Market Screening with Limited Records

Ayça Kaya*† and Santanu Roy‡

May 30, 2020

Abstract

Markets differ in the availability of past trading records of their participants. In a repeated sale model with adverse selection, we study the impact of the availability of such records on trading outcomes. We consider regimes varying with respect to the length of the available records. We characterize a class of equilibria in which the record length has direct welfare implications via the market's need to re-screen the seller, as well as indirect implications via the low quality seller's incentives to mimic the high quality seller. As the record length increases, the market needs to re-screen less frequently, which improves efficiency. In turn, less frequent screening makes mimicking more attractive and limits the market's ability to learn. These considerations lead to a non-monotonic relationship between record length and overall gains from trade.

Keywords: Repeated sales, adverse selection, lemons market, transparency, limited records.

JEL codes: D82, C73 , D61

*Corresponding author.

†Department of Economics, University of Miami, Coral Gables, FL 33146. E-mail: a.kaya@miami.edu.

‡Department of Economics, Southern Methodist University, Dallas, Texas 75275. E-mail: sroy@smu.edu

1 Introduction

Market participants form inferences about the quality of a seller's product through various channels. When the seller is a repeated participant, much can be inferred from his past trading behavior. For sellers that offer limited quantities for sale and trade through private transaction, trading records are often limited. Nevertheless, there is wide variation across markets in the availability of credible public records of past trading at the level of individual sellers.[1] Thus it is important to understand how transparency in this sense impacts trading outcomes and individual market participants' ability to capture surplus, both in order to understand individual incentives to disclose and to evaluate the desirability of policy interventions.

We study how the availability of past trading records impact the market's ability to screen a seller of a product with unknown quality, and welfare implications thereof. Specifically, we study the trading outcomes when a long-lived seller who can produce one unit of output in each period sequentially meets multiple buyers over an infinite horizon. The fully persistent (and binary) quality of the seller's output is his private information and the buyers make inferences about it by observing only a fixed-length history of the seller's past trades.[2] We characterize a class of equilibria and analyze how the length of the available records affect trading outcomes. Within this class of equilibria, we find that, contrary to what conventional wisdom would suggest, shorter record lengths in fact allow for finer screening of the seller's quality, but they require more frequent pauses of trade to re-screen. Due to this trade-off, the gains from trade is non-monotone in the length of publicly available records.

Our analysis sheds light on the impact of record lengths on the payoffs of market participants as well as total gains, and as such on both private and social incentives to provide long public records of past trading outcomes. Our results indicate that these incentives can differ widely across markets. One should therefore expect significant variation in the length and availability of such records across markets. Further, public policy interventions to facilitate or censure disclosure of past trading outcomes need to be nuanced and take

[1] For instance, when viewing rental properties on Airbnb or VRBO, the customers do not have access to the records of past bookings of the property. In contrast, when viewing profiles of service professionals on Thumbtack, a platform that connects customers to individual service providers in various fields, customers have access to the number of times each professional has been hired as well as for how long they have made their services available. In more organized labor markets, the full employment history of individuals is typically available to potential employers.

[2] We focus on the case where the trading prices are not observable, which is a natural assumption in markets where transactions are private. Besides, observable trading prices circumvent the limits imposed by record length by carrying information about the beliefs of previous periods' buyers. We comment on the impact of this on equilibrium outcomes in Section 6.

into account specific aspects of the economic environments. For instance, we show that long public records are more desirable in markets with high turnover where sellers are more likely to exit and in markets where buyers are unwilling to pay high quality premium.

Our main results are based on a class of equilibria, which we term one-step-separation (OSS) equilibria. These equilibria feature trading at only two different prices along its path at different histories: one lower price at which only the low quality goods trade, and one higher price at which a pool of high and low quality goods trade. This is a natural equilibrium structure when the seller type is binary. Further, an equilibrium from this class exists regardless of record length, initial market optimism and seller patience, which is a convenient feature that facilitates comparisons.

On the path of an OSS equilibrium, at certain histories, the market targets only the low type seller. Pausing of trade at such histories improves the market's assessment of the seller's quality, and allows him to trade at higher prices until further screening becomes necessary; i.e. when the latest trading pause disappears from the observable history.[3] Thus, the equilibrium path of a high quality seller (and with some probability, the low quality seller) cycles through single-period pauses of trade, followed by several periods of trading at relatively high prices. With positive probability, the low quality seller follows a path along which he trades every period at a price equal to the full value of the low quality good to the buyers.

The latter aspect of the equilibrium path implies that the low quality seller's equilibrium payoff is independent of record length. However, the record length affects the low quality seller's incentives to mimic the high quality seller as it determines the frequency of costly pauses in the latter's path. When the record length is short, the pauses are frequent, and thus mimicking is unattractive. Consequently, the market can achieve full screening of seller types. After screening, buyer competition drives price to the valuation of high quality. In the range where full screening remains possible, increasing the record length unambiguously benefits the high quality seller as the pauses for re-screening becomes less frequent and he can earn his full information surplus for a longer stretch of time.

As the record length becomes somewhat longer, high quality seller's trading frequency increases but the trading prices must come down from their highest level to render mimicking less attractive. This requires that buyers' expectation of the quality is lower, implying that some amount of pooling is necessary, and full screening is not feasible. Further, even though the increased frequency of trading and reduced price counter each other, we show that the high quality seller is now worse off with a longer record. Thus, the high quality seller's payoff is non-monotone in record length. The total surplus, i.e. the total gains

[3] As becomes clear, re-screening may be necessary earlier than that if the low type has sufficiently strong incentives to mimic the high type.

from trade moves in parallel with the high type seller's equilibrium payoff. In particular, increasing the record length improves welfare when records are short to begin with, but reduces welfare when they are already long.

The non-monotonicity of gains from trade in record length raises the question of optimal record lengths. In general, the optimal record length is reached when the records become just long enough that the low quality seller's incentives to mimic precludes full separation. This occurs at longer record lengths, and may even approach infinity, if the low quality seller's incentives to mimic are not particularly strong, for instance if the seller is not particularly patient or the value difference between the two product qualities (i.e. the quality premium) is not particularly large.

Though we derive our results based on equilibria involving one step separation, finer screening can be possible in other equilibria with multi-step screening. In Section 5, within a limited setting, we demonstrate that the existence of equilibria with finer screening does not alter the comparison of surplus across records of different lengths.

The rest of the paper is organized as follows. Section 1.1 discusses the related literature. Section 2 presents the model, Section 3 characterizes equilibria, Section 4 discusses welfare implications of record lengths, Section 5 considers equilibria other than OSS, and Section 6 concludes. All proofs and some technical details are collected in the Appendix.

1.1 Related literature

Our model features a dynamic lemons market in which sellers with persistent quality repeatedly participate. The main concern is the market's ability to learn about a persistent state, specifically to screen the seller for his quality. The variable whose impact we consider is information about the past trading activity of the seller, specifically the length of record about the seller's trading behavior.

A number of papers consider similar settings. Mauring (2017) and Mauring (forthcoming) study models where the market participants learn about the prior distribution of seller qualities and prior distribution of seller's reservation prices, respectively, when observing limited information about the outcomes of the last trading period only. In both papers, sellers are non-strategic. Consequently, the link between market efficiency or trading dynamics and observability of past outcomes is purely informational and thus is distinct from ours. Kovbasyuk and Spagnolo (2018), similar to us, studies the impact of record length on the efficiency of trading outcomes. Their setting is different from ours in two important aspects: in their case, the records include information that is directly about the seller's quality, based on, for instance, reviews left by previous users; and the seller's type evolves according to a Markov process. These two aspects create an experimentation motive of sales which drives their results. Thus, the economic mechanisms studied in the two papers are distinct.

Another literature considers the impact of information about past seller behavior on market outcomes when the seller has a single durable item to sell. For instance, Horner and Vieille (2009), and Fuchs et al. (2016) study the efficiency implications of observability of past (rejected) price offers, while Kim (2017) studies the implication of observability of time on the market. None of these papers consider the impact of limited records. Further, in single-sale models the impact of other forms of transparency is confined to the market's ability to achieve initial screening as maintenance of reputation is irrelevant. Thus, the trade-offs we study in this paper do not arise in that literature.

A separate literature considers the impact of erasing past records in environments with moral hazard, where reputation considerations discipline the market participant's behavior. In this literature Vercammen (1995) and Elul and Gottardi (2015) study credit markets with possibility of strategic default, incorporating both moral hazard and adverse selection. They demonstrate that erasing records of past defaults may improve incentives to repay. Relatedly, Liu and Skrzypacz (2014) study a model of reputation building in a trust game when a finite fixed-length record of past behavior of the long-run player is observable by his (sequence of) short-run partners. In their analysis, all equilibrium behavior can be described as a function of the time since the last incidence of exploitation, and generates cyclical paths of play. Though this structure is similar to our OSS equilibrium, their cycles are based on the incentive of the long-run player to exploit trust, while our cycles are based on the market's need to re-screen the seller.

More recently, Bhaskar and Thomas (forthcoming) demonstrate how, in the absence of adverse selection, limited but exact records lead to breakdown of cooperative behavior while providing coarse information about record lengths can be desirable. More broadly, optimal design of ratings, (i.e. public signals of past behavior or outcomes) have been taken up in contexts of reputation building (Ekmekci (2011)), career concerns (Horner and Lambert (forthcoming)) and adverse selection (Hopenhayn and Saeedi (2019)).

2 Model

A long-lived seller can produce one unit of output every period. Time is discrete and horizon is infinite. The calendar time varies over $T = \{\cdots, -1, 0, 1, \cdots\}$. As we clarify below, the assumption that the calendar time starts at $-\infty$ allows us to avoid tedium, and all our results extend if time starts at $t = 0$ at the cost of additional notation and arguments. In every period $t \in T$, the seller meets two or more buyers, each with unit demand, who submit simultaneous take-it-or-leave-it price offers to the seller. Seller either chooses one of the buyers to trade with or rejects all. Regardless, buyers leave the game, and the seller moves to the next period meeting with another set of buyers.

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

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

The instantaneous payoff for any party who does not trade is normalized to 0. The seller maximizes the expected discounted sum of all his future payoffs. Her discount factor is $\delta \in [0, 1]$. Seller's type is his private information. All players hold a common prior that assigns probability μ_0 to type $s = H$. The prior belief satisfies the static lemons condition:

$$\mu_0 v_H + (1 - \mu_0) v_L < c_H.$$

The observable history consists of whether seller has made a trade or not in the preceding k periods and not other details of transactions or information about behavior in earlier periods. Then, a k-length public history is a k-length vector of zeros and ones, a "one" indicating a trade and a "zero" indicating no trade. Our analysis is valid for all record lengths including the limit $k = \infty$. Let $H^k = \{0, 1\}^k$ be the set of all such histories.[4] Then a pure strategy for a buyer is a map from H^k to $\mathbf{R}_+$ representing his price offer. The seller's private history includes all of his past transactions, the price offers he has accepted and rejected and any current outstanding offer as well as the calendar time.

A pure strategy of the seller maps the set of these histories as well as his type into a decision to accept / reject the highest current offer. A belief system maps each public history into a probability that the buyers assign to the high type seller. Let $\mu(h^k)$ represent this probability at $h^k \in H^k$. A strategy profile and a belief system is a perfect Bayesian equilibrium if all buyers form beliefs based on observable histories and strategies of others using Bayes rule whenever possible, and all players maximize their payoff based on their beliefs and the strategies of others.

We characterize a class of perfect Bayesian equilibria which we refer to as one-step separation (OSS) equilibria. In these equilibria the buyers offer a unique price $P \geq c_H$ provided that the length of the most recent streak of trades are below a certain threshold, and otherwise offer v_L, targeting only the low type seller. On the path of play, the high

[4]Our assumption that the calendar time starts at $-\infty$ allows us to avoid describing shorter public histories. As becomes immediately clear in what follows, all equilibria we consider in this paper can easily be adapted to accommodate such histories.

type trades only at price P, while the low type either trades at price v_L every period, or mimics the high type's strategy.

To formalize the definition of OSS equilibria for each public history h^k, let $\tau(h^k)$ be the number of periods since the seller last failed to make a trade. In other words, $\tau(h^k)$ is the length of the latest streak of trades made by the seller (represented as ones in the history vector), capped at the observability limit k. We adapt the convention that for any history h^k of any length which features no rejection, $\tau(h^k) = \infty$. Let $0 \leq m \leq k$ be an integer, $v_H \geq P \geq c_H$ and $\alpha \in [0, 1)$ scalars. Next, we define a class of buyer strategies which we refer to as (m, P, α) offer strategies.

Definition 1 *An (m, P, α) offer strategy of a buyer is described as follows:*

- *If $\tau(h^k) > m$, offer v_L.*
- *If $\tau(h^k) = m$, randomize between offering v_L and P, assigning probability α to the latter.*
- *If $\tau(h^k) < m$, offer P.*

When buyers use an (m, P, α) offer strategy, the market attempts to screen the seller by making a low offer of v_L whenever the observed volume of trade is sufficiently high. A rejection of this price is rewarded by m or $m + 1$ periods of high price offers, after which the market re-screens the seller.

We note that since neither trading prices nor rejected price offers are observable, in all perfect Bayesian equilibria of this game, the sellers' strategies can be described relative to a type- and history-dependent reservation price: at any history, if the highest price offer strictly exceeds the reservation price, the seller trades, and if the highest offer is strictly below the reservation price, the buyer fails to trade. Further, these reservation prices are uniquely pinned down by buyer offer strategies. Then, the seller's best response to any offer strategy of the buyers can be identified by his history-dependent probability of accepting his reservation price if offered. In general, two extreme pure strategies of the seller are one where he always rejects his reservation price and one where he always accepts it. In an OSS equilibrium, seller randomizes over these two pure strategies.[5] Formally,

Definition 2 *An OSS equilibrium is defined by the following two restrictions:*

1. *buyers use an (m, P, α) strategy;*

[5]We conjecture that all best responses to an (m, P, α) offer strategy lead to equivalent outcomes. Restricting attention to this simple class of seller strategies allows us to describe outcomes and belief updating in the simplest possible way.

2. *the high type seller always accepts his reservation price, while the low type seller, with probability $\beta \in (0,1]$ adapts the pure strategy of always rejecting his reservation price, and with probability $1-\beta$ adapts the pure strategy of always accepting it.*

Not surprisingly, and as will be formally demonstrated below, in an OSS equilibrium, the high type seller trades if and only if the offer is P. For future reference it is convenient to introduce notation $Q_{m,\alpha}$ for the high quality seller's frequency of trade in such an equilibrium where buyers use (m, P, α) strategies. This frequency can be recursively defined as follows:

$$Q_{m,\alpha} = (1-\delta)(\delta + \cdots + \delta^m) + \delta^{m+1}[\alpha(1-\delta+\delta Q_{m,\alpha}) + (1-\alpha)Q_{m,\alpha}],$$

which yields

$$Q_{m,\alpha} = \frac{\delta + \cdots + \alpha\delta^{m+1}}{1+\delta+\cdots+\alpha\delta^{m+1}} \tag{1}$$

Naturally, $Q_{m,\alpha}$ is increasing in m, which is the consecutive number of times the high price P is offered—and thus the high type trades—between pauses.

The name OSS equilibrium refers to the fact that in these equilibria, the market screens the seller for one-period at a time, and thus on-path beliefs have only two values. As noted above, the path of an OSS equilibrium follows cycles of single periods of screening by the market, followed–in the case where the screening leads to higher beliefs–by the reward of high prices for the next several periods. Since we exogenously restrict the screening to take place in one period, the record length does not directly limit the market's ability to screen. However, since necessarily $m < k$, the record length constrains the amount of rewards the seller can receive in return for pausing the trade during screening. Thus, the focus on this class allows us to isolate the impact of constraints on "reputation maintenance" which is the novel feature brought on in a repeated sale environment relative to the better-understood single sale environment.

3 Characterization of OSS equilibria

This section characterizes the set of OSS equilibria. The results are presented in two lemmas leading to the main proposition characterizing equilibria. First, in Lemma 1, we address the seller's best responses to a (m, P, α) strategy of buyers. Next, in Lemma 2, we tackle the optimality of buyer offer strategies. Finally, Proposition 1 puts together these ingredients and characterizes the OSS equilibria. All formal proofs are relegated to Appendix A.

Lemma 1 (Reservation prices) *In an OSS equilibrium, the seller's best responses can be described with respect to a reservation price which depends only on τ. Fix the buyer strategy (m, P, α) and let $P_\theta(\tau)$ represent the reservation price of the type θ seller, $\theta \in \{L, H\}$. Then,*

- *for any τ,*
$$P_H(\tau) \begin{cases} < P & \text{if } P > c_H \\ = P & \text{if } P = c_H \end{cases}.$$

- *$P_L(\tau) \geq v_L$ for any τ, if and only if*
$$v_L - c_L \leq Q_{m,\alpha}(P - c_L). \tag{2}$$
If (2) does not hold, $P_L(\infty) < v_L$.

Further, for any τ, $P_H(\tau) > P_L(\tau)$, and $P_H(\tau) \geq c_H$.

Lemma 1 establishes a number of useful and intuitive properties of the seller's equilibrium behavior. First, it establishes that the best response of a high type seller to a (m, P, α) strategy is to always accept P and always reject v_L. Next, it establishes that the low type seller is less picky than the high type ($P_L(\tau) < P_H(\tau)$) and thus will also always accept P when offered. This is a familiar property and is generally satisfied in single-sale models but not guaranteed in a repeated sale setting.

Lemma 1 also establishes that, the low type's willingness to accept of a price v_L is related to the frequency $Q_{m,\alpha}$ of the offer P. If the anticipated frequency $Q_{m,\alpha}$ is sufficiently low, the low type's unique best response is to always accept the offer of v_L, and thus he never trades at P. In contrast, when $Q_{m,\alpha}$ is sufficiently high, low type's unique best response is to always reject v_L. Satisfaction of (2) with equality guarantees that $P_L(\infty) = v_L$, and at any other history $P_L(\tau) > v_L$. In particular, when (2) is satisfied with equality, then it is a best response of the low type to always reject the offer of v_L and always accept P, putting him on the same path as the high type. At the same time, when this holds, it is also a best response of the low type to accept v_L when $\tau = \infty$, and thereafter, follow a path along which he never receives an offer of P.

Next, we turn to the buyer strategies. Lemma 1 establishes that if $P > c_H$, at histories with $\tau \leq m$, the seller's reservation prices are strictly less than P. Then, for an (m, P, α) strategy with $P > c_H$ to be the buyers' equilibrium strategy, it is necessary that

$$\hat{P}(\mu) \equiv \mu v_H + (1 - \mu) v_L \geq P, \tag{3}$$

which in turn implies that the expected quality strictly exceeds the reservation prices of both types. The next proposition formalizes the intuitive conclusion that in this case, the unique equilibrium of the buyers' bidding game at a history with $\tau \leq m$ features an offer of $\hat{P}(\mu)$ with probability 1.

Lemma 2 (Bidding equilibrium) *Consider a history at which the buyers assign probability μ to the high type seller. Let P_θ represent the type θ seller's reservation price. If $\hat{P}(\mu) > P_H > P_L$, then the high price offer is equal to $\hat{P}(\mu)$ with probability 1.*

Lemmas 1 and 2 together imply that when $\tau \in \{0, \cdots, m\}$, the belief must satisfy $\hat{P}(\mu) = P$. When $m < k$, there are equilibrium path histories, for instance those with $\tau = m$, where the buyers know that they are facing the "pooling types," and thus their belief μ satisfies $\hat{P}(\mu) = P$, yet they are prescribed to offer P with probability less than 1. Then, once, again by Lemmas 1 and 2, this is only possible if $P = c_H$. Equivalently, (m, P, α) strategies with $P > c_H$ can be part of an equilibrium if and only if $m = k$. The next proposition makes use of this observation and characterizes the OSS equilibria. The characterization splits the economic environments into three groups based on the strength of the mimicking incentives when the high type is trading at the highest feasible frequency $Q_{k,0}$. Recall that $Q_{k,0}$ is the (average) frequency of trading if trade takes place with probability 1 whenever the record features at least one period of a pause. As such, for fixed δ, $Q_{k,0}$ varies from $\delta/(1+\delta)$ when $k = 1$ to δ when $k = \infty$.

Proposition 1 *There always exists an OSS equilibrium where buyers use an (m, P, α) strategy. Further, in any such equilibrium*

- *(weak incentives to mimic) If*

$$Q_{k,0}(v_H - c_L) < (v_L - c_L),$$

  then, $P = v_H$, $m = k$ and $\alpha = 0$.

- *(intermediate incentives to mimic) If*

$$Q_{k,0}(c_H - c_L) < v_L - c_L < Q_{k,0}(v_H - c_L),$$

  then $\alpha = 0$, and P satisfies

$$\frac{v_L - c_L}{P - c_L} = Q_{k,0}.$$

- *(strong incentives to mimic) If*

$$v_L - c_L < Q_{k,0}(c_H - c_L), ,$$

  then necessarily $P = c_H$ and all OSS equilibria are payoff equivalent to one where

$$Q_{m,\alpha} = \frac{v_L - c_L}{c_H - c_L}.$$

Proposition 1 partitions the economic environments into three sets. This partitioning highlights the role of the record length. When the record length is very short so that the reputation must be replenished often by a pause of trade, the low type does not find it profitable to mimic the high type, and thus, full separation becomes possible, and is indeed the unique outcome of an OSS equilibrium (weak incentives to mimic case).

As the record length, and thus $Q_{k,0}$ grows, the environment features intermediate or strong incentives to mimic. In this case, full separation is not possible, because it would imply that the buyers offer v_H at a frequency $Q_{k,0}$, which would make it very attractive for the low type to mimic a high type. Intermediate incentives to mimic case is when offering c_H at the maximum frequency $Q_{k,0}$ is not sufficiently attractive for the low type. Thus, if the buyers would follow a (m, P, α) strategy with $P = c_H$, the low type seller would always choose to accept v_L. But in that case the posited buyer strategy cannot be part of an equilibrium, since whenever the history features a rejection, the buyers would be compelled to offer v_H. That is why, in the intermediate incentives to mimic case all OSS equilibria feature $P \in (c_H, v_H)$.

Finally, in the strong incentives to mimic case, the trade frequency of the high type must be less than $Q_{k,0}$ because a higher frequency of trade even at the lowest price $P = c_H$ makes mimicking too attractive for the low type. Since trade cannot be slowed down when $P > c_H$ (Lemma 2), all OSS equilibria in this case feature $P = c_H$, and the frequency $Q_{m,\alpha}$ is pinned down by (2).

It immediately follows from the discussion of Section 3 that the equilibrium payoff of the low type seller is fixed at $v_L - c_L$ regardless of the parameter values. Moreover, the buyers always receive payoffs of 0. A direct corollary of Proposition 1 formally states these observations and describes the equilibrium payoff of the high type seller.

Corollary 1 (Equilibrium payoffs and surplus) *In any OSS equilibrium, the low type seller's payoff is $v_L - c_L$ while all buyers earn zero profits. The high type's equilibrium payoff varies with the environment as follows:*

- *In the case of weak incentives to mimic, high type seller's payoff is*

$$Q_{k,0}(v_H - c_H)$$

- *In the case of intermediate incentives to mimic, high type seller's payoff is*

$$v_L - c_L - Q_{k,0}(c_H - c_L)$$

- *In the case of intermediate incentives to mimic, high type seller's payoff is 0.*

Longer records move the environment from weak to intermediate to strong incentives to mimic cases. Since $Q_{k,0}$ increases in k, it is immediately observed that the high type seller benefits from longer records if the record length is short enough to begin with (weak incentives to mimic). But the high type's payoff declines in k within the case of intermediate incentives to mimic, and remains constant once the record length becomes long enough that the strong incentives to mimic case is reached. Since low type's and all buyers' payoffs are fixed across environments, this result demonstrates that the total gains from trade is also non-monotonic in the record length.

4 Trade observability and gains from trade

Panel (a) of Figure 1 represents the expected gains from trade conditional on each type of seller as well as the total expected gains from trade. It is not surprising that when the incentives to mimic are weak, longer record lengths are surplus-improving, as they reduce the frequency of the need to pause trade to re-screen the seller while full separation remains possible. In particular, as the figure demonstrates, with longer record lengths, the high type seller is able to trade more frequently, and thus generates more surplus, while the low type's trading frequency remains constant.

Within the intermediate incentives to mimic case, longer records have opposing impacts on the gains from trade generated by the low type and the high type. The high type's trading volume increases for the same reason as in the weak incentives to mimic case while the trading price must decline. Because of buyer competition, lower prices can be sustained only when the low type seller pools with the high type's path of less frequent trading. Longer record lengths require lower prices, and thus a bigger probability of pooling, further reducing the low type's overall trading volume. When considering the impact of longer records on the overall gains from trade, the lower trading volume of the low type swamps the higher trading volume of the high type seller, leading to lower surplus as k grows. The reason can be understood by noting that regardless of k, the low type's payoff from following the "pooling path" must be exactly $v_L - c_L$, his payoff from revealing himself. As k and therefore, frequency of trade along the pooling path, grows, the pooling outcome moves along a single indifference curve of the low type over quantity-price pairs, towards higher quantities and lower prices. Since the high type has a relatively stronger preference for higher prices/lower quantities, this move lowers the high type's payoff, and thus the overall surplus.

Finally, the environment switches to the strong incentives to mimic case once the record length becomes long enough that the price must go below c_H to keep the low type indifferent while trading at the highest possible frequency $Q_{k,0}$. Since the price cannot go below c_H,

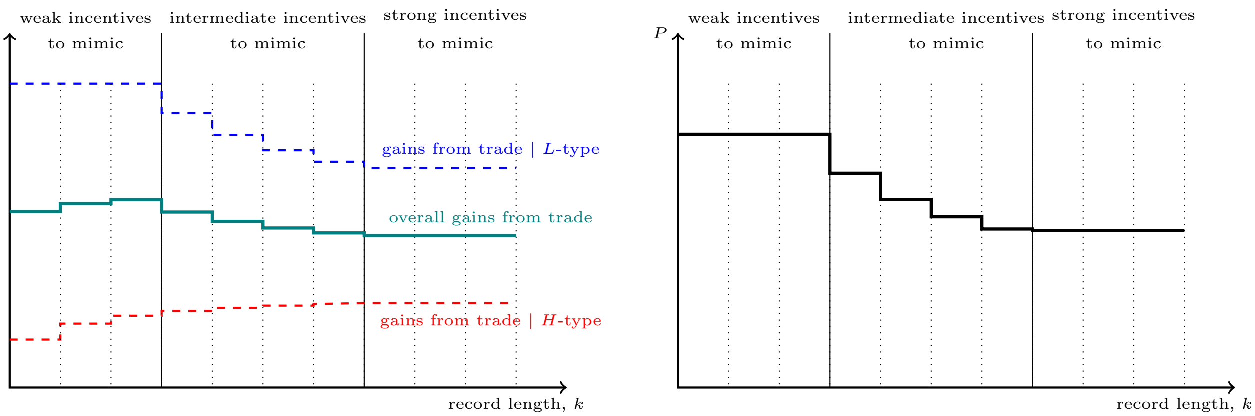

(a) Gains from trade (rescaled).

(b) High type's trading price (rescaled).

Figure 1: Impact of record length for parameter values $\delta = .9, v_H = 1.4, v_L = 1, c_H = 1.2, c_L = 0, \mu_0 = .5$.

now the frequency of trade cannot increase beyond $(v_L - c_L)/(c_H - c_L)$, and thus remains constant as the record length grows beyond this level. Buyer competition does not preclude slowing down trade in this case, as discussed above.

Panel (b) of Figure 1 represents the high type's trading prices as a function of the record length k. It is interesting to compare the evolution of gains from trade generated by the high type to the evolution of the prices he receives. Within the weak incentives to mimic case, the price remains constant at v_H as k increases and thus, the increase in the gains from trade due to the increase in record length is fully captured by the high type. Outside the weak incentives to mimic case, even though the trading volume of and therefore the trading surplus generated by the high type continues to increase, much of this additional surplus is allocated to the cross-subsidization of the low type's trades. This causes the prices to decline, and reduces the high type seller's payoff.

A second corollary of Proposition 1, by way of Corollary 1 reveals that the record length reaches its optimal level just around when the environment switches from the weak to intermediate incentives to mimic cases.

Corollary 2 (Optimal record length) *Let $\bar{k}$ be the longest record length for which the environment is one of weak incentives to mimic. That is,*

$$\bar{k} = \sup\left\{k \mid \frac{v_L - c_L}{v_H - c_L} \geq Q_{k,0}\right\}.$$

Then gains from trade is maximized either when the record length is $\bar{k}$ or when it is $\bar{k} + 1$.

In particular, when

$$\frac{v_L - c_L}{v_H - c_L} < \frac{\delta}{1+\delta},$$

$\bar{k} = 0$, *i.e. longer record lengths always reduce welfare, while when*

$$\frac{v_L - c_L}{v_H - c_L} > \delta,$$

$\bar{k} = \infty$, *i.e. longer records are always welfare improving.*

The mechanism that drives the desirability/undesirability of longer records in our model generates novel considerations when applied to the question of optimal levels of transparency. For instance, it is natural to conjecture that if the quality difference between the products of the two types of the seller is not large, transparency in terms of longer records is less likely to be beneficial, as there is little to learn. Our model presents a mechanism that goes against this reasoning. Specifically, Corollary 2 implies that when the quality premium v_H/v_L is small, the optimal record length is longer. By now, it is clear that this is due to the fact that when the quality premium is large, the gains from being known as high quality are large, and thus it is harder for the market to screen the seller, and this issue is exacerbated by longer records.

A similar conclusion can be derived when considering market turnover. By standard reasoning, it is possible to interpret δ as the survival rate of a seller in the market. Thus, larger δ indicates lower rates of turnover and thus lower volatility. Again, a natural conjecture would be that longer records are likelier to enhance outcomes in a market with slow turnover as in such markets "maintenance of reputation" is more valuable. Our model highlights a mechanism which, for this exact reason, in fact, leads to the conclusion that the optimal record length is longer when the turnover is faster.[6] Indeed, the increased survival rate or patience of the seller makes the ability to maintain reputations more valuable. This implies that, when records are long, building up the market's belief for high quality becomes disproportionately more valuable. This adversely affects the market's ability to screen the seller, and consequently reduces the optimal record length.

5 Discussion: other equilibria

Our main prediction that depending on the economic environment, longer records may increase or decrease the total gains from trade have been derived within the context of OSS equilibria. In this section we demonstrate that this prediction is valid when other

[6]This follows from Corollary 2 by noting that $Q_{k,0}$ is increasing in δ.

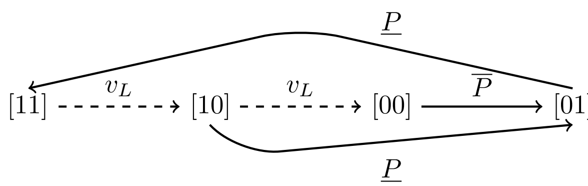

(a) high type seller and a portion of low type sellers

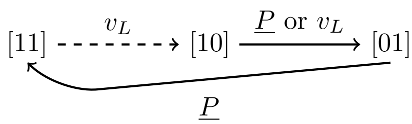

(b) a portion of low type sellers

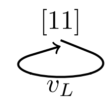

(c) a portion of low type sellers

Figure 2: Trading cycles in the 2SS equilibrium.

equilibria are considered. In particular, we focus on the comparison of two record lengths $k = 1$ and $k = 2$, while allowing a different equilibrium, which we refer to as the two-step separation (TSS) equilibrium, in the latter case.[7] The TSS equilibrium differs from the OSS in that it allows screening to continue for more than one period. For that reason, when considering a fixed record length, it is possible to achieve finer screening of the seller under a TSS equilibrium. This contributes to higher surplus, as it implies that a smaller portion of the low types follow an infrequent trading path to mimic the high type. However, this screening is achieved at the cost of longer pauses. Below when comparing $k = 1$ to $k = 2$, we demonstrate that in all cases where the finer screening of TSS improves the welfare relative to that of OSS with $k = 2$, the OSS equilibrium with $k = 1$ does even better, because it allows for even finer screening at similar frequency of pauses.

We first informally describe the path of play in the TSS equilibrium, and then justify the above assertions. We defer the formal construction and the derivation of payoffs to Appendix B.

[7]The existence of this equilibrium requires further restrictions on the environment, including an upper bound on the initial belief and a lower bound on δ. We defer the formal construction of the equilibrium and discussion of sufficient conditions for existence to the appendix. Further, we conjecture that OSS equilibrium is the unique equilibrium in the case of $k = 1$, and for the case of $k = 2$, TSS is the only possibility other than OSS.

TSS equilibrium Assume that $k = 2$. In this case, public histories are $\{[00], [01], [10], [11]\}$, where each 0 represents a period of no trade and each 1 represents a period of trade. In the TSS equilibrium, the trading follows a course similar to that of OSS equilibrium, but the seller types split into three instead of two cyclical paths, and consequently trade takes place at 3 different prices instead of 2. Figure 2 represents these cycles. In all panels, the solid curves represent trade while the dashed curves represent rejections. The price offers that lead to these outcomes are marked along the curves. In any such equilibrium, $\overline{P} > \underline{P} > c_H$.

Panel (a) of Figure 2 shows the path that the high type seller, possibly together with a portion of the low types follow in a TSS equilibrium. Along this path, screening is achieved in one or two steps: at the history [00] buyers offer v_L for sure, and if it is rejected (i.e. at [10]), they still offer v_L with positive probability. Thus, the history [00] occurs with positive probability. Panel (b) illustrates a path followed by the low type seller with some probability. Along this path, the low type seller mimics the high type seller up to a single-period pause, then reveals himself. That is, he rejects v_L at [11] but accepts it at [10]. Finally, with some probability, the low type seller always trades at price v_L, as demonstrated in panel (c) of Figure 2.

Notice that the high type visits all histories with positive probability. The belief that the buyers hold at each history is related to the relative frequency with which the two types visit that history. As clear by the ongoing discussion, the ratio of high type's frequency to low type's is highest at history [00] and lowest at [11], justifying the ranking of the price offers at these histories. The exact values of prices $\overline{P}, \underline{P}$ are pinned down by the seller's indifference conditions at various histories. In particular, the low type seller's reservation price must be no larger than v_L at history [10], as he may accept v_L at this history (if he is following path (b)), and must be equal if he follows path (a) with positive probability, as along that path low type rejects this offer. Focusing on the case of $\delta = 1$, this restriction amounts to

$$\overline{P} - v_L \leq v_L - c_L, \tag{4}$$

with equality if $2v_L - c_L > c_H$, yielding $\overline{P} = \min\{v_H, 2v_L - c_L\}$. At history [10], the high type's reservation price is necessarily $\underline{P}$.[8] For the case of $\delta = 1$ this restriction is equivalent to

$$\frac{\underline{P} - c_H}{\overline{P} - c_H} = \frac{3}{5}. \tag{5}$$

[8]Note that at this history buyers randomize over two different prices, v_L and $\underline{P} > c_H$. Such equilibrium of the bidding game is possible only if the high type seller's reservation price at that history is exactly $\underline{P}$. If $\underline{P}$ strictly exceeds the high type's reservation price, then by Lemma 2, the unique equilibrium is for the buyers to offer $\underline{P}$ with probability 1.

Combining (4) and (5) pins down the $\underline{P}$ as follows:

$$\underline{P} = \frac{2}{5}c_H + \frac{3}{5}\min\{v_H, 2v_L - c_L\}.$$

Payoffs In the TSS equilibrium, just as in the OSS equilibria, the low type seller's payoff is $v_L - c_L$ while all buyers' payoffs are 0. Thus, as before, the total surplus moves parallel with the changes in the high type's payoff. To calculate the high type's payoff, we note that at history [10] he is indifferent between accepting and rejecting the higher offer $\underline{P}$, and thus his equilibrium payoff can be calculated along the off-equilibrium path where he rejects both offers v_L and $\underline{P}$ at history [10].[9] Along this path, the high type trades at $\overline{P}$ a quarter of the time, at $\underline{P}$ a quarter of the time and fails to trade the remaining half of the time. Thus his equilibrium payoff when $\delta = 1$ is given by

$$\frac{\overline{P} + \underline{P}}{4} = \frac{2}{5}\min\{v_H - c_H\ ,\ 2v_L - c_L - c_H\}. \tag{6}$$

Limits of screening and impact of record length: Based on the ongoing discussion we observe the following parallels to the OSS equilibria. First, just as in OSS equilibria, the highest price that can be achieved is bound by the low type's incentives to mimic, this time given by (4). Second, the high type's payoff decreases as the low type's incentives to mimic becomes stronger, e.g. as his payoff $v_L - c_L$ from revealing his type becomes smaller.

Next, in comparing the equilibrium payoff under OSS vs. under TSS equilibria, we note that when $2/3(c_H - c_L) > v_L - c_L > 1/2(c_H - c_L)$, the high type receives a zero payoff under OSS but a positive payoff under TSS. Thus the ability to extend the screening stage may improve the high type's payoff, and thus the total surplus by allowing the high type to trade at higher prices at lower frequency.

Importantly, however, for the range where TSS generates more surplus than OSS with $k = 2$ (and in fact for the whole range where TSS generates positive payoff for the high type seller, i.e. when $v_L - c_L > (c_H - c_L)/2$), the OSS equilibrium with $k = 1$ does even better. Specifically, the high type seller's payoff in this case (for $\delta = 1$) is

$$\frac{1}{2}\min\{v_H - c_H\ ,\ 2v_L - c_L - c_H\},$$

which strictly exceeds the payoff under TSS when $k = 2$ and $\delta = 1$, given in (6). Thus, we conclude that the OSS equilibrium of $k = 1$ generates more surplus than both equilibria of $k = 2$ if and only if it generates more surplus than OSS equilibrium of $k = 2$. This implies that, even allowing for TSS equilibrium, $k = 1$ generates more surplus than $k = 2$

[9] On panel (a) of Figure 2, this path is represented by removing the lowest solid arrow.

if $(v_H - c_L)/(v_L - c_L)$ is sufficiently small, and the opposite ranking holds if this ratio is large.

The analysis of this section demonstrates that the main mechanism driving our conclusions when focusing only on OSS equilibria continue to be present when other equilibria are considered and thus our conclusions are likely to extend to more general settings. We have presented the results of this section for a limited subset of situations where our main results apply, largely because for longer record lengths the analysis quickly becomes intractable, as suggested by the already tedious analysis of this case. Nevertheless, it is reasonable to conjecture that similar conclusions are likely to hold for larger k when considering analogous k-step separation equilibria.

6 Conclusion

Our analysis highlights a particular channel via which availability of trading records may influence the market's ability to screen a seller's quality. In order to isolate the economic forces at play, we have made several modeling choices, some of which can be relaxed without qualitative changes to our conclusions.

Throughout, we assume that the seller's product quality is fully persistent. This renders the model stark, but at the same time, allows for a particularly parsimonious analysis. If the seller type followed a sufficiently persistent Markov process, it is likely that equilibria similar to the OSS equilibria will continue to exist with possibly non-constant prices during the "reward periods."

Our assumption that multiple sellers arrive each period well-represents a market that is more competitive on the demand side so that the seller captures most of the gains from trade. Analogous conclusions can be derived for a market where the demand side is thinner, so that the buyer captures more (or all) of the trading surplus. For instance, on the opposite extreme, where the buyers capture all surplus net of the low type seller's information rents, the high type's trading price remains at c_H regardless of record length. The lack of competition on the buyer side eliminates the need for partial pooling in order to reduce prices, and thus the total gains from trade continues to increase with record length over the range that corresponds to the "intermediate incentives to mimic" case in our model, and then abruptly falls to the same level as in the competitive buyer model that we consider. It is interesting to note that, for parameter values that place the environment in the "intermediate incentives to mimic" case, total surplus is larger when the demand side is thin. Intuitively, this is because when buyers have the bargaining power, prices remain low, and the low quality seller's incentives to mimic, which generally precludes efficient trading, remain weaker.

In this paper, we focus on situations where the records include only the trading volumes but no information about trading prices. This is a common property of many markets, where the transactions are private and prices are not posted. Even when proposed prices for future transactions are posted, they are only noisy indicators of the realized prices in historical transactions. Nevertheless, it is interesting to note that, if our model is modified so that only the price at which the latest transaction took place is available, an equilibrium can be constructed that is payoff-equivalent to the case of "strong incentives to mimic" of our main model. The aspect that allows this is that the last period's transaction price carries information about the belief of the market at that history. This eliminates the market's need to pause trading to re-screen the seller, which in turn renders initial mimicking very attractive for the low type seller. Note that the gains from trade in this equilibrium is the lower bound of such gains when prices are not observable and as the record length varies.

Finally, the fact that our model is particularly parsimonious and its analysis is particularly clear cut renders this model suitable for embedding into larger contexts to study questions of interest such as the determinants of private incentives for disclosure, sellers' incentives to invest in cost reduction or buyers' incentives to invest in information acquisition.

References

Bhaskar, V. and C. Thomas, "Community Enforcement of Trust," *Review of Economic Studies* (forthcoming).

Ekmekci, M., "Sustainable reputations with rating systems," *Journal of Economic Theory* 146 (March 2011), 479–503.

Elul, R. and P. Gottardi, "Bankruptcy: Is It Enough to Forgive or Must We Also Forget?," *American Economic Journal: Microeconomics* 7 (2015), 294–338.

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

Hopenhayn, H. and M. Saeedi, "Optimal Quality Ratings and Market Outcomes," *Working Paper* (2019).

Horner, J. and N. Lambert, "Motivational Ratings," *Review of Economic Studies* (forthcoming).

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

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

Kovbasyuk, S. and G. Spagnolo, “Memory and Markets,” *Working Paper* (2018).

Liu, Q. and A. Skrzypacz, “Limited records and reputation bubbles,” *J. Econ. Theory* 151 (2014), 2–29.

Mauring, E., “Learning from Trades,” *The Economic Journal* 127 (05 2017), 827–872.

———, “Informational Cycles in Search Markets,” *American Economic Journal: Microeconomics* (forthcoming).

Vercammen, J. A., “Credit Bureau Policy and Sustainable Reputation Effects in Credit Markets,” *Economica* 62 (1995), 461–478.

Appendix

A Omitted proofs

Proof of Lemma 1. That the seller uses a reservation price strategy immediately follows because neither the trading prices nor the rejected offers are observed by the buyers and thus future offers are not affected by the trading prices. Further, the future offers only depend on τ, thus so do the continuation values and, in turn, the reservation prices. To characterize these reservation prices, let $V_\theta(\tau)$ be the continuation payoff of type θ seller at the beginning of a period that starts at history associated with τ. Also, let, with a slight abuse of notation, $V_\theta(\tau, \tilde{p})$ be the seller’s continuation payoff, after buyer offers are made, and the highest offer is $\tilde{p}$. Then, these value functions are derived from a best response of the seller if and only if[10]

$$V_\theta(\tau, \tilde{p}) = \max\{\delta V_\theta(0), (1-\delta)(\tilde{p} - c_\theta) + \delta V_\theta(\tau+1)\},$$

and

$$V_\theta(\tau) = \begin{cases} V_\theta(\tau, P) & \text{if } \tau < m \\ \alpha V_\theta(\tau, P) + (1-\alpha) V_\theta(\tau, v_L) & \text{if } \tau = m \\ V_\theta(\tau, v_L) & \text{if } \tau > m \end{cases}.$$

[10]Since the value function is necessarily bounded and $\delta < 1$, the Bellman equation has a unique solution.

On the other hand, given the posited reservation prices (i.e. candidate strategies), one can calculate the implied payoffs $V^*_\theta(\tau)$ of the seller of type θ in order to characterize a candidate solution. Then, it remains to show that $V^*_\theta(\tau)$ satisfies the above Bellman equation. First, the case where (2) does not hold is trivial. Second, for $P = c_H$, it trivially follows that the high type seller's reservation price is c_H regardless of history.

Next, assume that (2) holds and $P > c_H$. Then, the posited reservation prices imply that it is optimal for the seller, regardless of his type, to always reject v_L and always accept P. To characterize the resulting value functions V^*_θ, define $Q_{m,\alpha}$ by

$$Q_{m,\alpha} = \frac{\delta + \cdots + \delta^m + \alpha\delta^{m+1}}{1 + \delta + \cdots + \delta^m + \alpha\delta^{m+1}}.$$

Given buyer strategies, $Q_{m,\alpha}$ is the maximum frequency with which the seller can trade at the high price P. Then V^*_θ is recursively defined as follows:

$$V^*_\theta(\tau) = \begin{cases} (1-\delta)(P - c_\theta) + \delta V^*_\theta(\tau+1) & \text{if } \tau < m \\ \alpha\left[(1-\delta)(P - c_\theta) + \delta V^*_\theta(\tau+1)\right] + (1-\alpha)\delta V^*_\theta(0) & \text{if } \tau = m \\ \delta V^*_\theta(0) & \text{if } \tau > m \end{cases},$$

with $\delta V^*_\theta(0) = Q_{m,\alpha}(P - c_\theta) = V^*(m+1)$.

Note that at any history, by always rejecting the offer of v_L, the seller can guarantee a frequency of the offer P no less than $Q_{m,\alpha}$ but at no history can he guarantee a frequency larger than $Q_{m,\alpha}/\delta$. Let $Q(\tau)$ represent the frequency with which an offer of P is made in the continuation game if the seller only trades at price P, and following a history associated with τ. Note that $Q(0) = Q_{m,\alpha}/\delta$, $Q(\tau)$ decreases in τ over $\tau \leq m$, and $Q(m) \geq Q(m+1) = Q(\infty) = Q_{m,\alpha}$. In particular, for any τ, $Q(\infty) \leq Q(\tau) \leq Q(0)$. Note that, the low type can guarantee himself a payoff of $v_L - c_L$ by always accepting the offer of v_L. Based on this, the seller's reservation price $P_\theta(\tau)$ is calculated by

$$\delta(V^*_\theta(0) - V^*_\theta(\tau+1)) = (1-\delta)(P_\theta(\tau) - c_\theta)$$
$$\delta(Q(0) - Q(\tau+1))(P - c_\theta) = (1-\delta)(P_\theta(\tau) - c_\theta).$$

Showing that $V^*_\theta(\tau)$ satisfies the Bellman equation boils down to showing that $P_\theta(\tau)$ is no less than v_L and no larger than P. Since $Q(\tau)$ is decreasing, the reservation price $P_\theta(\tau)$ is non-decreasing. Specifically, since $P > c_\theta$,

$$\frac{P_\theta(0) - c_\theta}{P - c_\theta} < \frac{P_\theta(1) - c_\theta}{P - c_\theta} < \cdots < \frac{P_\theta(m) - c_\theta}{P - c_\theta} = \frac{P_\theta(m+1) - c_\theta}{P - c_\theta} = Q_{m,\alpha}. \quad (7)$$

Here, the last equality follows because $V^*_\theta(m+2) = V^*_\theta(m+1) = \delta V^*(0)$, and thus, $\delta(V^*_\theta(0) - V^*_\theta(m+2)) = \delta V^*(0)(1-\delta)$. That $P_L(\tau) \geq v_L$ follows by the last equality and (2).

That $P_H(\tau) < P$ follows from the last equality and the fact that $Q_{m,\alpha} < 1$. That $P_H(\tau) \geq c_H$ follows because $V_H^*(0) \leq V_H^*(\tau)$ for any τ, and because after acceptance the continuation payoff is given by $V_H^*(0)$.

To establish $P_H(\tau) > P_L(\tau)$ first note that at any such equilibrium $P_H(\tau) \geq c_H$. Then, if (2) does not hold, the claim is trivially satisfied. For the case where $P = c_H$, the claim is again trivially satisfied by (7). When (2) holds, the reservation prices of both types are given by

$$\frac{\delta}{1-\delta}(Q(0) - Q(\tau+1))(P - c_\theta) + c_\theta = P_\theta(\tau).$$

The claim follows because

$$\frac{\delta}{1-\delta}(Q(0) - Q(\tau+1)) \leq \frac{\delta}{1-\delta}(Q(0) - Q(m+1)) = \frac{\delta}{1-\delta}(1-\delta)Q(0) = \delta Q(0) < 1,$$

and thus the left-hand-side of the previous equality is increasing in c_θ. ■

Proof of Lemma 2. Assume that $\hat{P}(\mu) > P_H > P_L$. Let Π be the support of the random variable describing the highest offer.

- First we claim that $\sup \Pi = \hat{P}(\mu)$.

  Suppose not. Let $\overline{P} = \sup \Pi$. Then any buyer can guarantee itself a payoff of approximately $\hat{P}(\mu) - \overline{P} > 0$.

  - First suppose $\overline{P} \leq v_L$. If $P_L > \overline{P}$ or $\overline{P} = v_L$, then buyers' profit is 0, a contradiction. Next, if $P_L \leq \overline{P} < v_L$, then there exists a buyer who makes a profit no larger than $(1-\mu_0)(v_L - \overline{P})/2$. However, each buyer can guarantee a payoff of approximately $(1-\mu_0)(v_L - \overline{P})$.
  - Next, suppose $\overline{P} \in (v_L, P_H)$. Then, the buyers cannot make positive profits (and may make negative), a contradiction.
  - Finally consider $\overline{P} \in [P_H, \hat{P}(\mu))$. There must exist a buyer whose profit is no larger than $(\hat{P}(\mu) - \overline{P})/2$, a contradiction.

  This establishes the claim.

- Next, we argue that the maximum price offer is $\hat{P}(\mu)$ with probability 1. Let P_{max} be the random variable representing the highest offer made. Suppose that there exists $\varepsilon_1, \delta > 0$ such that

$$Prob(P_{max} < \hat{P}(\mu) - \varepsilon_1) > \delta.$$

Since $\overline{P} = \hat{P}(\mu)$, for any $\varepsilon_2 > 0$, there exists a buyer who makes an offer exceeding $\hat{P}(\mu) - \varepsilon_2$ with positive probability, and thus whose payoff is no larger than ε_2. However, by offering $\hat{P}(\mu) - \varepsilon_1$, each buyer can gurantee an expected payoff of $\delta\varepsilon_1$ by offering $\hat{P}(\mu) - \varepsilon_1$. This is a profitable deviation, when ε_2 is sufficiently small.

This establishes that for any ε, δ,

$$Prob(P_{max} < \hat{P}(\mu) - \varepsilon_1) < \delta.$$

The claim follows by taking ε_1, δ to 0.

■

Proof of Proposition 1. The three cases are comprehensive. Thus, the proof of existence is by arguing that the described strategies form an equilibrium in each case, which we establish below. Before taking each case individually, we establish the following preliminary claims.

Claim 1: (m, P, α) strategies with $P > c_H$ can be part of an OSS equilibrium only if $m = k$ and $\alpha = 0$.

Proof of Claim 1: Note that at any on path history featuring any rejection, the buyers' belief μ must satisfy

$$\mu = \frac{\mu_0}{\mu_0 + (1-\mu_0)\beta} > \mu_0.$$

Letting P_θ represent the θ-type seller's reservation price, by Lemma 1, $P \geq P_H > P_L$. Then, by Lemma 2, for there to be a bidding equilibrium at this history where buyers offer P with positive probability and no strictly higher offer is made, it is necessary that

$$(v_H - v_L)\frac{\mu_0}{\mu_0 + (1-\mu_0)\beta} + v_L = P. \tag{8}$$

Further, Lemma 1 and Lemma 2 together imply that if $P > c_H$ and (8) is satisfied, at any on path history that features a rejection, the unique bidding equilibrium is for the buyers to offer P (which is the expected quality) with probability 1. Thus, (m, P, α) strategies with $P > c_H$ can be part of an OSS equilibrium only if $m = k$ and $\alpha = 0$, establishing the claim.

Claim 2: In any OSS equilibrium, at a history featuring no rejections, the high type seller does not trade, and the price cannot exceed v_L.

Proof of Claim 2: First, since at all other on path histories the belief exceeds μ_0, and the beliefs form a martingale, the belief at this history must be less than μ_0 and thus the expected quality is below c_H. Since at any history, high type's reservation price is no less than c_H while the low type's reservation price is strictly less than the high type's (Lemma 1), the price offer cannot exceed v_L, and in turn, the high type does not trade, establishing the claim.

Claim 3: In any OSS equilibrium, when choosing the strategy of always accepting it, the low type's payoff is $v_L - c_L$. If (2) holds, when choosing the strategy of always rejecting his reservation price, the low type receives a payoff of $Q_{m,\alpha}(P - c_L)$.

Proof of Claim 3: First part follows because at any history with $\tau = \infty$, the low type's reservation price is no larger than v_L (Lemma 1), thus accepting an offer of v_L is a best response. Further, this acceptance keeps him at a history with $\tau = \infty$. Next, when (2) holds, at any other history, low type seller's reservation price is no smaller than v_L but strictly smaller than P (Lemma 1.) Thus always rejecting v_L and accepting P is a best response of the low type to a (m, P, α) strategy of the buyers, and delivers the claimed payoff.

We note that by Claim 3, $v_L - c_L \geq Q_{k,\alpha}(P - c_L)$, with equality if $\beta > 0$. Now consider the three cases separately:

- Weak incentives to mimic: Since $P \leq v_H$ and $Q_{m,\alpha} \leq Q_{k,0}$, in this case (2) does not hold. Thus, by Lemma 1 the low type's reservation price is strictly less than v_L, which implies that $\beta = 0$. Thus, at on-path histories involving a rejection, the belief of the buyers is $\mu = 1$. Then, by Lemma 2, at any such history the high offer is v_H with probability 1. Since the history with no rejections is on path, any off-path histories features at least one rejection. At such histories the belief is not pinned down by Bayes rule. All buyers offering v_H with probability 1 as the bidding equilibrium is supported by beliefs assigning probability 1 to the high type at these histories.

- Intermediate incentives to mimic: First note that in this case $P > c_H$, because otherwise, by Claim 3, and since $v_L - c_L > Q_{k,0}(c_H - c_L)$, it must be that $\beta = 0$, which in turn implies by Lemma 2 that at any history featuring rejection, the bidding equilibrium must result in an offer of v_H with probability 1, a contradiction. Then, by claim 1, $m = k$ and $\alpha = 0$. Further, to satisfy (8), $\beta \in (0, 1)$. Thus it must be that $v_L - c_L = Q_{k,\alpha}(P - c_L)$. Next, when β is as stated in the proposition, at on path histories featuring rejections, the expected quality is P. Thus, it is optimal for the buyers to offer P. Off-path histories all of which feature at least one rejection are handled analogously to the weak incentives to mimic case.

- Strong incentives to mimic: First, we establish that the suggested strategies form an equilibrium. The argument for the optimality of the seller strategies is identical to the intermediate incentives to mimic case. To establish the optimality of the buyer strategies, note that, at any history, the high type's reservation price is exactly c_H. On the path of equilibrium, if a history features a rejection, the belief is such that the expected quality is c_H and the low type seller's reservation price is no less than v_L. Thus, the buyers receive zero payoff from offering c_H. Moreover, even when the probability that the other buyers offer c_H is less than 1, a buyer cannot get positive payoff because any offer less than c_H is rejected by the high type. This establishes the optimality of the buyer strategies.

  The fact that necessarily $P = c_H$ follows by Claim 3, and noting that for any $P \geq c_H$, $Q_{k,0}(P - c_L) \geq Q_{k,0}(c_H - c_L) > v_L - c_L$. In turn, $P = c_H$ implies that $\beta \in (0, 1)$, and thus it must be that $Q_{m,\alpha}(c_H - c_L) = v_L - c_L$, as claimed.

■

B Appendix: Two step separation

To formally describe the TSS equilibrium, note that when $k = 2$, the set of full length public histories is $\{[11], [10], [01], [00]\}$ where 1 represents an instance of trade (acceptance) and 0 represents an instance of no trade (rejection). Let $\mu_{xy}, P^{\theta}_{xy}, \overline{P}_{xy}$, respectively, be the belief, reservation price of type-θ seller, and the maximum price offer at history $xy \in \{[11], [10], [01], [00]\}$. The following proposition describes the equilibrium behavior when $(v_H - c_L)/2 > v_L - c_L > (c_H - c_L)/2$.

Proposition 2 (TSS equilibrium, intermediate incentives to mimic) *Assume that $(v_H - c_L)/2 > v_L - c_L > (c_H - c_L)/2$. There exists $\overline{\mu}$ and $\overline{\delta}$ such that when $\mu_0 < \overline{\mu}$, $\delta > \overline{\delta}$, and $k = 2$ there exists an equilibrium in which*

- *the buyers*
  - *offer v_L at $[11]$;*
  - *offer $\hat{P}(\mu_{[01]})$ at $[01]$;*
  - *offer $\hat{P}(\mu_{[00]})$ at history $[00]$;*
  - *randomize so that the high offer is $\hat{P}(\mu_{[10]})$ or v_L, each with positive probability at history $[10]$;(let α stand for the probability of the former);*

- *the high type seller's reservation prices satisfy*

$$c_H \leq P^H_{[11]} = P^H_{[01]} < \hat{P}(\mu_{[01]}) = P^H_{[10]} = P^H_{[00]} < \hat{P}(\mu_{[00]}),$$

and he always accepts his reservation price when offered;

- *low type seller's reservation price at each history is v_L and he randomizes over 3 pure strategies:*
  - *always rejects his reservation price (probability $\gamma_1 \geq 0$)*
  - *rejects his reservation price at* [11] *and* [01] *but accepts his reservation price at* [10] *and* [00] *(probability γ_2)*
  - *always accepts his reservation price (probability $\gamma_3 = 1 - \gamma_1 - \gamma_2$).*

Proof: The proof is by construction. First, we characterize the equilibrium beliefs as functions of $\gamma_1, \gamma_2, \gamma_3$ and α. Fix a path of play along which the trading probabilities depend only on the history $xy \in \{[11], [10], [01], [00]\}$. Let q_{xy} be the frequency with which this path visits history xy. Also let α_{xy} be the probability with which trade takes place at history xy along this path. Then the following must hold:

$$\begin{aligned}
q_{[11]} &= \alpha_{[11]} q_{[11]} + \alpha_{[01]} q_{[01]} \\
q_{[01]} &= \alpha_{[10]} q_{[10]} + \alpha_{[00]} q_{[00]} \\
q_{[10]} &= (1 - \alpha_{[01]}) q_{[01]} + (1 - \alpha_{[11]}) q_{[11]} \\
q_{[00]} &= (1 - \alpha_{[00]}) q_{[00]} + (1 - \alpha_{[10]}) q_{[10]}
\end{aligned}$$

Combining 1st and 3rd equality (or 2nd and 4th) we get

$$q_{[10]} = q_{[01]}.$$

Moreover,

$$\frac{q_{[01]}}{q_{[11]}} = \frac{q_{[10]}}{q_{[11]}} = \frac{1 - \alpha a_{[11]}}{\alpha_{[01]}} \quad \text{and} \quad \frac{q_{[01]}}{q_{[00]}} = \frac{q_{[10]}}{q_{[00]}} = \frac{\alpha_{[00]}}{1 - \alpha_{[10]}}.$$

With the strategies described in the proposition, the high type seller follows a path which features

$$\alpha_{[11]} = 0, \alpha_{[10]} = \alpha, \alpha_{[01]} = \alpha_{[00]} = 1.$$

Then, letting q^θ_{xy} represent the frequency with which type θ seller visits history xy, we have

$$q^H_{[11]} = q^H_{[10]} = q^H_{[01]} = \frac{1}{4 - \alpha} \text{ and } q^H_{[00]} = \frac{1 - \alpha}{4 - \alpha}.$$

Low type follows this same path with probability γ_1. With probability γ_2 the low type follows a path featuring $\alpha_{[11]} = 0, \alpha_{[10]} = \alpha_{[01]} = \alpha_{[00]} = 1$. Letting q'_{xy} represent the frequency with which this path visits history xy, we have

$$q'_{[11]} = q'_{[10]} = q'_{[00]} = \frac{1}{3}, q'_{[00]} = 0.$$

Finally, with probability γ_3, the low type follows a path featuring $\alpha_{xy} = 1$, for all $xy \in \{[11], [10], [01], [00]\}$, and thus this path visits only [11]. Then using Bayes rule and appealing to the improper uniform belief about calendar time, we obtain

$$\frac{\mu_{[00]}}{1-\mu_{[00]}} = \frac{\mu_0}{1-\mu_0}\frac{1}{\gamma_1} \tag{9}$$

$$\frac{\mu_{[10]}}{1-\mu_{[10]}} = \frac{\mu_{[01]}}{1-\mu_{[01]}} = \frac{\mu_0}{1-\mu_0}\frac{\frac{1}{4-\alpha}}{\gamma_1\frac{1}{4-\alpha} + \gamma_2\frac{1}{3}} \tag{10}$$

Further, since $\mu_{[00]}, \mu_{[01]}, \mu_{[10]} > \mu_0$, it must be that $\mu_{[11]} < \mu_0$.

Given the seller strategies, it is trivially observed that at [01] and [00] the suggested buyer strategies form a bidding equilibrium. That no buyer has a profitable deviation at [10] follows from the fact that the high type seller's reservation price is exactly $\hat{P}(\mu_{[10]})$. Finally, at [11], the expected quality is less than c_H, which is a lower bound on the high type's reservation price. Thus, all buyers offering v_L is a bidding equilibrium.

Let $\overline{P} = \hat{P}(\mu_{[00]})$ and $\underline{P} = \hat{P}(\mu_{[10]}) = \hat{P}(\mu_{[01]})$. Next, we characterize the seller's value based on the description of his reservation prices and buyer offer strategies, compute the implied reservation prices, and verify that $\overline{P}$, $\underline{P}$, α, $\gamma_1, \gamma_2, \gamma_3$ can be chosen to satisfy the description, establishing that if those indeed are the price offers, the seller is playing a best response.

Let V^θ_{xy} represent type-θ seller's reservation price at history $xy \in \{[11], [10], [01], [00]\}$. Since $P^H_{[10]} = \underline{P}$ and this is the highest equilibrium offer at [10], the high type's payoff can be calculated along the off-path history where he rejects $\underline{P}$ at [10]. Then,

$$\begin{aligned}
V^H_{[11]} &= \delta^2 V^H_{[00]} \\
V^H_{[10]} &= \delta V^H_{[00]} \\
V^H_{[01]} &= (1-\delta)(\underline{P} - c_H) + \delta^3 V^H_{[00]} \\
V^H_{[00]} &= (1-\delta)\frac{\overline{P} - c_H + \delta(\underline{P} - c_H)}{1-\delta^4}
\end{aligned}$$

It is immediately observed that $P^H_{[01]} < P^H_{[10]}$. The requirement that $P^H_{[10]} = \overline{P}$ is equivalent to $\underline{P} - c_H = \delta(V^H_{[00]} - V^H_{[01]})$. That is,

$$\underline{P} - c_H = \delta\left[\frac{(1-\delta^3)(\overline{P} - c_H) - (1-\delta)(\underline{P} - c_H)}{1-\delta^4}\right] \tag{11}$$

Since low type's reservation price is always v_L, it is possible to calculate his continuation values along the off equilibrium path where he accepts all offers of v_L. This yields,

$$\begin{aligned}
V^L_{[11]} &= v_L - c_L \\
V^L_{[10]} &= (1-\delta)(\alpha+\delta)(\underline{P} - v_L) + v_L - c_L \\
V^L_{[01]} &= (1-\delta)(\underline{P} - v_L) + v_L - c_L \\
V^L_{[00]} &= (1-\delta)(\overline{P} - v_L) + (1-\delta)\delta(\underline{P} - v_L) + v_L - c_L
\end{aligned}$$

The requirement that $P^L_{xA} - v_L$ is equivalent to $v_L - c_L = \delta(V^L_{[10]} - V^L_{[11]})$, which yields

$$v_L - c_L = (\alpha+\delta)(\underline{P} - v_L). \tag{12}$$

The requirement that $P^L_{xR} - v_L$ is equivalent to $v_L - c_L = \delta(V^L_{[00]} - V^L_{[01]})$, which yields

$$v_L - c_L = \overline{P} - v_L - (1-\delta)(\underline{P} - v_L). \tag{13}$$

To argue that there exists $\alpha, \underline{P}, \overline{P}, \gamma_1, \gamma_2, \gamma_3$ that solve (11),(12),(13), (9) and (10) for large enough δ and small enough μ_0, we show this existence for $\delta = 1$ and small enough μ_0. The claim follows because all conditions are continuous in δ. Solving for $\underline{P}, \overline{P}, \alpha$ from (11),(12),(13) with $\delta = 1$ yields Then, from the above analysis:

$$\begin{aligned}
\overline{P} &= 2v_L - c_L \\
\underline{P} &= \frac{6v_L - 3c_L + 2c_H}{5} \\
\alpha &= \frac{4v_L - 2c_L - 2c_H}{v_L + 2c_H - 3c_L}.
\end{aligned}$$

Note that $\alpha > 0$ because $v_L - c_L > (c_H - c_L)$. Further $\alpha < 1$ is equivalent to $(v_L - c_L) < 2(c_H - c_L)$, which is always satisfied. Next, we solve for γ_1 and γ_2. Note that

$$\frac{\overline{P} - v_L}{v_H - \overline{P}} = \frac{\mu_{[00]}}{1 - \mu_{[00]}} = \frac{\mu_0}{1-\mu_0}\frac{1}{\gamma_1}.$$

Thus,

$$\gamma_1 = \frac{\mu_0}{1-\mu_0}\left[\frac{v_H + c_L - 2v_L}{v_L - c_L}\right].$$

Next,

$$\frac{\underline{P} - v_L}{v_H - \underline{P}} = \frac{\mu_{[10]}}{1-\mu_{[10]}} = \frac{\mu_0}{1-\mu_0}\frac{\frac{1}{4-\alpha}}{\frac{1}{3}\gamma_2 + \frac{1}{4-\alpha}\gamma_1}.$$

Substituting $\underline{P}$ and γ_1 and solving for γ_2 yields

$$\gamma_2 = \frac{\mu_0}{1-\mu_0}(v_H - v_L)\frac{3}{5}\left\{\frac{2}{c_H - c_L} - \frac{1}{v_L - c_L}\right\}.$$

By inspection, $\gamma_2 > 0$, while $\gamma_1 > 0$ because $v_L - c_L < (v_H - c_L)/2$. Further, when μ_0 is sufficiently small, $\gamma_1 + \gamma_2 < 1$ so that $\gamma_3 = 1 - \gamma_1 - \gamma_2 > 0$. This completes the proof.

Proposition 3 (TSS equilibrium, weak incentives to mimic) *Assume that $(v_H - c_L)/2 < v_L - c_L$. There exists $\overline{\mu}$ and $\overline{\delta}$ such that when $\mu_0 < \overline{\mu}$, $\delta > \overline{\delta}$, and $k = 2$ there exists an equilibrium in which*

- *the buyers*
  - *offer v_L at [11];*
  - *offer $\hat{P}(\mu_{[01]})$ at [01];*
  - *offer $\hat{P}(\mu_{[00]}) = v_H$ at history [00];*
  - *randomize so that the high offer is $\hat{P}(\mu_{[10]})$, each with positive probability at history [10];(let α stand for the probability of the former);*
- *the high type seller's reservation prices satisfy*

$$c_H \leq P^H_{[11]} = P^H_{[01]} < \hat{P}(\mu_{[01]}) = P^H_{[10]} = P^H_{[00]} < \hat{P}(\mu_{[00]}),$$

  and he always accepts his reservation price when offered;
- *low type seller's reservation prices satisfy*

$$P^L_{[01]} = P^L_{[11]} = v_L > P_{[10]} = P_{[00]},$$

  and he randomizes over 2 pure strategies:
  - *rejects his reservation price at [11] and [01] but accepts his reservation price at [10] and [00] (probability γ_2)*
  - *always accepts his reservation price (probability $\gamma_3 = 1 - \gamma_2$).*

Proof: The proof is analogous. The modifications are to set $\gamma_1 = 0$ when calculating beliefs in (9) and (10), so that $\mu_{[00]} = 1$, and $\mu_{[01]} > \mu_0$ and $\underline{P}$ is calculated from (11) by setting $\overline{P} = v_H$. The rest of the arguments follow unchanged.

Equilibrium payoffs

Naturally, as $\delta \to 1$, for a given type θ, V_{xy}^{θ} becomes equal for any xy. Then, it immediately follows by setting $\delta = 1$ in the computed value functions that the low type's equilibrium payoff is $v_L - c_L$. The high type's equilibrium payoff in this limit is $\frac{\underline{P}+\overline{P}}{4}$.

For the case where $(v_H - c_L)/2 > v_L - c_L > (c_H - c_L)/2$,

$$\overline{P} = 2v_L - c_L \quad \text{and} \quad \underline{P} = \frac{3(v_L - c_L) + 2(c_H - v_L)}{5},$$

thus the high type's payoff at this limit is

$$\frac{2}{5}(2v_L - c_L - c_H).$$

For the case where $(v_H - c_L)/2 < v_L - c_L$,

$$\overline{P} = v_H \quad \text{and} \quad \underline{P} = \frac{3v_H + c_H}{5},$$

and thus the high type's payoff at this limit is

$$\frac{8v_H + c_H}{20} - \frac{c_H}{2} = \frac{2}{5}(v_H - c_H).$$