Repeated Trading: Transparency and Market Structure

Ayça Kaya* and Santanu Roy†

Abstract

We analyze the effect of transparency of past trading volumes in markets where an informed long-lived seller can repeatedly trade with short-lived uninformed buyers. Transparency allows buyers to observe previously sold quantities. In markets with intra-period monopsony (single buyer each period), transparency reduces welfare if the ex-ante expected quality is low, but improves welfare if the expected quality is high. The effect is reversed in markets with intra-period competition (multiple buyers each period). This discrepancy in the efficiency implications of transparency is explained by how buyer competition affects the seller's ability to capture rents, which, in turn, influences market screening.

Key words: repeated sales, adverse selection, competition, transparency, market efficiency
JEL codes: D82, D61, C73

1 Introduction

In many markets, sellers don't have the opportunity to form long-term relationships with their customers. Rather, they complete 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

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

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, resulting in an instance of the classical lemons problem.

Given that the lemons problem is an informational problem, and given that such sellers trade repeatedly over time, a natural question to ask is how transparency of the sellers' past trading behavior would impact the workings of such markets. Naturally, the past behavior of a seller would provide indirect clues about her private information, potentially alleviating the lemons problem. On the flip side, the understanding that information about past 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.[1] 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 study this question by focusing on a weak yet credible form of transparency. Namely, we consider how availability of information about sellers' past trading *volumes* (and not the trading prices or other aspects of history) may impact market efficiency. With recent technological developments even the smallest transactions are likely to be electronic, creating a verifiable source of information about past transactions of each seller, making this type of transparency feasible. Thus, our results may shed light on certain aspects of the optimal design and regulation of markets.

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 seller's ability to capture rents, and (ii) the initial market perception of quality. Figure 1

[1]It is well known that signaling (as well as screening) of private information can be associated with significant distortions (see, for instance, Spence (1973)). In the dynamic lemons market, inefficiency in trading arises when sellers withhold current trading in order to improve buyers' beliefs about their product quality; see, among others, Fuchs and Skrzypacz (2015, 2019). In our setting as well, for instance, if infrequent trade is associated with higher cost/higher quality, the sellers may be inclined to inefficiently slow down their sales to be able to improve the perception of the quality they are offering.

	low initial belief $\mu_0 < \mu^*$	intermediate initial belief $\mu^* < \mu_0 < \mu^{**}$
monopsony	$-$	$+$
competition	$+$	$-$

Figure 1: Impact of transparency on gains from trade.

summarizes our conclusions, with a $+$ indicating cases where transparency promotes more efficient trading and a $-$ indicating the opposite.

We obtain these results within the context of a formal model featuring a long-lived seller who has the capacity to sell one unit of a good each period. The binary quality ($\theta \in \{L, H\}$) of her output is exogenously fixed, is persistent and is her private information. The market's belief μ is the probability that is assigned to high quality ($\theta = H$), with μ_0 representing the initial belief. 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 are always positive. Further, 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.

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 say that transparency is welfare-reducing if all perfect Bayesian equilibria of a transparent market generate smaller gains from trade than any equilibrium of an opaque market, with at least one generating strictly less. The case where transparency is welfare-improving is analogously defined.

Naturally, in opaque markets the buyers cannot learn the seller's type, and the belief is never updated. Thus, the outcomes in such markets mirror those in a static setting. As such, the high quality can trade if and only if the market's initial belief is above a certain cutoff. In fact, when this is the case the high quality, as well as the low quality, trades efficiently, while below the cutoff only the low quality trades. The level of this "efficiency cutoff" belief depends on buyer competition. We consider two specifications with respect to buyer competition: 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. In Figure 1, what delineates lower from intermediate initial beliefs is the efficiency cutoff for an opaque market with competitive buyers (which we denote by μ^*), while the upper bound on the intermediate beliefs is the corresponding cutoff for markets with intra-period monopsony (which we denote by μ^{**}).

In opaque markets with intra-period buyer competition, the efficiency cutoff corresponds to the standard "lemons cutoff" and is exactly when the expected use value of the seller's offering is equal to the high quality's cost. The corresponding cutoff with intra-period monopsony is larger than the lemons cutoff, because in such markets the buyers can extract all gains from trade with the low quality seller, and are unwilling to target the high quality seller unless doing so delivers an expected payoff above a certain strictly positive bound. Introducing transparency in such a market remedies this additional inefficiency due to monopsony by allowing the low quality seller to capture information rents, and making it less attractive for buyers to go after only the low quality seller. This explains why transparency is welfare improving when initial beliefs are intermediate and the market features intra-period monopsony (top right cell in Figure 1).

The rest of the results summarized in Figure 1 are all explained by how the market's accuracy of screening the seller and the cost of such screening varies by buyer competition. Specifically, we find that markets are able to learn more accurately and at lower cost if the high quality seller captures positive rents. The high quality seller can capture rents in a market with buyer competition and not in a market with intra-period monopsony, driving the differential impact of transparency on these two types of markets.

To get a better understanding of how the seller's ability to capture rents informs market's ability to screen, and how this explains our results summarized in Figure 1, first consider a transparent market with *intra-period buyer competition*. In these markets the high quality seller can capture positive rents, and thus she may also fear losing them. In order to preserve *future* rents, she may be willing to forego some *current* rents by rejecting prices that exceed her cost. She would be incentivized to do so, for instance, if the market interprets (unexpectedly) frequent trading as bad news about quality, leading to lower price offers in the future. A crucial implication is that even when the market is quite convinced that the seller is of high quality (i.e. when the belief is above the lemons cutoff), trade can be slow on the equilibrium path. This observation explains why transparency can be harmful when the initial belief is high, since in those cases the opaque market necessarily delivers efficiency (bottom right cell of Figure 1).

On the flip side, the existence of slow-trade equilibrium when beliefs are high is precisely what allows credible screening of the seller *when the initial belief is low.* This is best understood by considering a complete learning equilibrium that exists in transparent markets with buyer competition (Propositon 2). In this equilibrium, all screening takes place in the first period in which the low quality seller reveals herself by trading with probability 1, and thus upon failure of trade in the first period, the market assigns all probability to the high quality. Since at each history, buyer competition drives the prices up to the expected quality, credible screening requires that high quality trades slowly, or else the low quality seller could not be incentivized to reveal herself. This highlights the crucial role that the high quality seller's ability to capture rents—and the resulting possibility of slow trade—plays in allowing the market to accurately screen the seller. The final piece is to note that this complete learning equilibrium improves upon the opaque market with buyer competition when initial belief is low, because this equilibrium features positive amount of trade by the high quality while featuring efficient trading by the low quality (bottom left cell of Figure 1).

The discussion so far leads to 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 $\mu_0 > \mu^*$.*

Next, consider a transparent market with *intra-period monopsony*. In such markets, in contrast to those with buyer competition, the high quality seller cannot capture any rents (Lemma 2). In this case, the high quality seller expects to never receive positive payoffs in the future. Thus, she will never walk away from a positive payoff in the current period, i.e., she will accept any offer that exceeds her cost c_H. Then, any buyer who arrives with belief exceeding the lemons cutoff (i.e. having an expectation of the use value in excess of c_H), can guarantee himself a positive payoff by offering a price (slightly above) c_H. This renders it impossible to slow down trade once such beliefs are reached, which in turn makes it exceedingly attractive for the low quality seller to reach histories where beliefs are this high (Lemma 2). This suggests that if the market starts at a belief lower than μ^*, it can never cross over this cutoff in any continuation history. Thus, a significant amount of pooling cannot be avoided (Proposition 4), or equivalently, the market cannot accurately screen the seller.

The unavoidability of extensive pooling limits the gains from trade, because it requires that the low quality seller sometimes trades at prices that exceed the value that his production creates, and thus must be cross-subsidized by the high quality's costlier production. To get a clearer picture of the impact of cross-subsidization, consider once again an equilibrium where all screening takes place in the first period. In this case, the low quality seller reveals herself with positive probability by trading in the first period, and with the remaining probability mimics the high quality's inefficient trading path. Then, upon trade in the first period, the market's belief is 0 while upon no trade in the first period, the market's belief updates to exactly the lemons cutoff so that thereafter trade takes place at price c_H . In such an equilibrium, the low quality seller's discounted average payoff is (at most) $v_L - c_L$ while the high quality seller as well as all short-run buyers receive zero payoffs. The resulting gains from trade are therefore no higher than in an opaque version of such a market in which only the low quality trades creating surplus equal to $v_L - c_L$ every period.

The final piece of the argument explains why a transparent market with intra-period monopsony may generate strictly less gains from trade than its opaque counterpart. This is due to the fact that in such markets, (and in contrast to the case with buyer competition), once revealed, the low quality seller may not receive any rents at all. Consequently, the alternative of mimicking the high quality's trading is very attractive and to deter such mimicking, the high quality's trading must be slowed down a lot further. It follows that, when the initial belief is sufficiently low, in a market with intra-period monopsony, transparency is welfare-reducing. This explains the the top left cell of Figure 1.

To recap, when the initial belief is low ($\mu_0 < \mu^*$), in opaque markets, all gains from trade comes from low quality's *efficient* trading. Transparency allows the high quality to trade some. But the amount of trade is limited by the low quality's incentives to mimic, which are very strong due to the fact that she may not be able to capture any rents when she reveals herself. Further, the low quality's trading is substantially reduced from its efficient level due to the need for extensive pooling—which, in turn, is due the high quality seller's inability to capture rents. Thus transparency brings in little gain in terms of increased trading of high quality at a large cost in terms of substantially reduced trading of low quality. In balance, it reduces the overall gains from trade.

Our second main result consists of these observations, along with the observation that when the opaque market is efficient (i.e. the initial belief is above μ^{**}) the transparent market is also efficient.

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 $\mu_0 > \mu^{**}$.*

So far, we explained our results in reference to specific equilibria or paths of play. Our proofs are, however, based on the properties that all equilibria must satisfy, and not on construction of specific sets of equilibria. One difficulty with this approach is that in this environment, it is not possible to establish the so-called "skimming property" from first principles.[2] In particular, it is not necessary that in every equilibrium, and at every history the high quality must trade with a smaller probability. In our proofs, we are able to overcome this challenge by referring to the properties of full equilibrium histories, and equilibrium probability distributions over these histories implied by weak non-mimicking conditions which must be satisfied by all equilibria.

[2]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 fail because the low quality seller may have a strong incentive to wait for frequent middling offers which would not be profitable for the high quality/cost seller, and thus the low quality seller 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.

We also construct sets of equilibria for each possible initial belief. This allows us to complete the arguments for Theorems 1 and 2 as well as to establish existence of equilibrium. Construction of these equilibria is non-trivial. Specifically, all equilibria we construct feature trading cycles, where the probability of trade conditional on quality varies across histories. These cycles are needed to guarantee that the seller's type-dependent reservation price remains within certain bounds, which in turn guarantees the existence of optimal price offer strategies for buyers. We are able to directly construct these trading cycles for markets featuring intra-period buyer competition. For markets featuring intra-period monopsony, this construction becomes insurmountably tedious. We overcome this difficulty by resorting to techniques introduced by Abreu 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.[3]

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 sell *a unit in each period*, sequentially meeting *short-lived* potential buyers. Our specific focus is 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, receives offers from uninformed parties and leaves the market

[3]The techniques developed by Abreu et al. (1990) in the context of repeated games of complete information with imperfect monitoring have been useful in other contexts. For instance, Fudenberg et al. (1994) derives a folk theorem for complete information games with imperfect public monitoring, Wiseman (2005) uses this approach to derive a partial folk theorem in a repeated game where players learn about a common state that determines their payoff distribution; Athey and Bagwell (2001, 2008) use it to establish the possibility of first best collusion in a dynamic Bertrand game with private (cost) information.

once trade takes place (e.g., Evans (1989), Vincent (1989), Deneckere and Liang (2006)).[4] 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 other 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.[5]) 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. This conclusion is consistent with our findings: in these "single-sale" models, once a sale is made the interaction ends. Thus, regardless of buyer competition, and similar to our repeated sale model with intra-period monopsony, the highest quality sellers have no future rents to protect. As in our model, this severely limits the market's ability to accurately screen the seller and leads to inefficiencies.

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. One important result 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.[6] However,

[4]See also Janssen and Roy (2002) for the analysis of a dynamic lemons market with decentralized equilibria. Several studies explore variations in this model. For instance, Vincent (1990) studies the impact of strategic buyer competition, Ortner (2023) 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.

[5]The comparison of Noldeke and Damme (1990) which studies public offers to 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.

[6]In the equilibria for our monopsony case, the ratchet effect manifests in the fact that a seller that reveals his type may earn zero surplus from then on; this deters learning. Freixas et al. (1985) and Laffont and Tirole (1987, 1988) establish incomplete learning due to the ratchet effect in the context of two-period principal-agent problems with spot contracting and adverse selection. In a recent contribution, Gerardi and Maestri (2020) consider an infinitely lived firm that offers short term contracts to an infinitely lived worker with private information about his persistent cost; somewhat similar to our results, the extent of learning or screening depends on beliefs; with low belief, there is no learning at all (also see other references on the ratchet effect cited in that paper).

in the context of that paper which does not feature lemons problem, failure to learn is beneficial, as it leads to fully efficient trading.

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, focusing on the characterization of trading patterns, and not on transparency. Finally, in that paper, as in Hart and Tirole (1988), the bargaining takes place between two long-lived players and our notion of transparency is not relevant. In a recent study Fuchs et al. (2022) analyzes 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 explores 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 and face no intertemporal capacity constraints. Fuchs et al. (2022) shows that 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 (2023) considers a repeated sale environment with moral hazard on the part of the seller, 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 (2022a) 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 related working paper (Kaya and Roy (2022b)), 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 observability can improve the outcomes by allowing the high quality seller to extract rents, playing a role similar to buyer competition in this paper.

2 Model

A long-lived seller can sell 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 $\theta \in \{L, H\}$ determines both the use value (v_θ) of his output and the cost (c_θ) of production. Seller's type is her private information. All buyers hold a common prior that assigns probability μ_0 to type $\theta = H$. If, in a given period, trade takes place at price P, the type-θ seller's payoff in that period is $P - c_\theta$ and her trading partner's payoff is $v_\theta - P$. Regardless of seller's type, gains from trade is strictly positive: $v_\theta - c_\theta > 0, \theta \in \{L, H\}$. We also assume that $c_H > v_L$ so that when the market's belief is low enough, the expected use value of the object for sale is less than the cost of producing high quality. The instantaneous payoff for any party who does not trade is 0. The seller maximizes the expected discounted sum of his future payoffs using discount factor $\delta \in [0, 1]$.

Assumption 1 *The seller is sufficiently patient:*[7]

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

The middle expression in Assumption 1 is the maximum payoff the low quality seller can receive once her type is revealed. Thus, the first inequality requires that she may be willing to reveal herself by not trading for one period, instead of receiving the price v_H one time. The second inequality requires that she prefers to wait two periods to receive the price c_H each period forever after, instead of revealing herself.

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 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 (infinite) 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 (infinite) 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 $\gamma_\theta(h|h')$, $\theta \in \{L, H\}$ that the seller of type s reaches history

[7] These inequalities are strictly satisfied as $\delta \to 1$. Further, if they hold for some δ, they will also hold for $\delta' > \delta$. Thus, they define a lower bound on the allowed values of δ.

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 the strategies whenever possible, and the strategies maximize each player's payoff based on their beliefs and the strategies of others.[8]

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 $h \in \mathcal{H}^{\infty}$, and any $(t' - 1)$-length history h' that h is a continuation history of, it is convenient to define $Q(h|h') = (1 - \delta) \sum_{i=t'}^{\infty} \delta^{i-t'} q_i(h)$ to be the expected discounted average trading volume along the continuation history h starting from history h'. Then, fixing an equilibrium and implied probabilities $\gamma_\theta(\cdot|\cdot)$, $\theta \in \{L, H\}$,

$$\overline{Q}_\theta(h') = \sum_{h \in \mathcal{H}^{\infty}(h')} \gamma_\theta(h|h') Q(h|h')$$

is the expected discounted average trading volume of type $\theta \in \{L, H\}$ in the continuation equilibrium. Note that, given an equilibrium, the expected gains from trade is given by

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

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 her private type. Fixing an equilibrium, throughout we let $V_\theta(h)$, $\theta \in \{L, H\}$, represent the type-s seller's continuation payoff at public history h. We express all payoffs in average per-period terms. We also note that since price offers are never observable by future buyers, in all specifications the seller strategies can be expressed with the aid of a type- and history-dependent reservation price. Let $P_\theta(h)$ represent this reservation price for type $\theta \in \{L, H\}$, at history $h \in \bigcup_{t=1}^{\infty} \mathcal{H}_t$.

[8]Formally, fixing a public history h, $\mu(h) = \frac{\mu_0 \gamma_H(h|h_0)}{\mu_0 \gamma_H(h|h_0) + (1-\mu_0)\gamma_L(h|h_0)}$, whenever the denominator is not zero.

Definition: In a given market (featuring either intra-period buyer competition or intra-period monopsony), we say that transparency is *welfare-improving* (*welfare-reducing*) if

- all equilibria of the transparent market generate larger (respectively, smaller) gains from trade than all equilibria of the opaque market, and
- at least one equilibrium of the transparent market generates strictly more (respectively, strictly less) gains from trade than any equilibrium of the opaque market.

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 her continuation payoff cannot depend on her actions. In turn, the outcome in an opaque market is the period-by-period repetition of the static market outcome.

The outcomes in such markets are shaped by two economic forces. First is the lemons problem: the high quality seller 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^* defined by

$$\mu^*(v_H - c_H) + (1-\mu^*)(v_L - c_H) = \underbrace{(1-\mu^*)(v_L - v_L)}_{=0}. \tag{2}$$

The cutoff μ^* is less than the corresponding belief cutoff μ^{**} for a market with intra-period monopsony defined by

$$\mu^{**}(v_H - c_H) + (1-\mu^{**})(v_L - c_H) = \underbrace{(1-\mu^{**})(v_L - c_L)}_{>0}. \tag{3}$$

The second economic force, "monopsony distortion," explains this discrepancy. Relative to competitive buyers, 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 high quality is higher. The next proposition formally states these observations.

The proof is omitted.

Proposition 1 *In an opaque market, regardless of buyer competition, 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 $\mu_0 > \mu^*$.*
- *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 $\mu_0 > \mu^{**}$.*

Naturally, transparency changes how the lemons problem and the monopsony distortion manifest. Next, we 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

This section provides the analysis that supports Theorem 1, reproduced below:

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 in which the market (almost) immediately learns the true quality of the seller's offering. In this equilibrium, the low quality seller trades with probability 1 each period and always at price v_L. The high quality seller trades only after market screening is complete, and therefore her trading price is always v_H. Because of this, the expected discounted volume Q_H of trade by the high quality seller must satisfy

$$v_L - c_L \geq Q_H(v_H - c_L), \tag{4}$$

so that the low quality seller does not prefer to mimic the high quality's path. The next proposition formally states the existence of such an equilibrium.

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

The proof of Proposition 2 formally describes the equilibrium strategies and beliefs and verifies that they form an equilibrium. The main thrust of the proof is the construction of a trading path for the high quality seller. This construction must imply a present discounted volume of trades that satisfy (4) while satisfying other restrictions that are implied by the dynamics. In our construction, the high quality seller's trading path features m-period long pauses of trade interspersed with n-period long streaks of trade. On the path of this equilibrium, after a first period pause of trade, the market is convinced that the product on offer is high quality. Nevertheless, trade must pause again and again. These pauses are achieved by the threat of "belief punishments." If trade unexpectedly occurs at a history when it is not supposed to, 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.[9]

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. When $\mu_0 > \mu^*$ the opaque market's outcome is efficient while this equilibrium of the transparent market inefficiently reduces the high quality's trading volume.[10] Thus, the claim of Theorem 1 for the case when $\mu_0 > \mu^*$ follows immediately. In contrast, when $\mu_0 < \mu^*$ this equilibrium improves upon the outcome of an opaque market, since in the latter, due to the lemons problem, the high quality never trades. For this range of beliefs the argument is completed by Proposition 3 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.

[9]The off-path beliefs upon observing unexpected trade need not jump to 0. There are equilibria starting with positive 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 of "maximally pooling equilibria" in the Online Appendix.

[10]The transparent market supports also an efficient pooling equilibrium. But with transparency, one cannot rule out the possibility that the market may coordinate on the inefficient separating equilibrium. 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 (e.g., Cho and Kreps (1987))—even when pooling equilibria provide higher payoffs to each party.

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 seller 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 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 Proposition 3, because neither the high quality seller nor the buyers can have negative equilibrium payoffs.

Remark (accuracy of screening and gains from trade): In addition to the complete learning equilibrium of Proposition 2, there are equilibria with partial learning where the high quality trades a positive amount. We provide a full characterization of these equilibria in the Online Appenix. In particular, we show that 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 seller 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.[11] 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 can be understood by noting that across these equilibria the low quality seller captures identical payoff $(v_L - c_L)$ while she generates less surplus (i.e. trades slower) the coarser

[11]The description of these equilibria is familiar and the constraints we present here are static. This obscures the difficulty of construction due to dynamics. We construct equilibria where the high quality's trading path features trading cycles similar to those in the separating equilibrium of Proposition 2.

is the learning. Thus, in equilibria with coarser learning, the low quality seller is cross-subsidized by the high quality's production. This reduces the high quality seller's payoff, and therefore the overall welfare since all buyers' payoffs are 0 across all these equilibria. In the coarsest possible one of these equilibria, the high quality's trading price is c_H and thus her payoff is 0. Consequently, the overall expected gains from trade is based solely on the low quality's payoff and is identical to that of the opaque market: $(1-\mu_0)(v_L-c_L)$.

Remark (seller's ability to capture surplus and accuracy of screening): In the complete learning equilibrium of Proposition 2 as well as the partial pooling equilibria, once the screening is complete, conditional on the true quality being high, the market's expectation of the use value exceeds the high quality's production cost. Nevertheless, high quality's path must feature pauses of trade. As discussed immediately following Proposition 2 these pauses are possible only when the high quality seller is willing to turn down prices that strictly exceed her cost of production, which in turn is possible because the high quality seller fears losing future surplus if she trades too frequently. Thus, the high quality seller's ability to capture surplus is a crucial factor in the market's ability to screen.

Put differently, the seller's strong bargaining power against competitive buyers allows her to capture surplus. This is what makes inefficient slow-trading equilibria possible even when the market's belief is high. In turn, when the initial belief is low, the market's ability to achieve credible accurate screening is thanks to the existence of these inefficient slow-trading equilibria at high beliefs.

5 Transparency with intra-period monopsony

Next, we turn to the markets with intra-period monopsony and present the analysis supporting Theorem 2 which is reproduced below.

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 $\mu_0 > \mu^{**}$.*

The impact of transparency in markets with intra-period monopsony is the exact opposite of its impact in markets with intra-period buyer competition. This discrepancy is explained by the differences in the two types of markets' ability to screen the seller and the cost of doing so. As discussed at the end of Section 4, the ability of a market with buyer competition to screen the seller is closely related to the high quality seller's ability to capture rents. In what follows, we demonstrate that in a market with intra-period monopsony, the high quality seller cannot capture any rents, and thus the market cannot effectively screen the seller. 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. When initial belief is low $(\mu_0 < \mu^*)$ so that an efficient pooling equilibrium does not exist in the opaque case, the market's inability to finely screen limits the gains from trade. Further, the costs of screening are inflated because the low quality seller may not receive any rents after revealing herself and thus has very strong incentives to pool with high quality, rendering transparency detrimental for gains from trade.

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 because of the monopsony distortion, even though there is no intrinsic lemons problem: 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.

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 this construction.

5.1 Learning in a transparent market with intra-period monopsony

We first formally show that in a market with intra-period monopsony, the high quality seller cannot extract any rents.

Lemma 1 *In any equilibrium of the transparent market with intra-period monopsony, at any history h, $V_H(h) = 0$. Thus, 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 when market's belief is below μ^*.

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 $\mu^* < \mu_0 \leq \mu^{**}$, then 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 three subsections, we separately take up the cases of low, intermediate and high initial beliefs.

5.2.1 Low initial beliefs: $(0, \mu^*)$

We show that when $\mu_0 < \mu^*$, 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.[12] Unlike that setting however, two factors preclude a transparent market from improving upon the level of gains from trade in an opaque market. First, 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 the case of markets with buyer competition. In fact, the belief cannot exceed μ^*, and therefore, the trading

[12] 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.5 for the formal construction.

price cannot exceed c_H. Second, 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 significantly smaller gains from trade along the pooling path.[13]

The intuition gained from this class of equilibria applies more generally. 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

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

[13]Recall that, as is the case for the maximally pooling equilibrium of a transparent market with buyer competition, when the pooling path features trading at frequency Q^* defined by $v_L - c_L = Q^*(c_H - c_L)$ while upon revealing herself the low quality seller trades efficiently, the gains from trade is exactly $(1 - \mu_0)(v_L - c_L)$, matching the opaque market's. Thus, when the trading frequency is smaller as defined in (5), the gains from trade is strictly less.

Further, $\overline{Q}_H$ is bounded by the low quality seller's incentives to mimic:[14]

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

Using the latter two inequalities to bound the total gains from trade $\mu_0\overline{Q}_H(v_H - c_H) + (1 - \mu_0)\overline{Q}_L(v_L - c_L)$ 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^*.[15]

5.2.2 Intermediate initial beliefs: (μ^*, μ^{**})

By Lemma 2, when $\mu > \mu^*$, and in particular when $\mu \in (\mu^*, \mu^{**})$, the low quality seller's payoff is no less than $\delta(c_H - c_L)$. Since the high quality seller's and the buyers' payoffs must be non-negative, this implies a lower bound $(1 - \mu_0)\delta(c_H - c_L)$ on the gains from trade in a transparent market. This, by Assumption 1, strictly exceeds $(1 - \mu_0)(v_L - c_L)$, where the latter is the gains from trade in an opaque market. However, this bound is loose. In the rest of this section we establish a strictly higher lower bound, which in the limit as $\delta \to 1$ approaches the first best gains from trade.

It is once again instructive to first discuss a specific class of equilibria for the case when $\mu_0 \in (\mu^*, \mu^{**})$. These equilibria feature partial pooling, but unlike in the case of low initial beliefs, in this case, the high quality—not the low 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 $\overline{Q}_L$. To deter the low quality seller from mimicking the high quality seller by rejecting offers in the first period, the frequency

[14] The left-hand-side reflects the highest payoff the low quality seller can receive upon revealing herself.
[15] This last assertion is established in Lemma 6 in the Appendix.

$\overline{Q}_L$ 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 first period. 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. Mimicking a path where high quality is exactly identified is very lucrative for the low quality seller, and deterring such mimicking requires that the alternative (in this case, the pooling outcome at belief μ^*) 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 $(1-\mu_0)(v_L - c_L)$.

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, completely eliminating the monopsony distortion.

5.2.3 High initial beliefs: $(\mu^{**}, 1)$

When a buyer has belief that exceeds μ^{**}, he prefers to target the high quality seller at price c_H even when the low quality seller's reservation price is as low as c_L. This observation leads to the following proposition.

Proposition 7 *If $\mu_0 > \mu^{**}$, the expected gains from trade is* $\mu_0(v_H - c_H) + (1-\mu_0)(v_L - c_L)$.

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.

The equilibrium construction is challenging. To illustrate why, consider a simple case with $\mu_0 = \mu^*$. By Proposition 4, equilibrium conditions require that the belief is never updated starting from this initial condition. Thus each type of the seller follows an identical trading path and therefore trade always takes place at price c_H. Let $\tilde{Q}$ be the (common) expected discounted frequency of trading along this path. There are many trading paths that can achieve this frequency, but not all of them can be part of an equilibrium due to the need to satisfy dynamic incentive constraints. For instance consider the "stationary" path along which each buyer offers c_H (and thus, trade takes place) with probability $\tilde{Q}$ regardless of history. In this case, each buyer's payoff is 0 and the low quality seller's reservation price is exactly c_L. But then each buyer has a profitable deviation to offering a price slightly above c_L, which would attract the low quality seller and generate a payoff close to $(1 - \mu^*)(v_L - c_L) > 0$, ruling out this stationary path. In fact, an analogous issue arises whenever the low quality's reservation price falls below v_L. Thus at any point along the equilibrium path, the future trading frequencies must be sufficiently separated after trade versus no trade so that $P_L(h) \geq v_L$. But this separation can also not be so large that the low quality seller's reservation price strictly exceeds c_H, because in that case the buyers would have a profitable deviation to targeting only the high quality seller with a price slightly above c_H. This latter requirement rules out a path along which all trade is "back-loaded," i.e. trade takes place with probability 1 each period after a sufficiently long initial pause. Thus, the equilibrium path must necessarily feature cyclical trading. These difficulties along with others that arise at beliefs different from μ^* makes direct construction of equilibrium strategies intractable. Because of this difficulty, our strategy of characterization appeals to dynamic programming arguments similar to the techniques developed in Abreu et al. (1990) and Fudenberg et al. (1994).

Our formal steps characterize sets of equilibrium payoffs for the low quality seller. In the process we also identify all equilibrium strategies and beliefs.[16]

The first step of our construction is to define "enforceable" payoffs for the low quality seller at each possible belief μ. For a payoff U to be enforceable at belief $\tilde{\mu}$ it must be an equilibrium payoff for the low quality seller in a one-shot game defined by continuation payoffs U^A and U^R so that if the seller of type L trades at price P, she receives a payoff of $(1-\delta)(P-c_L)+\delta U^A$, and if she fails to trade she receives a payoff δU^R. Of course, such an equilibrium will also specify an offer strategy for the buyer as well as a strategy for the H-type seller. Enforceability requires that the continuation payoffs U^A and U^R, under which U is an equilibrium payoff, are themselves equilibrium payoffs at some beliefs μ^A and μ^R respectively, *and* μ^A and μ^R are derived using Bayes rule from the strategies of the equilibrium enforcing U whenever, respectively, trade and no trade occur with positive probability in that equilibrium.

Since our setting is quite specific, instead of stating equilibrium conditions that define enforceability of payoffs in full generality, we take advantage of the specific structure of the setting. In particular, we make use of the following observations:

First, each pair of continuation payoffs U^A and U^R defines a reservation price P_L for the low quality seller by

$$(1-\delta)(P_L - c_L) = \delta(U^R - U^A). \tag{6}$$

Then, **the L seller's strategy** can be fully described by P_L and his probability of accepting P_L when offered. Let β stand for this probability.

Second, note that one can focus on buyer offer strategies whose support is contained in $\{P_L, c_H\}$.[17] Thus, **a buyer's strategy** can be fully described by the probability that he offers c_H. Let α stand for this probability. Further, for the optimality of the buyer strategy

[16] Since the high quality seller's payoff is always 0 and all buyers are myopic, the low quality seller is the only one with dynamic considerations.

[17] Any equilibrium where the buyer makes a losing offer is payoff equivalent to one where the buyer offers P_L which is rejected with probability 1.

it is necessary that

$$\alpha = \begin{cases} 1 & \text{if } \tilde{\mu}(v_H - c_H) + (1-\tilde{\mu})(v_L - c_H) > (1-\tilde{\mu})(v_L - P_L)\beta \\ \in [0,1] & \text{if } \tilde{\mu}(v_H - c_H) + (1-\tilde{\mu})(v_L - c_H) = (1-\tilde{\mu})(v_L - P_L)\beta \\ 0 & \text{if } \tilde{\mu}(v_H - c_H) + (1-\tilde{\mu})(v_L - c_H) < (1-\tilde{\mu})(v_L - P_L)\beta \end{cases} . \tag{7}$$

Finally, **the L-seller's payoff** can be calculated along the possibly off-equilibrium path where she rejects her reservation price when offered. Thus,

$$U = \alpha[(1-\delta)(c_H - c_L) + \delta U^A] + (1-\alpha)\delta U^R. \tag{8}$$

Because our setting features private information, and since equilibrium payoffs naturally vary with initial belief, the object that we are interested in characterizing is not simply a set of payoffs but is a correspondence that maps each belief into a set of payoffs for the low quality seller. Let $\{\mathcal{U}_\mu\}_{\mu\in[0,1]}$ be such a correspondence. Given a belief μ, if a payoff U for the low quality seller can be obtained with continuation beliefs and payoffs conforming to this correspondence, we say that U is enforceable with respect to $\{\mathcal{U}_\mu\}_{\mu\in[0,1]}$ and belief μ. We next provide the formal definition.

Definition: For each $\mu \in [0,1]$ let $\{\mathcal{U}_\mu\}_{\mu\in[0,1]}$ be a set of potential payoffs for the low quality seller. 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 $\alpha \in [0,1], P_L \leq c_H$, $\mu^A, \mu^R \in [0,1], U^A \in \mathcal{U}_{\mu^A}, U^R \in \mathcal{U}_{\mu^R}$ that satisfy (6), (7) and (8) and

$$\mu^A = \frac{\mu\alpha}{\mu\alpha + (1-\mu)(\alpha + (1-\alpha)\beta)} \text{ and } \mu^R = \frac{\mu(1-\alpha)}{\mu(1-\alpha) + (1-\mu)(1-\alpha)(1-\beta)} \tag{9}$$

whenever the denominators are positive.

Naturally, the sets of payoffs that are enforceable in our setting will depend on the market's current belief about the seller's type. Thus, the enforceability of a given payoff for the low quality seller is defined at a specific belief, and, as in Abreu et al. (1990), with respect to a set of continuation payoffs. Then, we will say that a correspondence

$\{\mathcal{U}_\mu\}_{\mu\in[0,1]}$ is self-generating, if all of its elements can be enforced with respect to itself. That is,

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

Then, it is easy to see that, similar to the case of Abreu et al. (1990), any such self-generating correspondence defines equilibrium payoffs, as formally stated in Proposition 8.

Proposition 8 *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. ■

5.3.1 Self-generating sets of payoffs

In the Appendix, Proposition 9 formally characterizes a self-generating correspondence $\{\mathcal{U}\}_{\mu\in[0,1]}$. In this section, we informally describe this correspondence, highlight some properties of this correspondence and the associated equilibrium behavior.

Equilibria when $\mu_0 = \mu^*$: The correspondence characterized in Proposition 9 has

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

Thus in particular, this set of payoffs is an interval.

As discussed above, along the path of an equilibrium starting with belief μ^*, the belief is never updated so that on-path continuation payoffs must all come from $\mathcal{U}_{\mu^*}$. Even so, we find that, often it is not possible to characterize $\mathcal{U}_{\mu^*}$ in isolation.[18] To see why note

[18] 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 needed. In the Online Appendix C we demonstrate that when $c_H - v_L > v_L - c_L$, there exists a strict subset $\mathcal{U}$ of $\mathcal{U}_{\mu^*}$ such that all $U \in \mathcal{U}$ are enforeceable at μ^* with respect to $\mathcal{U}$.

that, as discussed above, for belief to be never updated, the low quality seller's reservation payoff must be between v_L and c_H. That is, it is necessary that

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

In particular, the spread between U^R and U^A must be large enough. An implication of this requirement is that the payoffs that are close to the two ends of the interval $\mathcal{U}_{\mu^*}$ may not be enforceable with continuation payoffs $U^A, U^R \in \mathcal{U}_{\mu^*}$. In our construction, such values are enforced using belief punishments or rewards. In particular, an equilibrium that delivers a payoff close to the lower end of the interval $\mathcal{U}_{\mu^*}$ features zero probability of trade. Here, 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. 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.

Equilibria when $\mu_0 < \mu^*$ The multiplicity of equilibria when $\mu = \mu^*$ allows for multiple equilibria when $\mu = 0$, using belief rewards when needed. In particular, when $\mu = 0$, there exist fixed price equilibria where each buyer offers a specific P with $P \in [c_L, v_L]$ with probability 1 and P is accepted with probability 1. Such equilibria can be sustained by choosing $U^R = (P - c_L)/\delta$ ensuring that the low quality seller's reservation price is exactly P.[19] 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. 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

[19]If U is large, this may require off-path beliefs $\mu^R = \mu^*$. Please see the Appendix for formal details.

trade the belief is μ^*. Then, $U^A \in \mathcal{U}_0, U^R \in \mathcal{U}_{\mu^*}$ are chosen such that

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

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 $U^A = v_L - c_L$, and is equal to the gains from trade $(1-\mu_0)(v_L - c_L)$ in opaque markets established in Proposition 5. For any other U^A, the pooling path features less trade and thus the gains from trade is smaller. [20]

Equilibria when $\mu_0 > \mu^*$ Unlike in the case for low initial beliefs, we construct a single equilibrium for almost all initial beliefs in this range.[21] These equilibria feature finitely many periods of screening 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.

The convergence path is described with respect to an increasing sequence of cutoff

[20]It is interesting to note that, when $\mu_0 \in (0, \mu^*)$, there may exist equilibria where trade at the initial period is interpreted as good news about quality. In such an equilibrium, at the null history, the buyer offers c_H, which is accepted with probability 1 by the high quality seller and with appropriate probability by the low quality seller so that $\mu^A = \mu^*$, which also implies that $\mu^R = 0$. The incentives of the low quality seller can be satisfied by choosing $U^R \in \mathcal{U}_0, U^A \in \mathcal{U}_{\mu^*}$ such that $\delta(U^R - U^R) = (1-\delta)(c_H - c_L)$. Note that this is possible as long as $\delta^2(v_L - c_L) > (1-\delta)(c_H - c_L)$. Existence of such equilibrium highlights the implications of the failure of the "skimming property" in this model of repeated sales.

[21]For a countably many initial beliefs that constitute cutoffs in the screening process, there are multiple equilibria as described below.

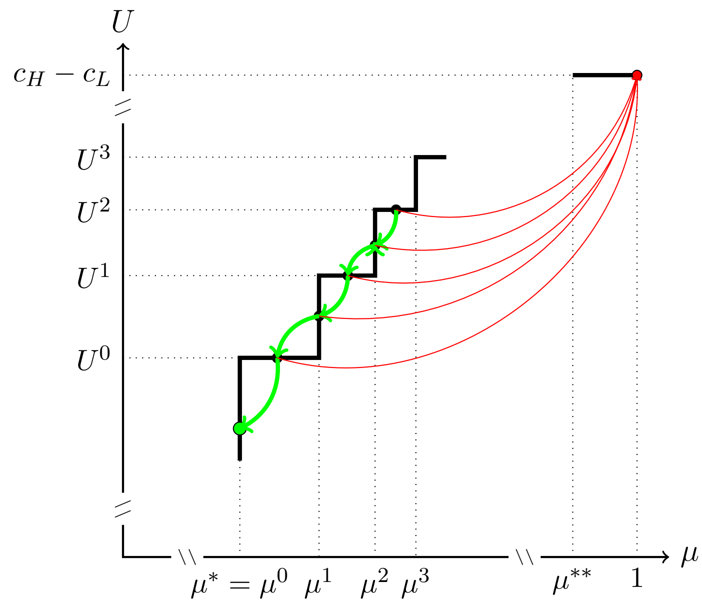

Figure 2: The heavy black step function maps a subset of beliefs $[\mu^*, \mu^{**})$ to equilibrium payoffs of the low quality seller at those beliefs. The (green) heavy arrows show how the (belief, payoff) pairs change after trade. The (red) light arrows indicate that after each failure of trade, the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

beliefs $\{\mu^k\}_{k=0}^{\infty}$ with $\mu^0 = \mu^*$ and $\lim_{k\to\infty} \mu^k = \mu^{**}$. Thus, this sequence defines a partition of the interval $[\mu^*, \mu^{**}]$ of beliefs. As illustrated in Figure 2, the low quality seller's payoff is constant, say U^{k-1} over beliefs $\mu \in (\mu^{k-1}, \mu^k)$, while at the cutoff beliefs μ^k there are equilibria that support all payoffs $[U^{k-1}, U^k]$.

We define the sequences $(\mu^0, \mu^1, \cdots)$ and $(U^0, U^1, \cdots)$ recursively, starting with $\mu^0 = \mu^*$. In all equilibria that deliver these payoffs, the buyer randomizes between offering P_L and c_H. The low quality seller accepts both offers with probability 1, while the high quality seller accepts c_H with probability 1 and rejects P_L with probability 1. Consequently, $\mu^R = 1$ and $U^R = c_h - c_L$. Further, for $\mu \in (\mu^0, \mu^1)$, the probability α that the buyer offers c_H is such that by (9), $\mu^A = \mu^0$. Note that since the buyer's payoff from offering c_H (as opposed to P_L) increases as μ increases, at higher beliefs, the buyer's indifference requires smaller P_L. Then, by (6), and because the continuation payoff upon no trade is fixed at $U^R = c_H - c_L$, U^A must increase as μ increases. Then, the next cutoff belief μ^1 is specified as the belief at which the required U^A is the upper boundary of the interval

$\mathcal{U}_{\mu^*}$ ($\equiv \mathcal{U}_{\mu^0}$).

To see how the payoff U^0 associated with beliefs $\mu \in (\mu^0, \mu^1)$ is constructed, we first note that using (6), equation (8) describing the low quality seller's enforced payoff can be re-written as:

$$U = (1-\delta)\alpha(c_H - P_L) + \delta(c_H - c_L) \tag{11}$$

Further, the buyer's indifference requirement along with the fact that $\mu^A = \mu^0$ yields[22]

$$\frac{\mu^0}{1-\mu^0}\frac{1}{\alpha} = \frac{\mu}{1-\mu} = \frac{c_H - P_L}{v_H - c_H}, \tag{12}$$

which uniquely pins down the quantity $\alpha(c_H - P_L)$. Combining (11) and (12) establishes that as μ varies over the interval (μ^0, μ^1)—even though α and P_L vary—the enforced payoff U^0 is uniquely pinned down.

Now, when the belief is exactly μ^1, as the buyer's probability of offering α varies, μ^A varies over (μ^0, μ^1), and thus $U^A = U^0$. This uniquely pins down P_L at the level which guarantees the buyer's indifference.[23] As α varies, by (11), the enforced payoff also varies. The payoff U^1 is specified as the largest payoff that can be enforced in this manner, which corresponds to $\alpha = 1$.

Iterating these steps, we recursively describe the sequences $(\mu^0, \mu^1, \cdots)$ and $(U^0, U^1, \cdots)$. Figure 2 illustrates (a portion of) the correspondence $\{\mathcal{U}_\mu\}_{\mu \in [\mu^*, \mu^{**}]}$. The arrows describe the equilibrium path starting from some belief $\mu \in (\mu^2, \mu^3)$. If trade takes place, belief updates to $\mu^A = \mu^2$. If trade continues to take place at each subsequent period, the belief declines as shown by the green arrows in Figure 2 and eventually reaches μ^* after five consecutive periods of trade. If trade fails to take place in any of those five periods, belief updates to $\mu^R = 1$. Once belief reaches μ^* it is never updated on path, and one of the equilibria described above for $\mu = \mu^*$ is played.

[22]The first equality guarantees $\mu^A = \mu^*$ and follows from (9). The second equality is the buyer's indifference condition and follows from (7).

[23]In our construction, U^0 is equal to the upper boundary of the interval $\mathcal{U}_{\mu^*}$.

6 Discussion

Our results highlight the crucial role that a seller's ability to capture rents plays in determining the markets' ability to distinguish her from sellers of lower quality and the cost of doing so. Specifically, Lemma 1 establishes that with monopsony buyers, the high quality seller cannot capture any rents, and consequently trade cannot be slowed down at high beliefs, rendering mimicking very profitable for the low quality seller (Lemma 2). This implies that substantial pooling is unavoidable. And finally, such pooling is inefficient due to the fact that it features cross-subsidization of the lower quality seller by the high quality seller. (Proposition 5). In contrast, in markets with competitive buyers, the high quality seller can capture rents. Thus, the above-mentioned mechanism which creates an upper bound on the gains from trade fails, allowing for higher gains from trade in transparent markets.

The starkness of our modeling choices—i.e., two quality types, extreme forms of buyer competition and observability of trades without noise— allows this mechanism to take center stage. We anticipate that this mechanism would continue to play a role in different, less stark, environments, even though it may be confounded by other strategic issues. Here we discuss some alternative specifications and our conjectures on how the insights from this paper may shed light on those.

Noisy observability of histories: An intermediate model between opaque and transparent market specifications is one where future buyers observe past trades with some noise. There are various ways in which this noise can be modeled. One reasonable way to do so, for instance, is to allow future buyers to sometimes fail to observe trades while not allowing them to mistakenly believe that there was trade when there was none. We conjecture that our results both for the competitive buyers case and for the monopsony case are robust to the introduction of such noise. For either type of market when the probability of such failure to observe is close to one, (i.e, the market is almost opaque) the time it would take for the market to be convinced to offer high enough prices that high quality seller may accept would diverge to infinity, essentially precluding the trade of high quality. This in turn almost allows the low quality to trade efficiently since it would eliminate the incentives to mimic. Thus, the outcome would be close to the outcome of an opaque market. Second, when the probability of the failure to observe is small (i.e., the market is almost transparent), learning from the pauses of trade will be sufficiently precise that analogues of all partial pooling equilibria can likely be constructed. An interesting observation is that in the monopsony case, with this type of noisy signal, keeping the belief from exceeding μ^* would require less inefficient pooling by the low quality seller, potentially allowing higher gains from trade.

The general case for intermediate levels of noise or other specifications of observability of past trades would likely be intractable.[24] Yet, it is easy to see that the high quality seller's inability to capture rents with monopsony buyers will readily extend (Lemma 1). Further, the market's inability to slow down trade when the belief exceeds μ^* should become even more severe if there is a chance that trading records will not be observed by future buyers. These observations suggest that the forces that lead to the payoff bounds in our model will continue to play a role.

Intermediate levels of competition: If the number of arriving buyers is stochastic (with support possibly including zero to capture thinner markets), questions about different forms of transparency arise, such as whether future buyers observe the realized number of buyers the seller encountered in prior periods. If they do, and if there is positive probability that multiple buyers arrive, the high quality seller would be able to capture rents. Thus, we conjecture that in this case the outcomes will be qualitatively similar to our competitive buyers case, with potentially stricter restrictions on the seller's discount factor. By the same token, if the market is very thin so that either a single buyer arrives or no buyers arrive, then the outcomes are likely to be qualitatively similar to our intra-period monopsony case. When the future buyers do not observe the past market conditions, similar conjectures apply but the analysis would be complicated due to considerations analogous to those in the case of noisy observability.

Limited records: In an earlier paper (Kaya and Roy (2022a)), we consider a repeated lemons market with *competitive* buyers who can observe finite-length records of past trad-

[24] However, see Dilme (2022) for an analysis of a case with noisy signals in a Coasian environment.

ing outcomes. We study the welfare implications of increasing the record lengths.[25] In the class of equilibria considered in that paper, if at some history beliefs exceed μ^*, the price is driven up to the expected value of the offering. Unlike in the competitive case considered in this paper, such prices are always accepted with probability 1. This is because, with finite record lengths, the loss of reputation due to too frequent trading—which in this paper is the threat that stops high quality from trading—is necessarily short-lived. On the flip side, any reputation built up must be frequently renewed. This exogenously slows down the frequency at which high prices are offered, dampening the low quality's incentives to mimic. As record lengths grow, this latter effect becomes weaker and it becomes necessary to lower the trading prices in order to avoid mimicking, which in turn is possible only via inefficient pooling.

Price observability: In a related working paper (Kaya and Roy (2022b)), we demonstrate that, even in a market with intra-period monopsony, if trading prices (in addition to the trading volumes) are observed, then the high quality buyers are able to collect rents, and as a result accurate screening of buyer types becomes possible, improving efficiency. This result further highlights the crucial role that the seller's ability to capture rents plays in the market's ability to screen.

More than two types: Extending our equilibrium construction and our arguments about payoff bounds to more than two types presents some novel challenges. Two features of the transparent markets with monopsony buyers that nevertheless easily generalize are (i) the inability of the *highest* type seller to capture rents when facing a sequence of monopsonist buyers (Lemma 1); and, due to this, (ii) the inability of the market to trade slowly when the expected value of the seller's offering exceeds cost of the highest quality (Lemma 2). These observations suggest that, similar to the case with two types, when starting in a lemon's market, the beliefs can never exceed this threshold, and thus all equilibria must feature some inefficient pooling. Different from the two-type case, the possibility of

[25] Thus, the competitive buyers version of the model of this paper can be viewed as the two extreme cases of that paper's model with record length being 0 (opaque markets) or infinity (transparent markets). Apart from focusing only on competitive buyers, that paper also only considers $\mu_0 < \mu^*$ and restricts attention to a specific class of equilibria.

"finer" separation of some intermediate types (who will necessarily receive information rents) cannot be ruled out by our mechanism. Thus, whether, in balance, transparency would increase or decrease the overall gains from trade is likely to depend on the prior distribution of types, which also determines the types that are able to trade in an opaque market.

Appendix

A Proofs for Section 4

A.1 Proof of Proposition 2

We construct an equilibrium in which the low quality trades with probability 1 each period at price v_L while the high quality's trading price is always v_H and her expected discounted volume of trade is some Q_H satisfying

$$\frac{v_L - c_L}{v_H - c_L} > Q_H > 1 - \delta. \tag{13}$$

Note that by Assumption (1), $(v_L - c_L)/(v_H - c_L) > 1 - \delta$, and therefore this interval is non-empty.

Next we construct a trading path for the high quality seller along which she trades at such frequency and then verify that this path is an equilibrium outcome. Define Q_{mn} to be the expected discounted frequency along a path that starts with m periods of no trade followed by n-period streaks of trade interspersed with m-period pauses. Then, $Q_{mn} = \delta^m \left[(1-\delta)(1+\delta+\delta^2+\cdots+\delta^{n-1}) + \delta^n Q_{mn}\right]$, so that

$$Q_{mn} = \frac{\delta^m + \delta^{m+1} + \cdots + \delta^{m+n-1}}{1+\delta+\cdots+\delta^{m+n-1}} = \delta^m \frac{1-\delta^n}{1-\delta^{m+n}}.$$

Next we claim that there exists m, n such that $Q_H = Q_{mn}$ satisfies (13). Note that for a fixed m, Q_{mn} is an increasing and convergent sequence (in n) with limit δ^m. Choose m^* such that $\delta^{m^*-1} \geq \frac{v_L - c_L}{v_H - c_L} > \delta^{m^*}$. By the first inequality, $\delta^{m^*} \geq \delta \frac{v_L - c_L}{v_H - c_L}$ which, by

Assumption 1, implies that $\delta^{m^*} > 1 - \delta$. Since $lim_{n \to \infty} Q_{m^* n} = \delta^{m^*} \in \left(1 - \delta, \frac{v_L - c_L}{v_H - c_L}\right)$, there exists large enough n^* such that $Q_{m^* n^*}$ satisfies (13). Fix m, n that satisfy (13).

Next, we describe strategies and beliefs that support the trading path described above as the high quality's equilibrium path. Define $h_\emptyset^s$ to be the s-length history featuring no trade and h_I^s to be the s-length history featuring trade in each period. Let (h, h') be a history obtained by appending history h' after history h. Therefore, $(h, h', h''', \cdots, h'''')$ is a history formed by appending the indicated histories after each other. Consider the following classification of histories:

- The null history: $h_\emptyset$.
- Histories that are shorter than m periods and feature no trade: $h_\emptyset^k$, $k < m$.
- Histories that feature trade in all previous periods: $\mathcal{H}^L = \{h_I^s \mid s = 1, 2, \cdots\}$.
- Histories that start with at least m periods of no trade and feature cycles between trading streaks of at most n periods interspersed with trade pause streaks of at least m periods, which are subcategorized below. In this representation each history has R streaks of trade pause, alternating with R or $R-1$ streaks of trading periods. The lengths of the streaks of trade pauses are represented by $s_1, \cdots, s_R$ and the lengths of streaks of trading are represented by $k_1, \cdots, k_R$.[26]
  - ending with a shorter than n-period streak of trade:

$$\mathcal{H}_{I,<n}^H = \left\{(h_\emptyset^{s_1}, h_I^{k_1}, h_\emptyset^{s_2}, h_I^{k_2}, \cdots, h_\emptyset^{s_R}, h_I^{k_R}) \mid k_R < n, s_i \geq m, 0 < k_j \leq n, 1 \leq i \leq R, 1 \leq j < R\right\}$$

  - ending with an n-period streak of trade:

$$\mathcal{H}_{I,n}^H = \left\{(h_\emptyset^{s_1}, h_I^{k_1}, h_\emptyset^{s_2}, h_I^{k_2}, \cdots, h_\emptyset^{s_R}, h_I^{k_R}) \mid k_R = n, s_i \geq m, 0 < k_j \leq n, 1 \leq i \leq R, 1 \leq j < R\right\}$$

  - ending with an m-period streak of no trade:

$$\mathcal{H}_{\emptyset,m}^H = \left\{(h_\emptyset^{s_1}, h_I^{k_1}, \cdots, h_I^{k_{R-1}} h_\emptyset^{s_R}) \mid s_R = m, s_i \geq m, k_i \leq n, i = 1, \cdots, R-1\right\}$$

[26]For instance $h = (0, 0, 0, 1, 1, 0)$ has $R = 2, s_1 = 3, s_2 = 1, k_1 = 2$, features $R = 2$ streaks of trade pause and $R - 1 = 1$ streak of trading, and it can be alternatively represented as $(h_\emptyset^3, h_I^2, h_\emptyset^1)$.

– ending with a shorter than m-period streak of no trade:

$$\mathcal{H}^H_{\emptyset,<m} = \left\{(h_\emptyset^{s_1}, h_I^{k_1}, \cdots, h_I^{k_{R-1}} h_\emptyset^{s_R}) \mid s_R < m, s_i \geq m, k_i \leq n, i = 1, \cdots, R-1\right\}$$

- $\mathcal{H}^{off}$: all histories that are not in $\{h_\emptyset\} \cup \{h_\emptyset^k \mid k < m\} \cup \mathcal{H}^L \cup \mathcal{H}^H_{I,<n} \cup \mathcal{H}^H_{I,n} \cup \mathcal{H}^H_{\emptyset,m} \cup \mathcal{H}^H_{\emptyset,<m}$.

Beliefs If $h \in \mathcal{H}^L \cup \mathcal{H}^{off}$, then $\mu(h) = 0$. Otherwise, $\mu(h) = 1$. Note that the set $\mathcal{H}^L$ contains all histories that are on the path for the low quality seller, and not for the high quality seller, justifying the belief on this set. All histories in $\mathcal{H}^{off}$ are off the equilibrium path and feature shorter streaks of pause and/or longer streaks of trade than expected on the path for the high quality seller. The beliefs at these histories assign all weight to the low quality. All other histories are either on the path for the high quality seller (justifying the belief via Bayes rule) or are off the equilibrium path and feature longer streaks of pause and/or shorter streaks of trade than expected on the path for the high quality seller.

Buyer strategies Buyers offer v_L if $h \in \{h_\emptyset\} \cup \mathcal{H}^L \cup \mathcal{H}^{off}$. Otherwise they offer v_H. That is, the buyers offer v_L at the null history and when the belief is 0. They offer v_H when their belief is 1.

Seller strategies The seller uses a type- and history-dependent reservation price $P_\theta(h)$, described below and always accepts any offer that weakly exceeds her reservation price.

- for $h = h_\emptyset$ or $h \in \mathcal{H}^H_{I,n}$: $(1-\delta)(P_\theta(h) - c_\theta) + \delta \max\{v_L - c_\theta, 0\} = Q_{mn}(v_H - c_\theta)$.

  Here, the left-hand-side is the payoff that the seller of type θ may obtain if she accepts an offer of $P_\theta(h)$, taking into account the buyer strategies which would offer v_L with probability 1 forever thereafter. The right-hand-side is the payoff the seller receives if she follows the high quality seller's trading path. Note that by Assumption (1) and construction of Q_{mn}, $P_H(h) > v_H$ and $P_L(h) \leq v_L$.

- for $h = h_\emptyset^s$, $s < m$: $(1-\delta)(P_\theta(h) - c_\theta) + \delta \max\{v_L - c_\theta, 0\} = \frac{Q_{mn}}{\delta^s}(v_H - c_\theta)$. This is a history which features no trade and is of length less than m. The left-hand-side of the equality takes into account that if trade takes place at this history, the buyers will offer v_L forever after. The right-hand-side is the payoff from continuing to follow the high quality seller's trading path. Again, by Assumption (1), $P_H(h) > v_H$.

- for $h \in \mathcal{H}^L \cup \mathcal{H}^{off}$, $P_\theta(h) = c_\theta$. This takes into account the fact that at any continuation history following h, the buyers will always offer v_L.
- for $h \in \mathcal{H}^H_{I,<n}$, letting k_R represent the length of the latest streak of trading,

$$(1-\delta)\left[P_\theta(h) - c_\theta + (\delta + \cdots + \delta^{n-k_R-1})(v_H - c_\theta)\right] + \delta^{n-k_R} Q_{mn}(v_H - c_\theta) = Q_{mn}(v_H - c_\theta).$$

  Here, the left-hand-side of the equality is the payoff from trading at $P_\theta(h)$ for one period and then trading at v_H for the next $n - k_R - 1$ periods, and then reverting to the high quality's trading path which alternates m periods of pause with n periods of trading at price v_H. The right-hand-side reflects the fact that if the seller chooses not to trade today, then this will count as the first period of reverting to the high quality seller's trading path. Here, it is easy to see that $P_\theta(h) < v_H$.
- for $h \in \mathcal{H}^H_{\emptyset,<m}$, letting s_R represent the length of the latest streak of trading pause,

$$(1-\delta)(P_\theta(h) - c_\theta) + \delta \max\{v_L - c_\theta, 0\} = \frac{Q_{mn}}{\delta^{s_R}}(v_H - c_\theta).$$

  The left-hand-side of the equality takes into account the fact that if trade takes place at this history, at all continuation histories, the buyers will offer v_L. The right-hand-side takes into account the fact that this is the s_R^{th} period of a trading pause. Again, since $Q_{mn} > (1-\delta)$, we have $P_H(h) > v_H$.
- for $h \in \mathcal{H}^H_{\emptyset,m}$:

$$(1-\delta)\left[P_\theta(h) - c_\theta + (\delta + \cdots + \delta^{n-1})(v_H - c_H)\right] + \delta^n Q_{mn}(v_H - c_H) = \frac{Q_{mn}}{\delta^{m-1}}(v_H - c_\theta).$$

  Here, the left-hand-side of the equality is the payoff from trading at $P_\theta(h)$ for one period and at v_H for the following $n-1$ periods, and then reverting to the high quality's trading path which alternates m periods of pause with n periods of trading at price v_H. The right-hand-side reflects the fact that if the seller chooses not to trade today, then he can start an n-period streak of trading at v_H next period. Here, it is once again easy to see that $P_H(h) > v_H$.

Equilibrium path: If the seller and the buyer use above strategies, the low quality seller trades in each period. High quality seller's trade follows a pattern that starts with m periods of no trade, followed by alternating between n periods of trade and m periods of no trade, as anticipated.

Verification: Given the equilibrium path, the beliefs are consistent with these strategies by construction. For the optimality of buyer strategies at $h = h_\emptyset$, we note that $P_L(h_\emptyset) \leq v_L < \mu_0 v_H + (1 - \mu_0) v_L < v_H < P_H(h_\emptyset)$. Thus, when all other buyers are offering v_L, each buyer's best response is also to offer v_L. At all histories either $\mu = 0$ or $\mu = 1$, and therefore it is trivial to see that all buyers offering v_L (or, respectively v_H) forms a bidding equilibrium. The seller's reservation prices are calculated directly from the buyer offer strategies as discussed throughout the construction, and are therefore optimal.

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. *Trading price is never less than v_L.*

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

Proof of Lemma 4.

1. Fix an on or off-path history h with $\mu(h) < 1$. Suppose for a contradiction that at this history, trade takes place at a price $P < v_L$ with positive probability. This implies that the buyers' payoffs are positive because an offer of $P + \varepsilon$ wins with positive probability and is less than v_L for $\varepsilon > 0$ and small. Then, (i) the infimum

$\underline{P}$ 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, and (ii) 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. Let $\eta = \Pi_{i \neq 1} F_i(\underline{P}) > 0$ be the probability that all buyers except buyer 1 tie at price $\underline{P}$. Next, we argue that buyer 1 has a profitable deviation to offer $\underline{P} + \varepsilon$ for some $\varepsilon > 0$.

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. Let $\gamma_\theta^\varepsilon$ be the corresponding probability when $\underline{P} + \varepsilon$ is the highest offer. First note that the probability that buyer 1 trades if he offers $\underline{P}$ is at most $\eta/2$. In contrast, if he offers $\underline{P} + \varepsilon$, then he trades with probability at least η. The latter is because all other buyers offer tie at $\underline{P}$ with probability η and for each θ, $\gamma_\theta \leq \gamma_\theta^\varepsilon$, which in turn is because each type of the seller uses a reservation price strategy.

Next, conditional on trading at price $\underline{P}$, buyer 1's payoff is

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

Conditional on trading at price $\underline{P} + \varepsilon$ buyer 1's payoff is

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

Thus, buyer 1's payoff is no larger than $\eta\Pi/2$ if he offers $\underline{P}$, and no less than $\eta\Pi^\varepsilon$ if he offers $\underline{P} + \varepsilon$. We note that

$$\begin{aligned}\Pi - \Pi^\varepsilon &= \mu(h)(\underbrace{\gamma_H - \gamma_H^\varepsilon}_{\leq 0})(\underbrace{v_H - \underline{P}}_{>0}) + (1 - \mu(h))(\underbrace{\gamma_L - \gamma_L^\varepsilon}_{\leq 0})(\underbrace{v_L - \underline{P}}_{>0}) + \varepsilon(\underbrace{\gamma_H^\varepsilon \mu(h) + \gamma_L^\varepsilon(1 - \mu(h))}_{\in(0,1]}) \\ &\leq \varepsilon(\gamma_H^\varepsilon \mu(h) + \gamma_L^\varepsilon(1 - \mu(h)).\end{aligned}$$

Thus, for small enough ε, $\Pi\eta/2 < \Pi^{\varepsilon}\eta$, i.e. buyer 1 strictly prefers to offer $\underline{P} + \varepsilon$ rather than $\underline{P}$. This is a contradiction and establishes item 1.

2. For item 2, assume that $P_L(h) < v_L$ and suppose that 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, as above, that the infimum $\underline{P}$ of the support of each buyer's bid distribution is the same. Further, $\underline{P} \leq P_L(h)$ because the low quality trades with probability less than 1. Finally, this offer is accepted with positive probability, because the buyers' payoffs are positive and thus they would not optimally make a losing offer. This contradicts item 1.

■

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 (acceptance) and no trade (rejection), respectively. Since $V_A \geq \underline{V}_L$, we have

$$(1-\delta)(v_L - c_L) + \delta \underline{V}_L \leq \delta V_R. \tag{14}$$

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

$$V_L(h) = \underline{V}_L + \varepsilon_2 \geq \delta V_R, \tag{15}$$

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

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

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 in both cases, $\underline{V}_L + \frac{\varepsilon_2}{1-\delta} \geq v_L - c_L$, and $\varepsilon > \varepsilon_2$ can be chosen arbitrarily close to zero, it 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 gains from trade in the market is at least $(1-\mu_0)(v_L - c_L)$. ■

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.

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

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

where Q_H is the high quality seller's trading frequency and, by construction, is larger than $1-\delta$, and in particular strictly positive. This establishes that transparency is welfare improving when $\mu_0 < \mu^*$.

B 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, $V_H(h) = 0$. Thus, the high quality seller accepts any offer that exceeds c_H.*

Proof. Let $\bar{V}_H = sup\{V_H(h)|h \in \mathcal{H}\}$. Suppose $\bar{V}_H > 0$. 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,

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

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 $V_H(h) = 0$ for any h. ■

In addition to Lemma 1, the following result which was not discussed in the text, plays a role in the proofs of Lemmas 2 and 3.[27]

Lemma 5 *In any equilibrium of the transparent market with intra-period monopsony, at any 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, a buyer arriving at this history never offers 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, $V_L(h, A) = c_H - c_L \geq V_L(h, R)$, 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 $\mu(h) > \mu^*$. This, together with Lemma 5, implies that the buyer's payoff is strictly positive, thus he makes no offers that will be rejected with probability 1. 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. Further, it must be true that $P = P_L(h)$. This in turn implies that low quality seller trades with probability 1: if $P < v_L$ were being rejected with positive probability, the buyer would have a profitable deviation to a slightly higher offer. Then, $\mu(h, R) = 1$, thus $V_L(h) \geq \delta(c_H - c_L)$.

Now consider histories where trade takes place only at price c_H. Let

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

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, $V_L(h^*) = (1-\delta)(c_H - c_L) + \delta V_L(h^*, A)$,

[27] This is a weak version of the so-called "skimming property," and unlike in the models of single sale, follows from equilibrium conditions rather than primitives of the model.

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. Thus,

$$\underline{V}_L > \delta(1-\delta)(c_H - c_L) + \delta V_L(h^*, A). \tag{16}$$

Let α_θ be the probability with which type-θ seller accepts c_H. If either α_H or α_L is different from 1, 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 (i.e. the buyer does not strictly prefer to do so) if and only if $\mu(h^*, R) \leq \mu^*$. This in turn implies that $\mu(h^*, A) > \mu^*$ since $\mu(h^*) > \mu^*$. If $\alpha_L = \alpha_H = 1$, then $\mu(h^*, A) = \mu(h^*) > \mu^*$. In either case, $V_L(h^*, A) \geq \delta(c_H - c_L)$. The claim follows by substituting the bound for $V_L(h^*, A)$ into (16). ■

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 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 at which (i) equilibrium probability of trade is positive, and (ii) $P_L(h_1) < c_H$.

The rest of the argument has two parts: (step 1) we argue that such h_1 exists; (step 2) we argue that trade at such h_1 reveals low quality.

To see that such h_1 exists (step 1), first suppose for a contradiction that it does not. Note that the buyer never offers $P > c_H$. Thus, along this path the low quality never trades either because the probability of trade is 0 or the buyer's offer is no larger than his reservation price. Then, $V_L(h, R) = 0$, which in turn implies that $P_L(h) \leq c_L$. But then, at h the buyer's payoff is strictly positive (which he can guarantee by offering $c_L + \varepsilon$ with $\varepsilon > 0$ and small), and thus trade must take place with positive probability, and thus $h = h_1$, a contradiction.

To argue step 2, we start by showing that $\mu(h_1) < \mu^*$. To see this first note that h_1 has the form $(h, R, \cdots, R)$ because until h_1, either the probability of trade is 0 or $P_L(h) \geq c_H$, and thus along the path where the low quality seller always rejects his reservation price, she does not trade before h_1. 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 does not offer c_H, because $\mu(h_1) < \mu^*$ and thus, such offer would generate negative payoff regardless of the high quality seller's acceptance probability. Thus high quality does not trade. Consequently, $\mu(h_1, A) = 0$, and thus by Lemma 3, $V_L(h_1, A) \leq v_L - c_L$. 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 consider $P_L(h) < c_H$. Then high type does not trade at h because buyer would make a loss offering c_H. This implies that, if (h, A) has positive probability 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. It is possible that $\mu(h,A) = \mu(h,R)$, in which case both these beliefs equal μ^*. Assume that $\mu(h,A) \neq \mu(h,R)$. Then, either $\mu(h,A) < \mu^* < \mu(h,R)$, or $\mu(h,A) > \mu^* > \mu(h,R)$. First, 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. Next, 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. Thus, $\mu(h,A) = \mu^* = \mu(h,R)$.

For item 3, we first note that since $P_H(h) = c_H$ and $\mu(h) > \mu^*$, the buyer's payoff is necessarily positive. Thus, the buyer never makes an offer that is rejected with probability 1. 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 $\tilde{Q}$. Further, since it is supposed that $P_L(h) = c_H$, and no offers exceeding c_H are made, $V_L(h) = \delta V_L(h,R)$. Then, (i) and (ii) 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 every period in the continuation play starting at (h,R), and trade takes place with probability 1. This implies that $V_L(h,R) = c_H - c_L = V_L(h,R,A)$. But this is a contradiction, because in that case, $P_L(h,R) \leq c_L$ while $\mu(h,R) = \mu^*$ and the buyers would have a profitable deviation to offering $P_L(h,R) + \varepsilon$ for small $\varepsilon > 0$. This establishes that $P_L(h) < c_H$ whenever (h,R) has positive probability while $\mu(h) > \mu^*$.

Since the buyer never makes offers that are rejected with probability 1 and $P_L(h) < c_H$, the low quality seller trades with positive probability. Suppose low quality seller trades with probability less than 1. Then, it must be that the buyer offers $P_L(h)$ with positive probability, which is rejected by the low quality seller with positive probability. Then, since the buyer is guaranteed a positive payoff, necessarily $P_L(h) < v_L$. But then the buyer has a profitable deviation to $P_L(h) + \varepsilon$ when $\varepsilon > 0$ is small, a contradiction. Thus, the low quality seller trades with probability 1 and therefore, $\mu(h, R) = 1$. 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. Now suppose that $\mu(h, A) < \mu^*$. 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 $h^t_\emptyset$ represent the t-length history that features no trading. It follows by Lemma 6 that for any t, if $(h^{t-1}_\emptyset, A)$ is on the equilibrium path, then $\mu(h^{t-1}_\emptyset, A) \in \{0, \mu^*\}$. Also, let $\gamma_s(h^t)$ be the probability with which the seller type $s \in \{L, H\}$ visits history h^t. Define $T_{\mu^*} = \{t | \mu(h^{t-1}_\emptyset, A) = \mu^*\}$ and $T_0 = \{t | \mu(h^{t-1}_\emptyset, A) = 0\}$. Then,

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

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

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

Further,

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

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 $\overline{Q}_H(h_\emptyset) \leq (v_L - c_L)/(c_H - c_L)$. 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

Assume that $\mu_0 \in (\mu^*, \mu^{**})$. Fix an equilibrium. Note that at any history h if $\mu(h) > \mu^*$, then (h, A) has positive probability. Because otherwise there exists sufficiently small $\varepsilon > 0$ such that a buyer arriving at this history has a profitable deviation to offering $c_H + \varepsilon$. Further, at each such history, the L seller trades with probability 1. This is because either trade takes place with probability 1, or if (h, R) is on the equilibrium path, by Lemma 6, $P_L(h) < v_L$, so that if $P_L(h)$ is rejected with positive probability when offered, the buyer would have a profitable deviation to offering $P_L(h) + \varepsilon$ for some sufficiently small $\varepsilon > 0$.

Let h_1^t represent the t-length history that features trading at each period. For any t, if h_1^t is on the equilibrium path, then so is $h_1^{t'}$ when $t' < t$. Further, if $\mu(h_1^t) > \mu^*$, by Lemma 6 $\mu(h_1^{t'}) > \mu^*$. Thus, there exists at most one T^* such that (i) $h_1^{T^*}$ is on the equilibrium path, (ii)$\mu(h_1^{T^*-1}) > \mu^*$ and (iii) $\mu(h_1^{T^*-1}, A) = \mu(h_1^{T^*}) = \mu^*$. Note that, conditional on having reached $h_1^{T^*-1}$, $(h_1^{T^*-1}, R)$ is on the equilibrium path, because otherwise belief cannot be updated. Further, by Lemma 6, $\mu(h_1^{T^*-1}, R) = 1$.

Consider two cases:

- If no such T^* exists, then for all t, h_1^t is on the equilibrium path and, thus, on the equilibrium path the L seller trades with probability 1 every period so that $\overline{Q}_L(h_\emptyset) = 1$. Further, by Lemma 6, along the equilibrium path the H seller fails to trade at most once. Thus, $\overline{Q}_H(h_\emptyset) \geq \delta$. Thus, the gains from trade in such an equilibrium is no less than

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

  which exceeds the bound stated in the proposition.

- If such T^* exists, the L seller reaches the history $h_1^{T^*}$ with probability 1, while the H seller reaches it with probability $\frac{1-\mu_0}{\mu_0} \frac{\mu^*}{1-\mu^*}$. Further, letting $\tilde{Q} \equiv \overline{Q}_L(h_1^{T^*-1}, A) = \overline{Q}_H(h_1^{T^*-1}, A)$ represent the continuation expected amount of trade if the belief reaches μ^* for the first time at time T^*, then the following must hold:

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

  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 $P_L(h^{t-1}) \leq v_L$. Defining $Q^* = (v_L - c_L)/(c_H - c_L)$, this is equivalent to

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

Then,

$$\begin{aligned}
\overline{Q}_L(h_\emptyset) &= (1-\delta^{T^*}) + \delta^{T^*}\tilde{Q} \geq (1-\delta) + \delta\tilde{Q} \geq 1 - (1-\delta)Q^* \\
\overline{Q}_H(h_\emptyset) &\geq [(1-\delta^{T^*}) + \delta^{T^*}\tilde{Q}]\frac{1-\mu_0}{\mu_0}\frac{\mu^*}{1-\mu^*} + \left(1 - \frac{1-\mu_0}{\mu_0}\frac{\mu^*}{1-\mu^*}\right)\delta \\
&\geq [1-(1-\delta)Q^*]\frac{1-\mu_0}{\mu_0}\frac{\mu^*}{1-\mu^*} + \left(1 - \frac{1-\mu_0}{\mu_0}\frac{\mu^*}{1-\mu^*}\right)\delta
\end{aligned}$$

Therefore, using the fact that $\mu^*/(1-\mu^*) = (c_H - v_L)/(v_H - c_H)$, the expected surplus is no less than

$$\underbrace{(1-\mu_0)(1-(1-\delta)Q^*)(v_L - c_L) + (1-\mu_0)(1-(1-\delta)Q^*)(c_H - v_L)}_{\equiv B}$$

$$+$$

$$\underbrace{\mu_0\delta\left(1 - \frac{1-\mu_0}{\mu_0}\frac{\mu^*}{1-\mu^*}\right)(v_H - c_H)}_{>0}.$$

Simple algebra yields the claimed expression in the proposition.

Next we show that the lower bound on the gains from trade strictly exceeds $(1-\mu_0)(v_L - c_L)$. 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

$$\frac{1}{2-\delta} \geq \delta \Leftrightarrow \delta^2 - 2\delta + 1 \geq 0 \Leftrightarrow (1-\delta)^2 \geq 0,$$

which holds.

B.4 Proof of Proposition 7

Fix an equilibrium. Take h such that $\mu(h) > \mu^{**}$. We claim that trade takes place with probability 1. If not, then (h, R) has positive probability. Then, 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 $P_L(h) \geq c_L$. But then, since $\mu(h) > \mu^{**}$, the buyer strictly prefers to trade with both types at price c_H rather than trading with only the L seller at price $P_L(h)$, and thus trade must take place with probability 1. Thus, at such histories, trade takes place with probability 1 at price c_H, and the belief is never updated, establishing the claim.

B.5 Equilibrium construction: transparent market with intra-period monopsony

As discussed in the main text, we construct equilibria using dynamic programming techniques in the spirit of Abreu 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.

We start by replicating the definitions and characterizing equations discussed in Section 5.3 of the main text. Our construction specifies parameters $(\alpha, \beta, U^A, U^R, P_L, U)$ where α is the probability with which the buyer offers c_H as opposed to the low quality seller's reservation price P_L, β is the probability with which the low quality seller accepts his reservation price when offered, U^A and U^R are the continuation payoffs of the low quality seller upon trade or no trade, respectively, while U is the equilibrium payoff of the low quality seller. The equilibrium conditions are

- Low quality seller's reservation price relate to U^A, U^R as follows:

$$(1-\delta)(P_L - c_L) = \delta(U^R - U^A). \tag{6}$$

- The optimality of the buyer strategy requires that

$$\alpha = \begin{cases} 1 & \text{if } \tilde{\mu}(v_H - c_H) + (1-\tilde{\mu})(v_L - c_H) > (1-\tilde{\mu})(v_L - P_L)\beta \\ \in [0,1] & \text{if } \tilde{\mu}(v_H - c_H) + (1-\tilde{\mu})(v_L - c_H) = (1-\tilde{\mu})(v_L - P_L)\beta \\ 0 & \text{if } \tilde{\mu}(v_H - c_H) + (1-\tilde{\mu})(v_L - c_H) < (1-\tilde{\mu})(v_L - P_L)\beta \end{cases}. \tag{7}$$

- The L-seller's payoff can be calculated along the possibly off-equilibrium path where she rejects her reservation price when offered. Thus,

$$U = \alpha[(1-\delta)(c_H - c_L) + \delta U^A] + (1-\alpha)\delta U^R. \tag{8}$$

Given these observations, we define enforceability as follows:

Definition: For each $\mu \in [0,1]$ let $\mathcal{U}_\mu \in \mathbb{R}_+$ be a set of potential payoffs for the low quality seller. 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 (6), (7) and (8) and

$$\mu^A = \frac{\mu\alpha}{\mu\alpha + (1-\mu)(\alpha + (1-\alpha)\beta)} \text{ and } \mu^R = \frac{\mu(1-\alpha)}{\mu(1-\alpha) + (1-\mu)(1-\alpha)(1-\beta)} \tag{9}$$

whenever the denominators are positive.

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

Then, Proposition 8 states that self-generating correspondences define equilibrium payoffs.

B.5.1 Constructing a self-generating set

Next we construct a self-generating correspondence in the above sense. This construction also describes the strategies and beliefs associated with these equilibria.

Overview: We demonstrate that for each $\mu \leq \mu^*$, $\mathcal{U}_\mu$ is an interval. We specify these intervals in our formal result Proposition 9. For $\mu > \mu^{**}$, we construct a unique equilibrium in which the L seller receives a payoff of $c_H - c_L$. When $\mu \in (\mu^*, \mu^{**})$, the construction of the sets $\mathcal{U}_\mu$ turns out to be quite intricate. Specifically, the interval (μ^*, μ^{**}) is partitioned into countably infinite sub-intervals, with cutoffs $\mu^0 = \mu^* < \mu^1 < \mu^2 < \cdots < \mu^{**}$ with $\lim_{i\to\infty} \mu^i = \mu^{**}$. For each i and $\mu, \mu' \in (\mu^i, \mu^{i+1})$, $\mathcal{U}_\mu = \mathcal{U}_{\mu'} = \{U^i\}$ for some properly chosen U^i while for the cutoff points μ^i, $\mathcal{U}_{\mu^i} = [U^{i-1}, U^i]$.

In our construction, we first describe the sequences $\mu^0, \mu^1, \cdots$ and $U^0, U^1, \cdots$ and specify the intervals that correspond to $\mathcal{U}_\mu$ for $\mu \leq \mu^*$. Then, in the proof of Proposition 9, we verify that these sets altogether are self-generating.

Construction: We first recursively define sequences μ^i, U^i described above that allow the characterization of $\mathcal{U}_\mu$ for $\mu \in (\mu^*, \mu^{**})$. Next, we demonstrate that $U \in [U^i, U^{i+1}]$ is enforceable when $\mu = \mu^i$ while U^i is enforceable when $\mu \in [\mu^i, \mu^{i+1}]$.

We first set $\mu^0 = \mu^*$ and $U^0 = c_H - c_L - (1-\delta)(v_L - c_L)$.[28] Then we define P^1 by

$$(1-\delta)(P^1 - c_L) = \delta\left[(c_H - c_L) - U^0\right].$$

Thus, P^1 is the smallest reservation price of the L-seller if $\mu^A = \mu^0$ so that $U^A \leq U^0$ and $U^R = c_H - c_L$. (Recall that for these beliefs any on path failure of trade leads to $\mu = 1$, justifying the specification of U^R.) Next, for $\mu > \mu^0$ we define $P(\mu)$ to be the reservation price of the L seller so that at belief μ the buyer is indifferent between targeting

[28] Below we demonstrate that $U^0 = c_H - c_L - (1-\delta)(v_L - c_L)$ is the upper bound of payoffs that are enforceable when $\mu = \mu^*$.

the L-seller alone at price $P(\mu)$ and targeting both types of the seller at price c_H.

$$\frac{\mu}{1-\mu} = \frac{c_H - P(\mu)}{v_H - c_H}.$$

Note that $P(\mu)$ is decreasing in μ. We define μ^1 by $P^1 = P(\mu^1)$. That is, μ^1 is the largest belief starting from which the above construction works, so that an equilibrium can be constructed with $\mu^A = \mu^0$ and $\mu^R = 1$. Next, for each $\mu \in (\mu_0, \mu^1)$ we define $\alpha(\mu)$ to be the probability with which c_H must be offered so that the updated belief upon trade satisfies $\mu^A = \mu^0$:[29]

$$\frac{\mu}{1-\mu}\alpha(\mu) = \frac{\mu^0}{1-\mu^0},$$

and we set $\alpha^1 = \alpha(\mu^1)$. We note that for any $\mu \in [\mu^0, \mu^1]$, the above construction leads to a payoff of U^0 for the L seller. This is observed by substituting the expressions for $P(\mu)$, $\alpha(\mu)$ into

$$U^0 = (1-\delta)\alpha(\mu)(c_H - P(\mu)) + \delta(c_H - c_L).$$

Next, starting at belief exactly μ^1, if P^1 is the reservation price of the L seller, the buyer is willing to randomize between offering c_H and offering P^1. Further, as long as the probability with which c_H is offered varies between α^1 and 1, we have $\mu^A \in [\mu^0, \mu^1]$, and therefore it is feasible to have $U^A = U^0$, justifying the reservation price P^1. Note that the L-seller's payoff with this construction varies with α and is given by

$$U^{0-1}(\alpha) \equiv (1-\delta)\alpha(c_H - P^1) + \delta(c_H - c_L).$$

We let $U^1 = U^{0-1}(1)$. That is, U^1 is the largest payoff that this construction delivers to the L seller, starting at belief μ^1.

Next, we can iterate this construction starting with (μ^1, U^1) instead of (μ^0, U^0) to obtain (μ^2, U^2). Iterating further allows us to recursively define $\mu^i, U^i, \alpha^i, P^i$ which can

[29]In these equilibria, H seller trades if and only if c_H is offered and the L seller trades with probability 1.

be simplified as follows: $\mu^0 = \mu^*$ and

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

Note that as $i \to \infty$, we have $P^i \to c_L$, $U^i \to c_H - c_L$ and $\mu^i \to \mu^{**}$. Thus, for each $\mu \in (\mu^*, \mu^{**})$, there exists i such that $\mu \in (\mu^{i-1}, \mu^i]$.

Now we are ready to define the sets of payoffs for each belief that we will subsequently show form a self-generating correspondence. The remaining details of the above construction as well as verification that the sets specified for other beliefs below are self-generating are presented in the proof of our formal result Proposition 9.

$$\mathcal{U}_\mu = \begin{cases} [0, v_L - c_L] & \text{if } \mu = 0 \\ [(1-\delta)(v_L - c_L), v_L - c_L] & \text{if } \mu \in (0, \mu^*) \\ [(1-\delta)(v_L - c_L), c_H - c_L - (1-\delta)(v_L - c_L)] & \text{if } \mu = \mu^* \\ [U^{i-1}, U^i] & \text{if } \mu = \mu^i \text{ for some } i \\ \{U^{i-1}\} & \text{if } \mu \in (\mu^{i-1}, \mu^i) \text{ for some } i \\ \{c_H - c_L\} & \text{if } \mu > \mu^{**} \end{cases} \tag{17}$$

Proposition 9 $\{\mathcal{U}_\mu\}_{\mu\in[0,1]}$ *where for each* μ, $\mathcal{U}_\mu$ *is as defined in* (17)*, is self-generating.*

Proof. We show that $\{\mathcal{U}_\mu\}_{\mu\in[0,1]}$ defined in (17) satisfies the conditions listed above.

Case 1: $\mu = 0$ Take $U \in \mathcal{U}_0 \equiv [0, v_L - c_L]$. Let $P_L = U + c_L$, $\alpha = 0$, $\beta = 1$, $\mu^A = 0$, and $U^A = U, U^R = U/\delta$. These choices satisfy (6) and (8). Since trade takes place with probability 1 ($\alpha = 0, \beta = 1, \mu = 0$), $\mu^A = 0$ satisfies (9) and (9) does not restrict μ^R. Buyer's offer of P_L with probability 1 also trivially satisfies (7). Finally, $U^A = U \in \mathcal{U}_{\mu^A}$ since $\mu^A = 0$. It remains to specify μ^R such that $U^R = U/\delta \in \mathcal{U}_{\mu^R}$.

- If $U \geq \delta(1-\delta)(v_L - c_L)$, let $\mu^R = \mu^*$. To establish $U^R = U/\delta \in \mathcal{U}_{\mu^*}$, it suffices to observe that the lowest value of U/δ is no less than $\inf \mathcal{U}_{\mu^*} = (1-\delta)(v_L - c_L)$ and the highest value is no more than $\sup \mathcal{U}_{\mu^*} = c_H - c_L - (1-\delta)(v_L - c_L)$. The former requirement is trivially satisfied because $\delta(1-\delta)(v_L - c_L)/\delta = (1-\delta)(v_L - c_L)$. The latter is equivalent to $(v_L - c_L)/\delta \leq c_H - c_L - (1-\delta)(v_L - c_L)$, which can be re-expressed as
$$\frac{v_L - c_L}{c_H - c_L}\left(\frac{1}{\delta} + 1 - \delta\right) \leq 1.$$
Since $\frac{v_L - c_L}{c_H - c_L} < \delta^2$ by Assumption 1, a sufficient condition is
$$\delta^2\left(\frac{1}{\delta} + 1 - \delta\right) \leq 1 \Leftrightarrow \delta^2(1-\delta) \leq 1 - \delta,$$
which holds.

- If $U < \delta(1-\delta)(v_L - c_L)$, let $\mu^R = 0$. For such U, $U/\delta = U^R \in [0, (1-\delta)(v_L - c_L)) \subset [0, v_L - c_L] \equiv \mathcal{U}_0$. This completes the argument.

Case 2: $\mu \in (\mathbf{0}, \mu^*)$ Take $U \in [(1-\delta)(v_L - c_L), v_L - c_L]$. Let $\alpha = 0$, $\beta = 1 - \frac{\mu^*}{1-\mu^*}\frac{1-\mu}{\mu}$, $\mu^A = 0$, $\mu^R = \mu^*$, $P_L = v_L$. Also choose $U^R = U/\delta$ and $U^A = U/\delta - \frac{1-\delta}{\delta}(v_L - c_L)$. By choice of these variable, (6),(7), (8) and (9) are satisfied. It remains to show that $U^R = U/\delta \in \mathcal{U}_{\mu^*}$ and $U^A = U/\delta - \frac{1-\delta}{\delta}(v_L - c_L) \in \mathcal{U}_0$. Note that both U^R and U^A are increasing in U. Then, for $U^R \in \mathcal{U}_{\mu^*}$ it suffices to show that

$$\frac{(1-\delta)(v_L - c_L)}{\delta} \geq (1-\delta)(v_L - c_L) \text{ and } \frac{v_L - c_L}{\delta} \leq c_H - c_L - (1-\delta)(v_L - c_L).$$

The first inequality holds by inspection. The second follows from Assumption 1 and is shown in the previous part of this proof.

Similarly, for $U^A \in \mathcal{U}_0$ it suffices to show that

$$\frac{(1-\delta)(v_L - c_L)}{\delta} - \frac{1-\delta}{\delta}(v_L - c_L) \geq 0 \text{ and } \frac{v_L - c_L}{\delta} - \frac{1-\delta}{\delta}(v_L - c_L) \leq v_L - c_L,$$

both of which hold with equality.

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

- First consider $\mathbf{U} \in [(\mathbf{1}-\delta)(\mathbf{v_L}-\mathbf{c_L}), (\mathbf{1}-\delta^{\mathbf{2}})(\mathbf{v_L}-\mathbf{c_L})]$. This is the lower end of the interval U_{μ^*}. The payoffs in this interval are enforced by belief punishments. Specifically, let $\alpha = 0$, $\beta = 0$, $\mu^A = 0$, $\mu^R = \mu^*$, $P_L = v_L$, $U^A = [U - (1-\delta)(v_L - c_L)]/\delta, U^R = U/\delta$. Since $\alpha = \beta = 0$, (9) does not restrict μ^A while $\mu^R = \mu^*$ satisfies (9). Further, by choice of these parameters, (6), (7) and (8) are satisfied. Also, for any U in this interval, $U^A = [U - (1-\delta)(v_L - c_L)]/\delta \in [0, (1-\delta)(v_L - c_L)] \subset \mathcal{U}_0$. It remains to show that $U^R \in \mathcal{U}_{\mu^*}$. That is,

  $$(1-\delta)(v_L - c_L) \le U/\delta \le c_H - c_L - (1-\delta)(v_L - c_L).$$

  The first inequality trivially follows because $U \ge (1-\delta)(v_L - c_L)$. The latter is equivalent to

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

  also holds by Assumption 1.[30]

- Next consider $\mathbf{U} \in [(\mathbf{1}-\delta^{\mathbf{2}})(\mathbf{v_L}-\mathbf{c_L}), \mathbf{c_H}-\mathbf{c_L}-(\mathbf{1}-\delta^{\mathbf{2}})(\mathbf{v_L}-\mathbf{c_L})]$. This is the middle component of the interval $\mathcal{U}_{\mu^*}$. In this construction, instead of specifying all parameters that enforce an arbitrary U in this interval, we specify a subset of parameters and then show that by varying the remaining parameters one can enforce all payoffs in this interval.

  Let $\mu^A = \mu^R = \mu^*$, $\beta = 0$ and $P_L = v_L$. 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 (6) and $U^A, U^R \in \mathcal{U}_{\mu^*}$. It is also readily verified that these parameters satisfy (7). Further, by varying α over $[0,1]$ and U^A over the specified

[30] Note that this is a weaker requirement than $\frac{v_L - c_L}{\delta} \le c_H - c_L - (1-\delta)(v_L - c_L)$, which we have verified above.

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 U enforced to be 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 payoff 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 $\mathbf{U \in [c_H - c_L - (1 - \delta^2)(v_L - c_L), c_H - c_L - (1 - \delta)(v_L - c_L)]}$. This is the upper end of the interval $\mathcal{U}_{\mu^*}$. The payoffs in this component are enforced using belief rewards.

  Specifically, let $\alpha = 1$, $\beta = 0$, $\mu^A = \mu^*$, $P_L = v_L$,

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

  Since $\alpha = 1$, (9) does not restrict μ^R while $\mu^A = \mu^*$ satisfies (9). Further, by choice of these parameters, (6), (7) and (8) are satisfied. Also, for any U in this interval.

  $$U^A = \frac{U - (1-\delta)(c_H - c_L)}{\delta} \in \left[c_H - c_L - \frac{1-\delta^2}{\delta}(v_L - c_L), c_H - c_L - \frac{1-\delta}{\delta}(v_L - c_L)\right] \subset \mathcal{U}_{\mu^*}.$$

  The last set inclusion follows because $c_H - c_L - \frac{1-\delta}{\delta}(v_L - c_L) < c_H - c_L - (1-\delta)(v_L - c_L)$ and $c_H - c_L - \frac{1-\delta^2}{\delta}(v_L - c_L) > (1-\delta)(v_L - c_L)$. The first inequality is apparent by observation. The second inequality is equivalent to $1 > [1 - \delta + (1-\delta^2)/\delta]\,(v_L - c_L)/(c_H - c_L)$. Since by Assumption 1, $(v_L - c_L)/(c_H - c_L) < \delta^2$, a sufficient condition is $1 > [1 - \delta + (1-\delta^2)/\delta]\,\delta^2$, which can be re-arranged as $(1-\delta)(1-$

$\delta^2) > -\delta^3$, which holds.

It remains to specify μ^R such that $U^R \in \mathcal{U}_{\mu^R}$. By construction,

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

Thus, there exists i such that $U^R \in (U^{i-1}, U^i] \subset \mathcal{U}_{\mu^i}$. It suffices to choose $\mu^R = \mu^i$.

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

- Fix $i > 0$ and consider $\mu \in (\mu^{i-1}, \mu^i)$.

Choose $\mu^R = 1$, $\mu^A = \mu^{i-1}$, $U^R = c_H - c_L$, $\beta = 1$. Also choose P_L, α, U^A as follows:

$$\begin{aligned}
\frac{\mu}{1-\mu} &= \frac{c_H - P_L}{v_H - c_H} \\
\frac{\mu}{1-\mu}\alpha &= \frac{\mu^{i-1}}{1-\mu^{i-1}} \equiv \frac{c_H - P^{i-1}}{v_H - c_H}. \\
(1-\delta)(P_L - c_L) &= \delta(c_H - c_L - U^A).
\end{aligned}$$

The first equality is the buyer's indifference condition and uniquely defines P_L and ensures that (7) is satisfied. The second uniquely defines α and ensures that (9) is satisfied. The third uniquely defines U^A and ensures that (6) is satisfied. The equivalence is due to the definition of μ_{i-1} and P^{i-1}. Using (6), the right-hand-side of (8) can be re-expressed as $(1-\delta)\alpha(c_H - P_L) + \delta(c_H - c_L)$. Substituting for α and μ from the first two equalities above yields $\delta(c_H - c_L) + (1-\delta)(c_H - P^{i-1}) = U^{i-1}$, 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 $P_L \in (P^i, P^{i-1})$. Since for any i, $(1-\delta)(P^i - c_L) = \delta(c_H - c_L - U^{i-1})$, we have

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

implying that $U^A \in \mathcal{U}_{\mu_{i-1}} = [U^{i-2}, U^{i-1}]$. This shows that U^{i-1} 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 $\mu^R = 1$, $U^A = U^{i-1}$, $U^R = c_H - c_L$, $P_L = P^i$, $\beta = 1$. By these choices (6) and (7) are satisfied. Also choose α, μ^A to satisfy

$$\delta(c_H - c_L) + (1-\delta)\alpha(c_H - P^i) = U, \text{ and } \frac{\mu^A}{1-\mu^A} = \frac{\mu_i}{1-\mu_i}\alpha.$$

The first equality uniquely pins down α and guarantees that (8) holds. The second equality guarantees that (9) holds. Further, 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 $\alpha \in (\alpha_i, 1)$. Thus, $\mu^A \in (\mu_{i-1}, \mu_i)$ and by construction, $U^{i-1} \in \mathcal{U}_{\mu^A}$. This shows that U is enforceable. ■

B.6 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 8 and 9 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 Proposition 7, 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 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

ABREU, D., D. PEARCE AND E. STACHETTI, "Toward a Theory of Discounted Repeated Games with Imperfect Monitoring," *Econometrica* (1990), 1041–1063.

ATHEY, S. AND K. BAGWELL, "Optimal Collusion with Private Information,," *RAND Journal of Economics* (2001), 428–465.

———, "Collusion with Persistent Cost Shocks,," *Econometrica* (2008), 493–540.

CHO, I. K. AND D. KREPS, "Signaling Games and Stable Equilibria," *Quarterly Journal of Economics* 56 (1987), 601–612.

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.

FREIXAS, X., R. GUESNERIE AND J. TIROLE, "Planning under incomplete information and the ratchet effect,," *Review of Economic Studies* (1985), 173–191.

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, "Government interventions in a dynamic market with adverse selection," *Journal of Economic Theory* (2015), 371–406.

———, "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. AND L. MAESTRI, "Dynamic contracting with limited commitment and the ratchet effect,," *Theoretical Economics* (2020), 583–623.

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

———, "Price Transparency and Market Screening," *working paper* (2022b).

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

LAFFONT, J.-J. AND J. TIROLE, "Comparative statics of the optimal dynamic incentive contract,," *European Economic Review* (1987), 901–926.

———, "The dynamics of incentive contracts,," *Econometrica* (1988), 1153–1175.

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* (2023).

PEI, H., “Reputation Building Under Observational Learning,” *Review of Economic Studies* (2023).

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).

WISEMAN, T., “A partial folk theorem for games with unknown payoff distributions,,” *Econometrica* (2005), 629 – 645.

Online Appendix: Partial pooling equilibria in transparent markets 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, these equilibria feature a positive amount of trade which is less than 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.

We construct trading paths that cycle through several periods of trade with single-period pauses.[31] 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.

[31] This construction is similar to the one-step separation equilibria constructed in Kaya and Roy (2022a). 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.

- 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 $\tau < k$, $\theta = L, H$,

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

  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 $Q_k < 1$.

  <u>For $\tau = k$</u>:

  $$\begin{aligned}(1-\delta)(P_L(\tau) - c_L) + \delta(v_L - c_L) &= Q_k(P_k - c_L)\\ (1-\delta)(P_H(\tau) - c_H) &= Q_k(P_k - c_H).\end{aligned}$$

  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.

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.[32]

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

[32]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.

- $P_H(\tau) = c_H$ 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, $\mu(h) = \mu^*$.

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.

B.6.1 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.

C Constructing self-generating sets of payoffs: A special case:

Our definition of self-generation requires that we specify payoff sets $\mathcal{U}_\mu$ for all possible μ. One may wonder if it is possible to construct some sets of enforceable payoffs for a subset of beliefs in isolation. For instance, it is trivial to see that at belief $\mu = 1$, $U = c_H - c_L$ is enforceable with respect to the set $\mathcal{U}_1 = \{c_H - c_L\}$ by choosing $\mu^A = \mu^R = 1$ and $U^A = U^R = c_H - c_L$.[33] Similarly, at belief $\mu = 0$, $U = 0$ is enforceable with respect to $\mathcal{U}_0 = \{0\}$. A more interesting case is when $\mu = \mu^*$. By Lemma 6, starting from μ^*, on the equilibrium path, belief is never updated. Thus, it is natural to wonder if a subset of enforceable payoffs at belief μ^* can be characterized in isolation. Here, we demonstrate that this is possible for some parameter values but not others. This exercise, in addition to clarifying our method of construction, also highlights some challenges we encounter.

Claim: If $\delta(c_H - v_L) > v_L - c_L$, then any $U \in [v_L - c_L, c_H - v_L]$ is enforceable with respect to $\mathcal{U}_{\mu^*} \equiv [v_L - c_L, c_H - v_L]$ at belief μ^*.

Proof of claim: Since on the equilibrium path belief is never going to be updated, it is necessary that $\beta = 0$ so that both types of the seller trade if and only if c_H is offered. Consider U that can be enforced by choosing $P_L = v_L$ and $\mu^A = \mu^R = \mu^*$ together with some U^A, and $U^R = U^A + (v_L - c_L)\frac{1-\delta}{\delta}$. These choices satisfy (6) and (9). Further, given that $\mu = \mu^*, P_L = v_L$ and $\beta = 0$, any $\alpha \in [0, 1]$ satisfies (7). For U^A, U^R as specified to be in $\mathcal{U}_{\mu^*}$, it is necessary that $U^R \in \left[\frac{v_L - c_L}{\delta}, c_H - v_L\right]$, which is non-empty since by assumption $\delta(c_H - v_L) > v_L - c_L$. Then by (8), any U satisfying the following for some $\alpha \in [0, 1]$ and $U^R \in \left[\frac{v_L - c_L}{\delta}, c_H - v_L\right]$ can be enforced:

$$U = \alpha(1 - \delta)(c_H - v_L) + \delta U^R.$$

[33] Here, we are abusing terminology since our formal notion of enforceability is required to specify $\mathcal{U}_\mu$ for all possible μ.

Thus, $v_L - c_L$ is enforced by choosing $\alpha = 0$ and $U^R = (v_L - c_L)/\delta$, while $c_H - v_L$ is enforced by choosing $\alpha = 1$ and $U^R = (c_H - v_L)$. Since α and U^R can vary continuously over their respective ranges, all $\mathcal{U}_{\mu^*} = [v_L - c_L, c_H - v_L]$ is self-generating regardless of how $\mathcal{U}_\mu$ are satisfied for other μ.

If $\delta(c_H - v_L) > v_L - c_L$ does not hold, the above construction fails. Typically, it becomes impossible to characterize self-generating sets of payoffs at belief μ^* without characterizing $\mathcal{U}_\mu$ for other μ since off-path punishments and off-path rewards become necessary. Below we construct $\mathcal{U}_\mu$ for each μ without imposing any restrictions on the parameters.