say U_1, that the low quality seller can obtain in this construction is obtained when $\alpha = 1$.

As suggested by the above construction, and as depicted in Figure 2, the map between beliefs in (μ^*, μ^{**}) and the equilibrium payoffs of the low quality seller forms a step-function. Similarly to μ^1, μ^k for $k > 1$ in this figure, is the highest belief from which the belief can reach μ_{k-1} in one step. The (blue) arrows in Figure 2 represent the path of a specific equilibrium that delivers the low quality agent a payoff of U starting with belief μ_3. In this equilibrium, convergence to μ^* takes 6 periods of screening, during which the low quality seller trades with probability 1. Thus, conditional on low quality, the belief-payoff combinations follow the declining path indicated by the heavier (blue) arrows. Conditional on high quality, the path depends on the realization of the buyers' randomization. As soon as a buyer fails to make the offer of c_H, the belief jumps to 1, as indicated by the lighter (blue) arrows in Figure 2. Note that these outcomes are off-the-equilibrium path conditional on low quality.

A Omitted proofs for Section 4

A.1 Proof of Proposition 2

Throughout this section we assume that

$$(1-\delta) < \delta \frac{v_L - c_L}{v_H - c_L}. \tag{9}$$

We construct an equilibrium that has the following equilibrium path:

- Conditional on low quality, trade takes place with probability 1 at each period $t = 1, 2, \cdots$.
- Conditional on high quality, trade takes place with probability zero at $t = 1, 2, \cdots, k(\delta)$, and takes place with probability 1 at $t = k(\delta), k(\delta)+1, \cdots$.

In particular, this equilibrium features no trade in periods $t = 2, \cdots, k(\delta)$ if there was no trade at $t = 1$. In our construction, if there is unexpected trade during these periods following no trade at $t = 1$, all the subsequent buyers hold belief $\mu' = 0$, regardless of the other details of the history, and make offers of v_L.[22]

We start by formally describing the strategies. For this purpose, define $h^s_\emptyset$ to be the s-length history featuring no trade and h^s_I be the s-length history featuring trade in each period. Let (h, A) be the continuation history of h obtained by adding a period of trade. We partition non-null histories into the following comprehensive and mutually exclusive cases:

- Case 1: features trade in the first period (other specifics of the history are irrelevant)
- Case 2: $h = h^s_\emptyset$ for some $1 \leq s < k(\delta)$, i.e. a history of length shorter than $k(\delta)$ periods that features no trade.
- Case 3: $h = h^{k(\delta)}_\emptyset$ or is any continuation history of $h^{k(\delta)}_\emptyset$, i.e. a history of length $k(\delta)$ that features no trade or a continuation history thereof.

[22]Alternative possible with belief updating to μ^*. See Section A.4.1 for construction of such an equilibrium.

- Case 4: h is a continuation history of $(h_\emptyset^s, A)$ for some $s < k(\delta)$, i.e. a history that features first no trade for fewer than $k(\delta)$ periods and then trade. These histories are off the equilibrium path.

Buyer strategies: Consider the following buyer strategies:

- At $t = 1$, offer v_L.
- At $t \geq 2$,
  - Case 1: offer v_L.
  - Case 2: make a losing offer (an offer less than v_L.)
  - Case 3: offer v_H.
  - Case 4: offer v_L.

Beliefs: Buyers have the following beliefs:

- Case 1: $\mu(h) = 0$.
- Cases 2 and 3: $\mu(h) = 1$.
- Case 4: $\mu(h) = 0$.

Seller strategies: Each type of the seller uses a reservation price strategy $P_\theta(h)$, $\theta = L, H$, satisfying

- At $t = 1$,

$$(1-\delta)(P_H(h) - c_H) = \delta^{k(\delta)}(v_H - c_H)$$
$$(1-\delta)(P_L(h) - c_L) + \delta(v_L - c_L) = \delta^{k(\delta)}(v_H - c_L)$$

  We note that, by choice of $k(\delta)$, $P_L(h) < v_L$ and $\delta^{k(\delta)-1} > \frac{v_L - c_L}{c_H - c_L}$. The latter is equivalent to $\delta^{k(\delta)} > \delta \frac{v_L - c_L}{c_H - c_L}$, which implies that

$$(1-\delta)(P_H(h) - c_H) > \delta \frac{v_L - c_L}{v_H - c_L}(v_H - c_H).$$

  By Assumption (9) this implies that $P_H(h) > v_H$.

- At $t \geq 2$
  - Case 1, 3 and 4: $P_\theta(h) = c_\theta$, $\theta = H, L$.
  - Case 2: If $h = h_\emptyset^s$ for some $1 \leq s < k(\delta)$,

$$\begin{aligned}(1-\delta)(P_H(h) - c_H) &= \delta^{k(\delta)-s}(v_H - c_H) \\ (1-\delta)(P_L(h) - c_L) + \delta(v_L - c_L) &= \delta^{k(\delta)-s}(v_H - c_L)\end{aligned}$$

    We note that in this case, $P_H(h) > v_H$ and by choice of $k(\delta)$, $P_L(h) > v_L$.

At any history each type of the buyer accepts any offer that weakly exceed their reservation price with probability 1, and rejects all other offers.

Optimality of buyer strategies To establish the optimality of buyer strategies we take up the following cases:

- At $t = 1$, the high quality's reservation price strictly exceeds v_H, while the low quality's reservation price is strictly less than v_L. Thus, there is a bidding equilibrium where each buyer bids v_L.
- At $t \geq 2$,
  - Case 1: Since the belief is $\mu(h) = 0$, it is a bidding equilibrium for the buyers to bid v_L.
  - Case 2: Since $\mu(h) = 1$, but $P_H(h) > v_H$, it is a bidding equilibrium for the buyers to each make a losing offer.
  - Case 3: Since $\mu(h) = 1$, and $P_H(h) = c_H < v_H$, it is a bidding equilibrium for the buyers to bid v_H.
  - Case 4: Since $\mu(h) = 0$ and $P_L(h) = c_L < v_L$, it is a bidding equilibrium for the buyers to bid v_L.

Optimality of seller strategies: The reservation prices are calculated using buyer offer strategies. Thus, accepting offers weakly exceeding them and rejecting offers that are strictly less are best responses to those strategies.

Belief consistency: Follows trivially by the equilibrium strategies.

A.2 Proof of Proposition 3

Proposition 3 relies on the following lemma:

Lemma 4 *Assume that $N > 1$. In any equilibrium, at any history h, on or off the equilibrium path, the following are true:*

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

2. *Trading price is never less than v_L.*

Proof of Lemma 4. Fix a history h with $\mu(h) < 1$.

1. Assume that $P_L(h) < v_L$. Suppose the low quality trades with probability less than 1 at this history. Then, necessarily each buyer's expected payoff is positive, because for $\varepsilon > 0$ and small, an offer of $P_L(h) + \varepsilon$ will be the winning offer with positive probability and it will be accepted with probability 1 by the low quality seller. This implies that (i) the lower bound $\underline{P}$ of the support of each buyer's bid distribution is the same because otherwise at least one buyer would be making an offer that wins with zero probability, (ii) $\underline{P} \leq P_L(h)$ because the low quality trades with probability less than 1; and (iii) each buyer offers $\underline{P}$ with an "atom," i.e. if F_i is buyer i's bid distribution, $F_i(\underline{P}) > 0$ for each i.

   Since each buyer offers $\underline{P}$ with positive probability, there is a positive probability of ties. Conditional on a tie, there exists at least one buyer who trades with probability at most 1/2. Without loss of generality, label this buyer, buyer 1.

   Now, we claim that, there exists $\varepsilon > 0$ such that buyer 1 receives a strictly higher payoff by offering $\underline{P} + \varepsilon$ than by offering $\underline{P}$. Let γ_θ be the probability with which the seller with quality $\theta \in \{L, H\}$ accepts $\underline{P}$ when it is the highest offer. [note that this event has positive probability] Let $\gamma_\theta^\varepsilon$ be the corresponding probability when $\underline{P} + \varepsilon$ is offered. Note that for each θ, $\gamma_\theta \leq \gamma_\theta^\varepsilon$. Then, conditional on winning at price $\underline{P}$, the payoff is

$$\Pi \equiv \mu(h)\gamma_H(v_H - \underline{P}) + (1 - \mu(h))\gamma_L(v_L - \underline{P}) > 0.$$

   Conditional on winning at price $\underline{P} + \varepsilon$ the payoff is

$$\Pi^\varepsilon \equiv \mu(h)\gamma_H^\varepsilon(v_H - \underline{P} - \varepsilon) + (1 - \mu(h))\gamma_L^\varepsilon(v_L - \underline{P} - \varepsilon).$$

Note that for any $\delta > 0$, there exists $\bar{\varepsilon} > 0$ such that for any $0 < \varepsilon < \bar{\varepsilon}$, $\Pi - \Pi^{\varepsilon} < \delta$ and $\Pi_{\varepsilon} > 0$. Finally, note that there exist $\omega > 0$ such that for any $\varepsilon > 0$, buyer 1's probability of winning with offer $\underline{P} + \varepsilon$ exceeds his probability of winning with offer $\underline{P}$ by ω, which establishes the claim.

Thus, there cannot be any equilibrium where the low quality trades with probability less than 1 provided that $P_L(h) < v_L$.

2. Suppose at history h, trade takes place with positive probability at a price $P < v_L$. Then each buyer's expected payoff is positive. Then, by the same arguments as above, one can construct a profitable deviation from the lowest price offer in the support of the bid distribution.

■

Now we are ready to prove Proposition 3.

Proof of Proposition 3. Let $\underline{V}_L$ be the infimum of the continuation equilibrium payoffs of the low quality seller on or off the equilibrium path. Fix $\varepsilon > 0$ and choose h such that $V_L(h) < \underline{V}_L + \varepsilon$.

If at history h, the low quality seller trades with probability less than 1, then, by Lemma 4 it must be that $P_L(h) \geq v_L$. Equivalently,

$$(1 - \delta)(v_L - c_L) + \delta V_A \leq \delta V_R,$$

where V_A and V_R are continuation payoffs after trade and no trade, respectively. Since $V_A \geq \underline{V}_L$, we have

$$(1 - \delta)(v_L - c_L) + \delta \underline{V}_L \leq \delta V_R.$$

Further, since the low quality seller always has the option to reject any offer at history h,

$$V_L(h) = \underline{V}_L + \varepsilon_2 \geq \delta V_R,$$

where $\varepsilon_2 \geq 0$ is chosen such that $V_L(h) = \underline{V}_L + \varepsilon_2$ and thus, $\varepsilon_2 < \varepsilon$. Combining the latter two inequalities yields

$$(1 - \delta)(v_L - c_L) + \delta \underline{V}_L \leq \underline{V}_L + \varepsilon_2 \Leftrightarrow v_L - c_L \leq \underline{V}_L + \varepsilon_2.$$

Next, suppose the low quality trades with probability 1 at h. Then, again by Lemma 4 (since price is at least v_L),

$$\underline{V}_L + \varepsilon_2 \geq (1-\delta)(v_L - c_L) + \delta V_A \geq (1-\delta)(v_L - c_L) + \delta \underline{V}_L \Leftrightarrow \underline{V}_L + \frac{\varepsilon_2}{1-\delta} \geq v_L - c_L.$$

Since $\varepsilon > \varepsilon_2$ can be chosen arbitrarily close to zero, is not possible that $v_L - c_L > \underline{V}_L$. Thus, the low quality seller's equilibrium payoff is no less than $v_L - c_L$. Since the equilibrium payoff of neither the high quality seller nor the buyers can be negative, the total surplus in the market is no less than $v_L - c_L$.

The claim of Proposition 3 follows because the buyers and the high quality seller's payoffs must be non-negative. ■

A.3 Proof of Theorem 1

First consider $\mu_0 > \mu^*$. By Proposition 1, in an opaque market with intra-period buyer competition there is a unique equilibrium which is efficient. The fully separating equilibrium constructed in Proposition 2 does not achieve full efficiency and exists for this range of initial beliefs. This establishes that when $\mu_0 > \mu^*$, transparency is welfare reducing when $\mu_0 > \mu^*$.

Next, consider $\mu_0 < \mu^*$. By Proposition 1, in an opaque market with intra-period buyer competition, there is a unique equilibrium in which high quality never trades and low quality trades efficiently. Thus, overall gains from trade is $(1-\mu_0)(v_L - c_L)$, which is the lower bound on gains from trade in a transparent market with buyer competition by Proposition 3. Further, the fully separating equilibrium constructed in Proposition 2 exists for this belief range and generates gains from trade

$$(1-\mu_0)(v_L - c_L) + \mu_0 \delta^{k(\delta)}(v_H - c_H) > (1-\delta)(v_L - c_L),$$

establishing that transparency is welfare improving when $\mu_0 < \mu^*$.

A.4 Partial pooling equilibria with intra-period buyer competition

Proposition 2 constructs a fully separating equilibrium. In this section we construct a class of partial pooling equilibria. Similar to the fully separating equilibrium, conditional on high quality, there is a positive amount of trade which is distorted down from its efficient level. Further, high quality's trade takes place always at the same price. Let Q_H be the expected discounted frequency with which the high quality trades, and P_H be the price at which she trades. Unlike in the fully separating equilibrium, the low quality now pools with the high quality along the said path with positive probability. With the remaining probability, the low quality trades efficiently (with probability 1 each period) at price v_L.

Unlike in the case of the fully separating equilibrium, a high-quality trading path with an initial pause followed by efficient trading may not be feasible. Instead, we construct trading paths that cycle through several periods of trade with single-period pauses.[23] For this purpose, for each k define the frequency Q_k by

$$Q_k = \frac{\delta + \delta^2 + \cdots + \delta^k}{1 + \delta + \cdots + \delta^k},$$

and the price P_k by

$$v_L - c_L = Q_k(P_k - c_L).$$

We show that as long as $Q_k > (1 - \delta)$ and $P_k \in [c_H, v_H]$, there exists an equilibrium where $Q_H = Q_k$ and $P_H = P_k$.

To construct such an equilibrium, define $\tau(h)$ to be the number of periods since the last pause of trade. Let $\tau(h) = \infty$ if every previous period involved trade or it is the null history, and naturally $\tau(h) = 0$ if the last period outcome was trade. We describe beliefs and strategies as functions of τ. We partition non-null histories into two groups:

- Case 1: There has been no previous streaks of trade exceeding k consecutive periods.
- Case 2: There has been at least one previous streak of trade exceeding k consecutive period.

Buyer strategies: In case 2, offer v_L. In case 1, if $\tau(h) = k$, offer v_L, otherwise offer P_k.

Seller strategies: The seller uses a type- and history-dependent reservation price. With an abuse of notation we write these reservation prices as functions of τ. They satisfy:

- Case 1: For $\tau < k$, $\theta = L, H$,

$$(1-\delta)\left[(P_\theta(\tau) - c_\theta) + \delta(P_k - c_\theta) + \cdots + \delta^{k-\tau-1}(P_k - c_\theta)\right] + \delta^{k-\tau} Q_k(P_k - c_\theta) = Q_k(P_k - c_\theta).$$

[23]This construction is similar to the one-step separation equilibria constructed in Kaya and Roy (2022). That paper considers limited records of past trading, and thus it cannot appeal to belief punishments for unexpected trading. In the current paper, such punishments are possible, and this makes it possible to construct different trading cycles than those discussed here.

For this case, we note that $P_\theta < P_k$. To see this substitute $P_\theta(\tau) = P_k$ to yield

$$(1-\delta)\left[(P_k - c_\theta) + \delta(P_k - c_\theta) + \cdots + \delta^{k-\tau-1}(P_k - c_\theta)\right] + \delta^{k-\tau}Q_k(P_k - c_\theta) = [(1-\delta^{k-\tau}) + \delta^{k-\tau}Q_k](P_k - c_\theta),$$

on the left-hand-side, which is larger than the right-hand-side since $Q_k < 1$.

For $\tau = k$:

$$(1-\delta)(P_L(\tau) - c_L) + \delta(v_L - c_L) = Q_k(P_k - c_L)$$
$$(1-\delta)(P_H(\tau) - c_H) = Q_k(P_k - c_H).$$

We note that in this case by choice of Q_k, P_k, $P_L(\tau) = v_L$. Further, since $(1-\delta) < Q_k$, $P_H(\tau) > P_k$.

- Case 2: $P_\theta(\tau) = c_\theta$.
- At $t = 1$: the reservation prices are identical to the case where $\tau = k$.

At all histories, the high quality seller accepts all offers that weakly exceed his reservation price, and rejects others. At $t = 1$ the low quality seller accepts his reservation price v_L with probability β satisfying

$$\frac{\mu_0}{1-\mu_0} = \frac{\mu_k}{1-\mu_k}(1-\beta),$$

where μ_k is defined by

$$\mu_k(v_H - P_k) + (1-\mu_k)(v_L - P_k) = 0.$$

At all other histories in Case 1, the low quality seller rejects all offers weakly less than his reservation price and accepts those that are strictly higher. In Case 2, she accepts all offers that weakly exceeds her reservation price and rejects all others.

Beliefs: In Case 2, $\mu(h) = 0$, in Case 1, $\mu(h) = \mu_k$.

Optimality of buyer strategies:

- In case 2, all buyers offering v_L is a bidding equilibrium because the belief is 0.

- In case 1, when $\tau < k$, we have $P_L(\tau) < P_H(\tau) < P_k$ and the expected quality is P_k. Therefore, it is a bidding equilibrium for all buyers to offer P_k. When $\tau = k$, we have $P_L(\tau) = v_L < P_k$ and $P_H(\tau) > P_k$. Thus offering v_L is a bidding equilibrium.

Optimality of seller strategies: The reservation prices are calculated using buyer offer strategies. Thus the decisions based on these reservation prices are optimal.

Belief consistency: Follows trivially from Bayes rule, when possible.

A.4.1 Maximally pooling equilibria when $\mu_0 \leq \mu^*$.

In the partial pooling equilibria constructed above, the buyers are always making pure strategy offers, and the belief remain strictly above μ^* except in a potential knife-edge case where there exists k with Q_k equal to

$$\frac{v_L - c_L}{c_H - c_L} \equiv Q^*$$

Here, we construct an equilibrium in which the high quality seller trades only at price c_H and at an expected discounted frequency $Q^* \equiv \frac{v_L - c_L}{c_H - c_L}$. In addition to being of interest for comparisons, it can also serve as an alternative punishment equilibrium to support partial and full pooling equilibria discussed so far.

In this equilibrium, the low quality seller follows this path with probability β satisfying

$$\frac{\mu_0}{1-\mu_0} = \frac{\mu^*}{1-\mu^*}(1-\beta),$$

and trades efficiently otherwise. The construction is almost identical to the pure-offer partial pooling equilibria above with the following modifications.

Fix k and α such that

$$\frac{\delta + \cdots + \delta^k}{1 + \delta + \cdots + \delta^k} \geq \frac{v_L - c_L}{c_H - c_L} \geq \frac{\delta + \cdots + \delta^{k-1}}{1 + \delta + \cdots + \delta^{k-1}},$$

and

$$\frac{v_L - c_L}{c_H - c_L} = \frac{\delta + \cdots + \delta^{k-1} + \alpha\delta^k}{1 + \delta + \cdots + \delta^{k-1} + \alpha\delta^k}.$$

As above define $\tau(h)$ to be the number of periods since the last pause of trade. Let

$\tau(h) = \infty$ if every previous period involved trade or it is the null history, and naturally $\tau(h) = 0$ if the last period outcome was trade.

Buyer strategies: Offer c_H if $\tau(h) < k$, offer v_L if $\tau(h) > k$, offer c_H with overall probability α if $\tau(h) = k$, and v_L otherwise.[24]

Seller strategies: As above, each type of the seller uses a reservation price strategy. Once again, we express reservation prices as functions of τ.

- $P_H(\tau) = c_H$ for any τ.
- $P_L(h)$ satisfies
  - If $\tau \geq k$
  $$(1-\delta)(P_L(\tau) - c_L) + \delta Q^*(c_H - c_L) = Q^*(c_H - c_L),$$
  therefore $P_L(h) = v_L$.
  - If $\tau < k$:
  $$(1-\delta)(P_L(\tau) - c_L) + \alpha\left\{\delta\left[1 + \delta + \cdots + \delta^{k-\tau}\right](1-\delta)(c_H - c_L) + \delta^{k-\tau+1}Q^*(c_H - c_L)\right\}$$
  $$+(1-\alpha)\left\{\delta\left[1 + \delta + \cdots + \delta^{k-\tau-1}\right](1-\delta)(c_H - c_L) + \delta^{k-\tau}Q^*(c_H - c_L)\right\}.$$
  In this case we note that $P_L(\tau) < c_H$. This is because, substituting c_H instead of $P_L(\tau)$ would yield the following left-hand-side:
  $$\left[(1 - \alpha\delta^{k-\tau+1} - (1-\alpha)\delta^{k-\tau}) + (\alpha\delta^{k-\tau+1} + (1-\alpha)\delta^{k-\tau})Q^*\right](c_H - c_L),$$
  which is larger than the right-hand-side since $Q^* < 1$.

The high quality seller accepts all offers weakly exceeding c_H. At $t = 1$, the low quality seller accepts his reservation price with probability β defined above. At $t \geq 2$, the low quality seller accepts his reservation price with probability 1 if $\tau = \infty$. Otherwise, he rejects his reservation price with probability 1.

Beliefs: If $\tau = \infty$, $\mu(h) = 0$. Otherwise, $\mu(h) = \mu^*$.

[24]Note that these strategies do not punish unexpected trade with a forever switch to low prices. Instead, after each pause of trade, the buyers offer c_H again for the next consecutive k or $k+1$ periods.

Optimality of buyer strategies: When $\tau = \infty$, the belief is 0, thus it is a bidding equilibrium for all buyers to offer v_L. When $k \leq \tau < \infty$, since the belief is μ^*, $P_L(h) = v_L$ and $P_H(h) = c_H$, all buyers offering c_H, all buyers offering v_L as well as buyers randomizing across c_H and v_L are bidding equilibria.

Optimality of seller strategies: Reservation prices are calculated using buyer offer strategies, and are therefore optimal.

Belief consistency: Follows trivially using Bayes rule from equilibrium strategies.

A.4.2 Accuracy of screening and gains from trade

Each of the partial pooling and the fully separating equilibria discussed so far are characterized by the price P_H at which the high quality trades and the expected discounted frequency Q_H with which she trades. In all these equilibria, the screening of the seller is completed in the first period, and thereafter, the belief is not updated on the equilibrium path. These equilibria can be ranked with respect to how accurate their screening is. In fact, take $P_H > P'_H$ and associated $Q_H < Q'_H$, a partial pooling equilibrium featuring (P_H, Q_H) is more informative in the sense of Blackwell than an equilibrium featuring (P'_H, Q'_H). The finer learning allows the high quality seller to trade at higher prices, but at lower frequency to ensure the credibility of learning. We note that in spite of this trade-off, the equilibria with more accurate learning feature higher gains from trade. To see this first note that in all these equilibria buyers' payoff is 0 and the low quality seller's payoff is $v_L - c_L$. Thus, the higher the high quality seller's payoff, the higher is the gains from trade (since the total gains from trade is equal to the sum of the payoffs of all players). The high quality seller's payoff can be expressed as

$$Q_H(P_h - c_H) = (v_L - c_L)\frac{P_H - c_H}{P_H - c_L},$$

because $Q_H = (v_L - c_L)/(P_H - c_L)$. It is easy to see that this expression increases in P_H.

B Omitted proofs of Section 5

Fix a history h. In what follows, (h, A) represents the continuation history of h obtained by adding one period of trade (acceptance) and (h, R) represents the continuation history of h obtained by adding one period of no trade (rejection).

B.1 Proof of Proposition 4

Here we replicate Lemmas 1, 2, 3 and present their proofs, as well as other preliminary results. Then we show how they come together to prove Proposition 4.

Lemma 1 *In any equilibrium of the transparent market with intra-period monopsony, at any history h, $V_H(h) = 0$. Therefore, the high quality seller accepts any offer that exceeds c_H.*

Proof. 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 $V_H(h) = 0$. ■

In addition to Lemma 1, the following result which was not discussed in the text, also plays a role in the proofs of Lemmas 2 and 3.

Lemma 5 *At any on or off-path history h, $P_H(h) = c_H \geq P_L(h)$.*

Proof. That $P_H(h) = c_H$ immediately follows by Lemma 1. This, in turn, implies that, buyers never offer any price exceeding c_H. If $P_L(h) > c_H$, low quality trades with probability 0 while high quality trades with probability 1, implying that $\mu(h, A) = 1$. If $\mu(h, A) = 1$, each subsequent buyer offers c_H which is accepted with probability 1, and the belief is never updated. Thus, at such history, $V_L(h, A) = c_H - c_L \leq V_L(h, R)$, where the latter inequality is because no price offers exceeding c_H is made. This, in turn implies $P_L(h) \leq c_L$, a contradiction. ■

Lemma 2 *In any equilibrium of the transparent market with intra-period monopsony, at any history h, if $\mu(h) > \mu^*$, then $V_L(h) \geq \delta(c_H - c_L)$.*

Proof of Lemma 2. Assume that $\mu(h) > \mu^*$. Buyer's payoff is strictly positive, thus he makes no losing offers. If trade takes place at $P < c_H$ with positive probability, then necessarily $P < v_L$, as othewise the buyer's payoff from offering P would be non-positive. This in turn implies that low quality trades with probability 1, because necessarily $P = P_L(h)$, and if it were being rejected with positive probability, the buyer would have a profitable deviation to a slightly higher offer. 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 α_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

$$\mu(h^*)(1 - \alpha_H)(v_H - c_H) + (1 - \mu(h^*))(1 - \alpha_L)(v_L - c_H),$$

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. ■

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

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 $P_L(h) \leq c_L$. But then, at h the buyer's payoff is strictly positive, and thus trade must take place with positive probability, and thus $h = h_1$, a contradiction.

Next, note that $h_1 = (h, R, \cdots, R)$. We claim that along this path, the belief remains strictly below μ^*. 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 $\mu(h_1) \leq \mu(h) < \mu^*$, establishing the claim.

Since $P_L(h_1) < c_H$, 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. ■

The next lemma is the last step in the proof of Proposition 4

Lemma 6 *Fix an equilibrium and a history h.*

1. *If $\mu(h) < \mu^*$ and (h, A) has positive probability conditional on having reached h, 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 $\mu(h, x) = \mu^*$, whenever (h, x) has positive probability conditional on having reached h, $x \in \{A, R\}$.*

3. *If $\mu(h) > \mu^*$ and (h, R) has positive probability conditional on having reached h, then $1 = \mu(h, R) > \mu(h, A) \geq \mu^*$. Further, in this case, necessarily $P_L(h) < v_L$.*

Proof of Lemma 6 . For item 1, first assume that $P_L(h) < c_H$. Then high type does not trade at h because buyer would make a loss offering c_H. Then, $\mu(h, A) = 0$. If at the same time $\mu(h, R) > \mu^*$, by Lemmas 2 and 3, we have $P_L(h) > v_L$. Since only the low quality trades, the buyer makes a loss. Thus, $\mu(h, R) \leq \mu^*$.

Next consider $P_L(h) = c_H$. Then trade takes place necessarily 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 $P_L(h) < v_L$, a contradiction. Thus, $\mu(h, A) = \mu^*$.

For item 2, first note that if (h, R) (respectively, (h, A)) is not on the equilibrium path, then necessarily $\mu^* = \mu(h) = \mu(h, A)$ (respectively, $\mu^* = \mu(h) = \mu(h, R)$). Assume both (h, A) and (h, R) are on the equilibrium path. Then, 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.

For item 3, we first note that in since $P_H(h) = c_H$ and $\mu(h) > \mu^*$, the buyer's payoff is necessarily positive. Thus, the buyer never makes a losing offer. Next, we show that $P_L(h) < c_H$. Suppose, for a contradiction, that $P_L(h) = c_H$. Then, $\mu(h, R) \leq \mu^*$ because otherwise the buyer would have a profitable deviation to $c_H + \varepsilon$ where $\varepsilon > 0$ is sufficiently small. Further, $\mu(h, R) = \mu^*$ because if $\mu(h, R) < \mu^*$, then by Lemmas 2 and 3, $P_L(h) < c_L$, which would lead to a contradiction. Further, (i) by Lemma 2, $V_L(h) \geq \delta(c_H - c_L)$, and (ii) by item 2, $V_L(h, R) = \tilde{Q}(c_H - c_L)$, for some $\tilde{Q}$. The latter assertion follows because once belief reaches μ^* it is never updated and thus high and low quality always trades with the same probability, and thus trade takes place only at price c_H at some frequency $\tilde{Q}$. Then, (i) and (ii) together with the supposition that $P_L(h) = c_H$ imply that $\delta\tilde{Q}(c_H - c_L) \geq \delta(c_H - c_L)$, and thus $\tilde{Q} = 1$. The latter means that the buyers offer c_H with probability 1 at every continuation history of (h, R). But this is a contradiction, because in that case, for such a continuation history h', $P_L(h') \leq c_L$ while $\mu(h') = \mu^*$ and the buyers would have a profitable deviation to offering $P_L(h) + \varepsilon$ for small $\varepsilon > 0$. This establishes that $P_L(h) < c_H$ whenever (h, R) has positive probability while $\mu(h) > \mu^*$.

Since $P_L(h) < c_H$, necessarily $\mu(h, R) = 1$. Further, if c_H is offered, it is accepted with probability 1 because otherwise the buyer would have a profitable deviation to $c_H + \varepsilon$ with $\varepsilon > 0$ sufficiently small. This also implies that $P_L(h) < v_L$ because otherwise the buyer would optimally offer c_H with probability 1, and $\mu(h, R)$ would have zero probability. Also note that if $\mu(h, A) < \mu^*$, then $\mu(h, R) > \mu^*$ and by Lemmas 2 and 3, we have $P_L(h) > v_L$, a contradiction. Thus, $\mu(h, A) \geq \mu^*$. ■

Proposition 4 directly follows from Lemma 6.

B.2 Proof of Proposition 5

Fix an equilibrium. Let $h_\emptyset^t$ represent the t-length history that features no trading. It follows by Lemma 6 that that for any t, if $(h_\emptyset^{t-1}, A)$ is on the equilibrium path, then $\mu(h_\emptyset^{t-1}, A) \in \{0, \mu^*\}$. Also, let $\gamma_s(h^t)$ be the probability with which the seller type $s \in \{L, H\}$ visits history h^t. Define $T_{\mu^*} = \{t | \mu(h_\emptyset^{t-1}, A) = \mu^*\}$ and $T_0 = \{t | \mu(h_\emptyset^{t-1}, A) = 0\}$. Then,

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

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 $t \in T_{\mu^*}$, $\overline{Q}_L(h_\emptyset^{t-1}, A) = \overline{Q}_H(h_\emptyset^{t-1}, A)$. It follows that

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

Next, since the low quality seller can always mimic the high quality, by Lemma 3, $v_L - c_L \geq V_L(h_0) \geq \overline{Q}_H(h_\emptyset)(c_H - c_L)$, or equivalently $\overline{Q}_H(h_\emptyset) \leq Q^*$. Plugging this in the above bounds and also using the fact that $\mu^*/(1-\mu^*) = (c_H - v_L)/(v_H - c_H)$, the result follows by simple algebra.

B.3 Proof of Proposition 6

We first show that when $\mu_0 > \mu^{**}$, there is a unique equilibrium where both types of the seller trade with probability 1 each period, which proves the second claim in Proposition 6. First note that whenever $\mu(h) > \mu^{**}$ for some h, $P_L(h) \geq c_L$. This is because by item (iii) of Lemma 6, at any such h, $\mu(h, R) = 1$, and therefore $V_L(h, R) = c_H - c_L \geq V_L(h, A)$. The latter inequality is because the buyers never offer a price exceeding c_H, and thus the maximum level of the continuation payoff is $c_H - c_L$. By definition,

$$
(1-\delta)(P_L(h) - c_L) = \delta(V_L(h, R) - V_L(h, A)),
$$

which is thus non-negative, and therefore $P_L(h) \geq c_L$. Next, note that a buyer arriving at such history h is better off offering c_H and trading with both types than targeting only the low quality, since $\mu(h) > \mu^{**}$. Thus, at such histories, trade takes place with probability 1 at price c_H, and the belief is never updated, establishing the last claim in the proposition.

Next assume that $\mu_0 \in (\mu^*, \mu^{**})$. Fix an equilibrium. Let h_1^t represent the t-length history that features trading at each period. By Lemma 6, for any such history, if (h_1^t, R) is on the equilibrium path, $\mu(h_1^t, R) = 1$ and $\mu(h_1^t, A) \geq \mu^*$. Further, if (h_1^t, R) is not on the equilibrium path, $\mu(h_1^t, A) \geq \mu^*$, as in this case belief is not updated. Define $T_{\mu^*} = \{t | \mu(h_\emptyset^{t-1}, A) = \mu^*\}$ and $T_1 = \{t | \mu(h_\emptyset^{t-1}, R) = 1\}$. Since no learning takes place once belief reaches μ^*, for each $t \in T_{\mu^*}$, $\overline{Q}_L(h_\emptyset^{t-1}, A) = \overline{Q}_H(h_\emptyset^{t-1}, A) < 1$. As usual, let $\gamma_s(h^t)$ be the probability with which the seller type $s \in \{L, H\}$ visits history h^t. Then,

$$\begin{aligned}
\overline{Q}_L(h_\emptyset) &= \sum_{t \in T_{\mu^*}} \left[(1-\delta^t) + \delta^t \overline{Q}_L(h^{t-1}, A)\right] \gamma_L(h_1^{t-1}, A) \\
&\geq (1-\delta) + \delta \min_{t \in T_{\mu^*}} \overline{Q}_L(h^{t-1}, A), \\
\overline{Q}_H(h_\emptyset) &= \sum_{t \in T_{\mu^*}} \left[(1-\delta^t) + \delta^t \overline{Q}_H(h^{t-1}, A)\right] \gamma_H(h_1^{t-1}, A) \\
&\quad + \sum_{t \in T_1} \left[(1-\delta^t) + \delta^{t+1}\right] \gamma_H(h_1^{t-1}, R) \\
&\geq (1-\delta) + \delta \min_{t \in T_{\mu^*}} \overline{Q}_H(h^{t-1}, A) \underbrace{\frac{\mu^*}{1-\mu^*} \frac{1-\mu_0}{\mu_0}}_{=\sum_{t \in T_{\mu^*}} \gamma_H(h_1^{t-1}, A)} \\
&\quad + \delta \left(1 - \frac{\mu^*}{1-\mu^*} \frac{1-\mu_0}{\mu_0}\right).
\end{aligned}$$

Further, letting $\tilde{Q}_t \equiv \overline{Q}_L(h^{t-1}, A) = \overline{Q}_T(h^{t-1}, A)$ represent the continuation expected amount of trade if the belief reaches μ^* for the first time at time t, then at each $t \in T_{\mu^*}$, the following must hold:

$$(1-\delta)(v_L - c_L) + \delta \tilde{Q}_t (c_H - c_L) \geq \delta (c_H - c_L).$$

This is because for such belief updating to be possible, both (h^{t-1}, A) and (h^{t-1}, R) must have positive probability, which is possible only when $P_L(h^{t-1}) \leq v_L$. This is equivalent to

$$\delta(1 - \tilde{Q}_t)(c_H - c_L) \leq (1-\delta)(v_L - c_L) \Leftrightarrow \tilde{Q}_t \geq 1 - \frac{1-\delta}{\delta} Q^*.$$

Then, the surplus conditonal on low quality is bounded below by

$$[(1-\delta) + \delta(1 - \frac{1-\delta}{\delta} Q^*)](v_L - c_L) = (1 - (1-\delta) Q^*)(v_L - c_L).$$

And, the surplus conditional on high quality is no less than

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

Then the expected surplus is no less than μ_0 times the latter plus $(1-\mu_0)$ times the former. These can be re-organized as

$$\underbrace{(1-\mu_0)(1-(1-\delta)Q^*)(v_L-c_L)+(1-\mu_0)(1-(1-\delta)Q^*)(c_H-v_L)}_{\equiv B}$$
$$+$$
$$\underbrace{\mu_0\delta\left(1-\frac{1-\mu_0}{\mu_0}\frac{\mu^*}{1-\mu^*}\right)(v_H-c_H)}_{>0}.$$

Simple algebra yields the claimed expression in the proposition.

Next we show that the lower bound on the gains from trade strictly exceeds $(1-\delta)(v_L-c_L)$. Re-organizing the expression for B we get

$$B=(1-\mu_0)[(c_H-c_L)-(1-\delta)(v_L-c_L)]$$

We show that $B\geq(1-\mu_0)(v_L-c_L)$. This is equivalent to

$$(c_H-c_L)-(1-\delta)(v_L-c_L)\geq v_L-c_L\Leftrightarrow\frac{1}{2-\delta}\geq Q^*.$$

Since $Q^*\leq\delta$ by Assumption (1), a sufficient condition is

$$\frac{1}{2-\delta}\geq\delta\Leftrightarrow\delta^2-2\delta+1\geq0\Leftrightarrow(1-\delta)^2\geq0,$$

which holds.

B.4 Equilibria of the transparent market with intra-period monopsony

We construct equilibria using dynamic programming techniques in the same spirit as Abreau et al. (1990) and Fudenberg et al. (1994). We borrow terminology and techniques from these papers and adjust them to account for the fact that our game features private information. Because our construction is specific to this setting and because we do not seek to characterize all equilibria but simply a subset, this exercise remains tractable.

For each $\mu \in [0,1]$ let $\mathcal{U}_\mu \in \mathbb{R}_+$ be a set of potential payoffs for the low quality seller.[25] We say that $U \in \mathbb{R}_+$ is **enforceable with respect to** $\{\mathcal{U}_\mu\}_{\mu \in [0,1]}$ **at belief** μ if there exists $\alpha \in [0,1], P \leq c_H$, $\mu^A, \mu^R \in [0,1], U^A \in \mathcal{U}_{\mu^A}, U^R \in \mathcal{U}_{\mu^R}$ that satisfy

$$U = \alpha[(1-\delta)(c_H - c_L) + \delta U^A] + (1-\alpha)U^R. \tag{10}$$

$$(1-\delta)(P - c_L) = \delta(U^R - U^A) \tag{11}$$

- If $\mu = 0$, then $\alpha = 0$ and $\mu^A = \mu = 0$, and $P \leq v_L$.
- If $\mu \in (0, \mu^*)$, then $\alpha = 0$, $P = v_L$, $\mu^A = 0$ and $\mu^R = \mu^*$
- If $\mu = \mu^*$ then $P \geq v_L$,
  - if $\alpha > 0$, $\mu^A = \mu^*$
  - if $\alpha < 1$, $\mu^R = \mu^*$
- If $\mu > \mu^*$, then $\mu^R = 1$ and $\frac{\mu^A}{1-\mu^A} = \frac{\mu}{1-\mu}\alpha$, and
  - if $\alpha > 0$, then

$$\mu(v_H - c_H) + (1-\mu)(v_L - c_H) \geq (1-\mu)(v_L - P)$$

  - if $\alpha < 1$, then

$$\mu(v_H - c_H) + (1-\mu)(v_L - c_H) \leq (1-\mu)(v_L - P)$$

Intuitively, here α is to the probability with which the buyer offers c_H in an equilibrium, with the understanding that the with the remaining probability the buyer offers the low quality seller's reservation price, represented by P. The values U^A and U^R respectively represent the continuation payoffs of the low quality seller after trade and no-trade and $\mu^A, \mu^R \in [0,1]$ represent the corresponding continuation beliefs. The condition (10) guarantees that indeed with such α and such continuation payoffs, the low quality seller's equilibrium payoff is U, while (11) requires that P is indeed the low quality's continuation

[25] In all equilibria, the high quality seller receives a payoff of zero. Since the buyers are short-lived, their incentives do not need to be considered in the equilibrium conditions. All information stemming from dynamics and is relevant to buyer payoffs is summarized by the seller's type-specific reservation price.

payoff. The conditions listed for each separate belief restrict the buyer strategy summarized by α, the low quality seller's reservation price P and the continuation beliefs μ^A and μ^R so that (i) it is optimal for the buyer to choose this strategy, and (ii) continuation beliefs are derived using Bayes rule whenever possible.

Next, we say that $\{\mathcal{U}_\mu\}_{\mu\in[0,1]}$ is **self-generating** if for all μ, all $U \in \mathcal{U}_\mu$ is enforceable with respect to $\mathcal{U} \equiv \bigcup_{\mu\in[0,1]} \mathcal{U}_\mu$ at belief μ.

Proposition 7 *If $\{\mathcal{U}_\mu\}_{\mu\in[0,1]}$ is self-generating, then for each μ and $U \in \mathcal{U}_\mu$, when the belief is μ, there exists an equilibrium that delivers the low quality seller a payoff of U.*

Proof. The proof follows by iterative construction of equilibria. ■

B.4.1 Constructing a self-generating set

Let $\mu^0 = \mu^*$, $P^0 = v_L$, $U^0 = (1-\delta)(v_L - c_L) + \delta(c_H - c_L)$. For $i = 1, 2, \cdots$ define $\mu_i, U^i, P^i, \alpha_i$ as follows:

$$U^i = c_H - c_L - (1-\delta)\delta^i(v_L - c_L)$$
$$P^i = c_L + \delta^i(v_L - c_L)$$
$$\frac{\mu^i}{1-\mu^i} = \frac{c_H - P^i}{v_H - c_H}$$
$$\alpha_i = \frac{c_H - P^{i-1}}{c_H - P^i}.$$

Note that as $i \to \infty$, we have $P^i \to c_L$, $U^i \to c_H - c_L$ and $\mu_i \to \mu^{**}$

For each $i = 1, 2, \cdots$, let

$$\mathcal{U}_{\mu_i} = [U^{i-1}, U^i],$$

and for each $\mu \in (\mu_{i-1}, \mu_i)$, let

$$\mathcal{U}_\mu = \{U^{i-1}\}.$$

Also let

$$\mathcal{U}_0 = [0, v_L - c_L]$$
$$\mathcal{U}_{\mu^*} = [(1-\delta)(v_L - c_L), c_H - c_L - (1-\delta)(v_L - c_L)]$$

Finally, for all $\mu \in (0, \mu^*)$, let $\mathcal{U}_\mu = [(1-\delta)(v_L - c_L), v_L - c_L]$ and for all $\mu \geq \mu^{**}$, let $\mathcal{U}_\mu = \{c_H - c_L\}$. Note that

$$\bigcup_{\mu \geq \mu^*} \mathcal{U}_\mu = [(1-\delta)(v_L - c_L), c_H - c_L].$$

Proposition 8 $\{\mathcal{U}_\mu\}_{\mu \in [0,1]}$ *where for each* μ, $\mathcal{U}_\mu$ *is as defined above, is self-generating.*

Proof. We take up each U_μ and show that they satisfy the appropriate conditions listed above.

Case 1: $\mu = 0$: Take $U \in [0, v_L - c_L]$. Let $\alpha = 0, \mu^A = 0$ and $U^A = U$, $U^R = U/\delta$ so that (10) is satisfied. Let P be such that $U = P - c_L$ so that (11) is satisfied. Note that $U \in \mathcal{U}_0$. Since $\alpha = 0$, μ^R can be chosen arbitrarily. If $U^R = U/\delta < (1-\delta)(v_L - c_L)$, choose $\mu^R = 0$, otherwise choose $\mu^R = \mu^*$. It just remains to argue that whenever $U/\delta \geq (1-\delta)(v_L - c_L)$, $U/\delta \in \mathcal{U}_{\mu^*}$. This is equivalent to

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

which holds by Assumption 1 ($Q^* < \delta^2$).

Case 2: $\mu \in (0, \mu^*)$: Take $U \in [(1-\delta)(v_L - c_L), v_L - c_L]$. Let $\alpha = 0$, $\mu^A = 0$, $\mu^R = \mu^*$, $P = v_L$. Also, by (10), $U^R = U/\delta$. Note that this is necessarily in $\mathcal{U}_{\mu^*}$ [shown above]. It suffices to show that $U^A \in \mathcal{U}_0$ can be chosen to satisfy (11):

$$(1-\delta)(v_L - c_L) = \delta(U/\delta - U^A)$$

Note that U^A must be larger the larger is U. If $U = (1-\delta)(v_L - c_L)$, then $U^A = 0 \in \mathcal{U}_0$. If $U = v_L - c_L$, then $U^A = v_L - c_L \in \mathcal{U}_0$. Thus, for any $U \in [(1-\delta)(v_L - c_L), v_L - c_L]$, $U^A \in [0, v_L - c_L] = \mathcal{U}_0$.

Case 3: $\mu = \mu^*$ We partition $\mathcal{U}_{\mu^*}$ into three components:

- First consider $\mathbf{U} \in [(\mathbf{1} - \delta)(\mathbf{v_L} - \mathbf{c_L}), (\mathbf{1} - \delta^{\mathbf{2}})(\mathbf{v_L} - \mathbf{c_L})]$. Let $\alpha = 0$, $\mu^A = 0$, $\mu^R = \mu^*$, and $P = v_L$. Further, let

$$U^A = \frac{U - (1-\delta)(v_L - c_L)}{\delta}, U^R = \frac{U}{\delta}, \alpha = 0.$$

By choice of U^A, U^R, (10) and (11) are satisfied. Moreover, $U^A \in [0, (1-\delta)(v_L - c_L)] \subset \mathcal{U}_0$. Finally, we need $U^R \geq (1-\delta)(v_L - c_L)$, and

$$U^R \leq \frac{1-\delta^2}{\delta}(v_L - c_L) \leq c_H - c_L - (1-\delta)(v_L - c_L).$$

The former trivially follows because $U \geq (1-\delta)(v_L - c_L)$. The latter also holds by Assumption 1. To see this, re-arrange the second inequality as

$$\left[\frac{1-\delta^2}{\delta} + (1-\delta)\right](v_L - c_L) \leq c_H - c_L.$$

which is equivalent to

$$\frac{v_L - c_L}{c_H - c_L} \leq \frac{\delta}{1-\delta^2 + \delta(1-\delta)}.$$

By Assumption 1, the left-hand-side is less than δ^2. Then, a sufficient condition is

$$\delta^2 \leq \frac{\delta}{1-\delta^2 + \delta(1-\delta)}.$$

which is equivalent to

$$-\delta^3 \leq (1-\delta)(1-\delta^2),$$

which always holds. Thus, $U^R \in \mathcal{U}_{\mu^*}$. This establishes that such U is enforceable.

- Next consider $\mathbf{U} \in [(\mathbf{1} - \boldsymbol{\delta}^\mathbf{2})(\mathbf{v_L} - \mathbf{c_L}), \mathbf{c_H} - \mathbf{c_L} - (\mathbf{1} - \boldsymbol{\delta}^\mathbf{2})(\mathbf{v_L} - \mathbf{c_L})]$. Let $\mu^A = \mu^R = \mu^*$ and $P = v_L$. Consider

$$U^A \in [(1-\delta)(v_L - c_L), c_H - c_L - \frac{1-\delta^2}{\delta}(v_L - c_L)], U^R = U^A + \frac{1-\delta}{\delta}(v_L - c_L)$$

Such U^A, U^R satisfy (11) and $U^A, U^R \in \mathcal{U}_{\mu^*}$. Further, by choice of α they enforce

$$U \in [\delta U^R, (1-\delta)(c_H - v_L) + \delta U^R] = [\delta U^A + (1-\delta)(v_L - c_L), \delta U^A + (1-\delta)(c_H - c_L)]$$

When varying U^A over the allowed range we obtain the minimum U enforced to be when $U^A = (1-\delta)(v_L - c_L)$ and $\alpha = 0$. This enforces

$$\underline{U} = \delta(1-\delta)(v_L - c_L) + (1-\delta)(v_L - c_L) = (1-\delta^2)(v_L - c_L).$$

We obtain the maximum enforced U when $U^A = c_H - c_L - \frac{1-\delta^2}{\delta}(v_L - c_L)$ and $\alpha = 1$. This enforces

$$\overline{U} = \delta(c_H - c_L - \frac{1-\delta^2}{\delta}(v_L - c_L)) + (1-\delta)(c_H - c_L) = c_H - c_L - (1-\delta^2)(v_L - c_L).$$

Since U^A and α can be varied continuously, and the enforced quantity continuously increases in both, all $U \in [\underline{U}, \overline{U}] = [(1-\delta^2)(v_L - c_L), c_H - c_L - (1-\delta^2)(v_L - c_L)]$ are enforceable.

- Finally consider $\mathbf{U} \in [\mathbf{c_H} - \mathbf{c_L} - (\mathbf{1} - \delta^{\mathbf{2}})(\mathbf{v_L} - \mathbf{c_L}), \mathbf{c_H} - \mathbf{c_L} - (\mathbf{1} - \delta)(\mathbf{v_L} - \mathbf{c_L})]$. Consider

$$U^R \in [(c_H - c_L - (1-\delta)(v_L - c_L), c_H - c_L] \equiv \mathcal{U}_I,$$

and

$$U^A = U^R - \frac{1-\delta}{\delta}(v_L - c_L), \text{ and } \alpha = 1$$

These choices enforce $U \in [\underline{U}, \overline{U}]$ where

$$\begin{aligned}
\underline{U} &= \delta[c_H - c_L - (1-\delta)(v_L - c_L) - \frac{1-\delta}{\delta}(v_L - c_L)] + (1-\delta)(c_H - c_L) \\
&= c_H - c_L - (1-\delta^2)(v_L - c_L) \\
\overline{U} &= \delta[c_H - c_L - \frac{1-\delta}{\delta}(v_L - c_L)] + (1-\delta)(c_H - c_L) \\
&= c_H - c_L - (1-\delta)(v_L - c_L)
\end{aligned}$$

Thus, $[\underline{U}, \overline{U}] = [c_H - c_L - (1-\delta^2)(v_L - c_L), c_H - c_L - (1-\delta)(v_L - c_L)]$.

Case 4: $\mu > \mu^*$ We will consider two subcases.

- Fix $i > 0$ and consider $\mu \in (\mu_{i-1}, \mu_i)$. For such μ, $\mathcal{U}_\mu \equiv \{U^{i-1}\}$.

We show that U^{i-1} is enforceable at μ with respect to $\mathcal{U}$. Choose $\mu^R = 1$, $\mu^A = \mu_{i-1}$, $U^R = c_H - c_L$. Also choose P, α, U^A as follows:

$$\frac{\mu}{1-\mu} = \frac{c_H - P}{v_H - c_H}$$
$$\frac{\mu}{1-\mu}\alpha = \frac{\mu_{i-1}}{1-\mu_{i-1}} \equiv \frac{c_H - P^{i-1}}{v_H - c_H}.$$
$$(1-\delta)(P - c_L) = \delta(c_H - c_L - U^A).$$

The first equality uniquely defines P, the second uniquely defines α and the third uniquely defines U^A. The equivalence is due to the definition of μ_{i-1} and P^{i-1}. By choice of U^A, (11) is satisfied. By choice of P,

$$\mu(v_H - c_H) + (1-\mu)(v_L - c_H) = (1-\mu)(v_L - P)$$

And the right-hand-side of (10) becomes

$$[(1-\delta)(P - c_L) + \delta U^A] + \alpha(c_H - P)(1-\delta)$$

By (11), the term in brackets is equal to $\delta(c_H - c_L)$. Plugging this in and substituting for α and μ from the first two equalities above

$$\delta(c_H - c_L) + (1-\delta)(c_H - P^{i-1}) = U^{i-1},$$

where the equality follows by the definition of P^{i-1} and U^{i-1}. It remains to show that $U^A \in \mathcal{U}_{\mu_{i-1}} = [U^{i-2}, U^{i-1}]$. To see this we note that, since $\mu \in (\mu_{i-1}, \mu_i)$, we have $P \in (P^i, P^{i-1})$. Since for any i, $(1-\delta)(P^i - c_L) = \delta(c_H - c_L - U^{i-1})$, we have

$$\underbrace{\delta(c_H - c_L - U^{i-1})}_{(1-\delta)(P^i - c_L)} < \underbrace{\delta(c_H - c_L - U^A)}_{(1-\delta)(P - c_L)} < \underbrace{\delta(c_H - c_L - U^{i-2})}_{(1-\delta)(P^{i-1} - c_L)}.$$

This shows that U is enforceable.

- Now consider $\mu = \mu_i$ for some i. Let $U \in \mathcal{U}_{\mu_i} = [U_{i-1}, U_i]$.

We show that U is enforceable at μ_i with respect to $\mathcal{U}$. Choose $\mu^R = 1$, $U^A = U_{i-1}$, $U^R = c_H - c_L$, $P = P^i$. Also choose α, μ^A as follows:

$$\delta(c_H - c_L) + (1-\delta)\alpha(c_H - P^i) = U,$$

so that (10) is satisfied. Since $\delta(c_H - c_L) + (1-\delta)(c_H - P^i) = U^i > U$ and $\delta(c_H - c_L) + (1-\delta)\alpha_i(c_H - P^i) = U^{i-1} < U$, we have $\alpha \in (\alpha_i, 1)$. Let μ^A be given by

$$\frac{\mu^A}{1-\mu^A} = \frac{\mu_i}{1-\mu_i}\alpha.$$

Thus, $\mu^A \in (\mu_{i-1}, \mu_i)$ and by construction, $U^{i-1} \in \mathcal{U}_{\mu^A}$. Further, by construction,

$$\mu(v_H - c_H) + (1-\mu_i)(v_L - c_H) = (1-\mu)(v_L - P_i).$$

Finally, (11) is satisfied by construction of P_i because

$$(1-\delta)(P^i - c_L) = \delta(c_H - c_L - U^{i-1}).$$

■

B.5 Proof of Theorem 2

First consider $\mu_0 < \mu^*$. By Proposition 5, the gains from trade in a transparent market is never larger than that from an opaque market. By Propositions 7 and 8 there exists equilibria of the transparent market that generate strictly less gains from trade than the opaque market. This establishes that transparency is welfare reducing in a market with intra-period monopsony when $\mu_0 < \mu^*$.

Next, consider $\mu_0 > \mu^{**}$. By the second part of Proposition 6, the gains from trade in a transparent market is necessarily the same as that in an opaque market.

Finally, consider $\mu_0 \in (\mu^*, \mu^{**})$. By the first part of Proposition 6, the gains from trade in a transparent market is strictly larger than that in an opaque market. This establishes that transparency is welfare-improving in a market with intra-period monopsony when $\mu_0 \in (\mu^*, \mu^{**})$.

References

Abreau, D., D. Pearce and E. Stachetti, "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* (1990), 1041–1063.

Admati, A. and M. Perry, "Strategic Delay in Bargaining," *Review of Economic Studies* (1987), 345–364.

Deneckere, R. and M.-Y. Liang, "Bargaining with Interdependent Values," *Econometrica* 74 (2006), 1309–1364.

Dilme, F., "Repeated bargaining with imperfect information about previous transactions," *working paper* (2022).

Evans, R., "Sequential Bargaining with Correlated Values," *The Review of Economic Studies* 56 (1989), 499–510.

Fuchs, W., P. Gottardi and H. Moreira, "Time trumps quantity in the market for lemons," *working paper* (2022).

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

Fuchs, W. and A. Skrzypacz, "Costs and benefits of dynamic trading in a lemons market," *Review of Economic Dynamics* (2019), 105–127.

Fudenberg, D., D. Levine and E. Maskin, "The Folk Theorem with Imperfect Public Information," *Econometrica* 62 (1994), 997–1039.

Gerardi, D., L. Maestri and I. Monzón, "Bargaining over a Divisible Good in the Market for Lemons," *American Economic Review* 112 (May 2022), 1591–1620.

Hart, O. and J. Tirole, "Contract Renegotiation and Coasian Dynamics," *Review of Economic Studies* LV (1988).

Hörner, J. and N. Vieille, "Public vs. Private Offers in the Market for Lemons," *Econometrica* 77 (2009), 29–69.

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 S. Roy, "Price Transparency and Market Screening," *working paper* (2020).

———, "Market Screening with Limited Records," *Games and Economic Behavior* (2022), 106–132.

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

Noldeke, G. and E. V. Damme, “Signaling in a Dynamic Labour Market,” *The Review of Economic Studies* (1990), 1–23.

Ortner, J., “Bargaining with Evolving Private Information,” *Theoretical Economics* (forthcoming).

Pei, H., “Reputation Building Under Observational Learning,” *Review of Economic Studies* (forthcoming).

Spence, M., “Job Market Signaling,” *The Quarterly Journal of Economics* 87 (1973), 355–374.

Swinkels, J. M., “Education signaling with preemptive offers,” *The Review of Economic Studies* (1999), 949–970.

Vincent, D., “Bargaining with Common Values,” *Journal of Economic Theory* 48 (May 1989), 47–62.

———, “Dynamic Auctions,” *Review of Economic Studies* 57 (1990).