Trading Dynamics in the Market for Lemons

Ayça Kaya[*] and Kyungmin Kim[†]

April 2013

Abstract

We present a dynamic model of trading under adverse selection. A seller faces a sequence of randomly arriving buyers, each of whom receives a noisy signal about the quality of the asset and makes a price offer. We show that there is generically a unique equilibrium and fully characterize the resulting trading dynamics. Buyers' beliefs about the quality of the asset gradually increase or decrease over time, depending on the initial level. The rich trading dynamics provides a way to overcome a common criticism on dynamic adverse selection, thereby broadening its applicability. We also show that improving asset transparency may lead to gains or losses in efficiency.

JEL Classification Numbers: C73, C78, D82.

Keywords : Adverse selection; market for lemons; inspection; time-on-the-market.

1 Introduction

When an asset has been for sale for a while, what would buyers infer about its quality? The answer depends on the perceived source of the delay. First, it could be just that no potential buyer has shown up yet (search frictions). This source is independent of the quality of the asset. Second, the seller might have a high reservation value for the asset and, therefore, be unwilling to settle at a low price. If this is the perceived source of the delay, a longer delay indicates a higher quality. Finally, potential buyers might have observed an unfavorable attribute of the asset and, therefore, decided not to purchase it. In this case, clearly buyers get more pessimistic about the quality of the asset over time.

We develop a model that incorporates all of these three sources and study the resulting equilibrium dynamics. Conceivably, depending on the conditions in a given market, each one of them

[*]University of Iowa. Contact: ayca-kaya@uiowa.edu
[†]University of Iowa. Contact: kyungmin-kim@uiowa.edu

can be the overriding factor dictating trading dynamics. It is for this reason that all have been extensively, yet typically separately, studied in the literature: the first in the literature on sequential search, going back to Stigler (1961), the second in the recent literature on dynamic adverse selection,[1] and the last in the literature on observational learning, pioneered by Benerjee (1992) and Bikhchandani, Hirsleifer and Welch (1992). However, the existing literature is silent on how these sources are linked together and to the other aspects of the market environment. Understanding these links is not only of theoretical interest but is also crucial in informing policies aimed at alleviating inefficiencies due to adverse selection. We think that by clarifying the interplay among various sources of delay, our paper provides a deeper understanding of dynamic adverse selection, as well as a richer context for policy evaluation.

We consider the problem of the seller who possesses an indivisible asset and faces a randomly arriving sequence of buyers. The asset is either of low quality (low type) or of high quality (high type). There are always gains from trade, but the quality of the asset is known only to the seller. Each buyer, upon arrival, receives a noisy signal about the quality of the asset, which we interpret as the outcome of inspection, and makes a take-it-or-leave-it offer. Importantly, we assume that each buyer observes how long the asset has been up for sale (time-on-the-market).

We show that there is generically a unique equilibrium in this dynamic trading problem, and equilibrium trading dynamics crucially depend on the asset's initial reputation (i.e. buyers' prior beliefs about the quality of the asset). If an asset is likely to be of high quality, then delay typically results from unfavorable inspection outcomes, and thus the asset's reputation (the probability that the asset is of high quality) declines over time. In contrast, if the asset is likely to be of low quality, then delay stems mainly from the seller's rejecting low prices. In this case, the seller's reputation increases over time. Intuitively, the higher the seller's reputation is, the more likely do buyers offer a high price. Therefore, if the seller's reputation is high, even the low-type seller is unwilling to accept a low price. This means that trade is delayed only when, despite the seller's high reputation, buyers are unwilling to offer a high price, which is when they received unfavorable inspection outcomes. The low type is more likely to generate an unfavorable inspection outcome than the high type, and thus the asset is less likely to be of high quality, the longer it stays on the market. If the seller's reputation is low, then the low type trades even at a low price, while the high type still accepts only a high price. Since delay arises due to the high type's insistence on a high price, the seller's reputation improves over time. Interestingly, buyers' beliefs converge to a certain level, whether the seller's initial reputation is high or low. This is the level at which the seller's reputation is such that the two effects are exactly balanced.

The role of search frictions is worth explaining. First, as mentioned earlier, search frictions are

[1] See Evans (1989), Vincent (1989, 1990), Janssen and Roy (2002), Deneckere and Liang (2006), Hörner and Vieille (2009), and Moreno and Wooders (2010) for some seminal contributions

neutral to the *direction* of the evolution of beliefs, because they affect both types equally. However, they affect the *speed* of the evolution. Buyers can never exclude the possibility that the seller has been so unfortunate that no buyer has ever contacted the seller yet. This forces buyers' beliefs to change gradually. Second, search frictions are responsible only for a portion of delay. As is true in other dynamic models of adverse selection, the expected time to trade remains bounded away from zero even if search frictions are arbitrarily small. In the limit, which we study in Section 5, buyers' beliefs immediately jump to a stationary level, but trade does not necessarily take place immediately.

The richness of our equilibrium dynamics provides a way to overcome a common criticism on the literature on dynamic adverse selection, thereby contributing to its applicability. The literature is growing fast,[2] mainly because it has the potential to provide a synthetic theory of several forms of market inefficiencies, such as trading delay (liquidity), market freeze (breakdown), and inefficient assignments, and thus can be used to address various policy issues, including the policies that have been implemented or stipulated after the recent financial crisis. Yet, most existing studies present only one form of equilibrium dynamics: it has been repeatedly found that either trading is immediate (when the initial reputation is above some threshold) or buyers' beliefs only increase over time (below the threshold). Casual observations, however, suggest that a high-quality asset tends to trade faster than a low-quality asset (in other words, the longer an asset stays on the market, the more likely is its quality to be low) in various markets. Although it might be controversial whether information asymmetries are indeed present in such markets, the inability to generate such dynamics clearly limits the applicability of dynamic adverse selection. Our results suggest that with only one additional but plausible modeling innovation (buyers' receiving noisy signals about the quality of the asset), trading dynamics under adverse selection can be very much enriched, and a broader set of empirical patterns can be accommodated.

Our model, with its unique equilibrium and clean characterization, is particularly suited for evaluating government policies that influence the market structure.[3] Within our framework, probably the most intriguing exercise would have to do with the informativeness of buyers' signals (i.e., the quality of buyers' inspection technology). It is widely accepted that asset (corporate) transparency improves market efficiency by facilitating socially desirable trade. Such beliefs have been reflected in recent government policies, such as the Sarbanes-Oxley act passed in the aftermath of the Enron scandal and the Dodd-Frank act passed in the aftermath of the recent financial crises,

[2]A non-exhaustive list includes Camargo and Lester (2011), Chang (2010), Chari, Shourideh and Zetlin-Jones (2010), Chiu and Koeppl (2011), Choi (2013), Daley and Green (2012), Guerrieri and Shimer (2012), Hörner and Vieille (2009), Kurlat (2010), Lauermann and Wolinsky (2013), Moreno and Wooders (2012), Roy (2012), and Zhu (2012).

[3]For instance, the government may mitigate search frictions by centralizing matching and/or trading mechanisms. It may provide subsidies to trading parties, so as to alleviate their incentive problems. Our complete and neat characterization make it fairly straightforward to analyze the effects of those and related policies.

both of which include provisions for stricter disclosure requirements on the part of the sellers. Presumably, the main goal of such policies is to help buyers assess the merits and risks of financial assets more accurately. In our model this corresponds to an increase in the informativeness of buyers' signals.

We demonstrate that enhancing asset transparency does not necessarily lead to efficiency gains. In particular, we show that an increase of the precision of buyers' signals can increase or decrease the expected time to trade for each type, depending on the seller's initial reputation. More precisely, if buyers believe that the asset is likely to be of low quality, then more precise signals speed up trade for both types, while they unambiguously slow down trade in the opposite case.[4] This mixed result arises, because an increase of the informativeness influences buyers' inferences from the length of delay as well as from the signal that they receive. Intuitively, each buyer must take into account the fact that not only his signals, but also all other buyers' signals have become more precise. This indirect effect can work in the opposite direction to the direct one and be significant, which makes the overall effects ambiguous.

Related Literature

As explained above, most existing studies on dynamic adverse selection generate one form of equilibrium dynamics (that buyers' beliefs that the asset is of high quality increase over time). One notable exception is Taylor (1999). He studies a two-period model in which the seller faces a random number of buyers and conducts a second-price auction in each period. He considers several settings which differ in the observability of first-period trading outcomes (in particular, inspection outcome and price history) by second-period buyers. In all settings, buyers assign a lower probability to the event that the asset is of high quality in the second period than in the first period (that is, buyers' beliefs decrease over time). The logic behind the evolution of beliefs is similar to ours: The high type generates good signals more often than the low type. Therefore, a high-quality asset is more likely to be traded in the first period than a low-quality asset. However, the opposite form of trading dynamics (buyers' beliefs increase over time) is absent in his model.[5] In addition, since his model has only two periods, trading dynamics is not as rich and complete as ours.

Two papers study a similar model to ours. Lauermann and Wolinsky (2013) investigate the ability of prices to aggregate dispersed information in a setting very close to ours: an informed player (buyer in their model) faces an infinite sequence of uninformed players, who each receives

[4]There is an intermediate range in which trade of the high type speeds up, while that of the low type slows down. For a more complete picture, see Figure 4.

[5]Taylor (1999) assumes that there are no gains from trade of a low-quality asset. Therefore, buyers never make an offer that can be accepted only by the low type. This forces the high type to always trade faster than the low type.

a noisy signal about the informed player's type. There are two important modeling differences. First, in their model, the informed player makes a take-it-or-leave-it offer to each uninformed player. They adopt the bargaining protocol, because they look for a condition under which the price fully reflects the underlying value of the good (translating into our framework, the price must be equal to v_H (v_L) if high (low) quality). This, of course, creates a severe equilibrium multiplicity problem, as is typical in signalling games. They circumvent the problem by restricting attention to undefeated equilibria. Second, more importantly, in their model, uninformed players do not observe the informed player's time-on-the-market (more generally, the informed player's past behavior) and, therefore, make inferences only based on their own signals. This creates a non-trivial inference problem on the part of uninformed players, as their actual beliefs do not have to coincide with their initial beliefs. However, it essentially forces buyers' beliefs (and strategies) to be always stationary (that is, buyers' beliefs do not evolve over time), while the evolution of uninformed players' beliefs and the resulting trading dynamics are the main focus of this paper.

Zhu (2012) considers a model in which an informed seller can contact only a finite number of buyers. As in Lauermann and Wolinsky (2013), he does not allow buyers to observe the seller's time-on-the-market. However, he assumes that the seller can contact an identical buyer repeatedly, and each buyer knows whether the seller has visited him before or not. Naturally, the seller revisits a buyer only after she has contacted all other buyers. Therefore, buyers make different inferences about the seller's outside options, depending on whether she has visited before or not, which affects their optimal offer strategies. Indeed, this new type of inference is the main focus of the paper. Consequently, the resulting trading dynamics is qualitatively different from ours.

Finally, Daley and Green (2012) study the role of arrival of exogenous information ("news") about the quality of the asset in a setting similar to ours. The most crucial difference from ours is that news is public information to all buyers.[6] Therefore, buyers do not face any inference problem regarding other buyers' signals. This makes their trading dynamics quite distinct from ours.[7] Similarly to us, they also explore the effects of increasing the quality of news and find that it is not always efficiency-improving. However, the mechanism leading to the conclusion is quite different from ours.

The rest of the paper is organized as follows. We formally introduce the model in Section 2. We present several useful properties of the players' optimal strategies and beliefs in Section 3 and provide a formal characterization in Section 4. We then present the limit equilibrium outcome as search frictions vanish in Section 5. We investigate the effects of improving the informativeness of

[6] There are two other modeling differences. First, the seller always faces a competitive pool of buyers (in other words, at least two buyers in each time). Second, there are no search frictions: the seller continuously receives price quotes.

[7] The difference persists even in the limit as search frictions disappear in our model. See Section 5 for our frictionless market outcome.

buyers' signals in Section 6. Omitted proofs are collected in the Appendix.

2 The Model

A seller wishes to sell an indivisible asset. Time is continuous, and the time the seller comes to the market is normalized to 0. Potential buyers arrive sequentially according to a Poisson process of rate $\lambda > 0$. Once a buyer arrives, he receives a private signal about the quality of the asset and offers a price. If the seller accepts the price, then they trade and the game ends. Otherwise, the buyer leaves, while the seller waits for subsequent buyers. The seller discounts future payoff at rate $r > 0$.

The good is either of low quality (L) or of high quality (H). If low quality, the seller derives a flow payoff of rc_L from owning the asset, while a buyer, once he acquires it, receives a flow payoff of rv_L. The corresponding values for high quality are rc_H and rv_H, respectively. There are always gains from trade: $c_L < v_L$ and $c_H < v_H$. However, the quality of the asset is private information to the seller. It is commonly known that the probability that the asset is of high quality at time 0 is equal to $\widehat{q} \in (0, 1)$.

Upon arrival, each buyer receives a private signal about the quality of the asset. A signal s is drawn from the set $S = \{s_1, ..., s_N\}$. The signal generating process depends on the quality of the asset. If the quality is low (respectively, high), then the probability that a buyer receives signal s_n is given by $\gamma_L(s_n)$ (respectively, $\gamma_H(s_n)$). Without loss of generality, assume that the likelihood ratio $\frac{\gamma_H(s_n)}{\gamma_L(s_n)}$ is strictly increasing in n, so that the higher a signal is, the more likely is it that the asset is of high quality. For later use, let $\Gamma_a(s_n)$ ($\Gamma_a^-(s_n)$) be the probability that a buyer receives a signal *weakly* (*strictly*) below s_n from a type a seller. Formally, for each $a = L, H$, $\Gamma_a(s_n) \equiv \sum_{n' \leq n} \gamma_a(s_{n'})$, while $\Gamma_a^- \equiv \sum_{n' < n} \gamma_a(s_{n'})$.

We assume that buyers observe (only) how long the asset has been up for sale (i.e. time t).[8] This is certainly consistent with the arrangement in the housing market and also seems to be plausible in several other markets. In addition, it allows us to study trading dynamics under adverse selection, which is the main focus of the paper, without raising additional complications.[9] The information structure also has an important technical advantage. For any t, there is a positive probability ($e^{-\lambda t}$) that no buyer has arrived and, therefore, trade has not occurred by time t. Therefore, there are no off-the-equilibrium paths, which implies that all beliefs can be obtained through Bayesian

[8]It is well-known that the information buyers have about past histories of the game plays a crucial role in this type of games. See Nöldeke and Van Damme (1990), Swinkels (1999), Hörner and Vieille (2009), Kim (2012), Kaya and Liu (2012), and Fuchs, Öry and Skrzypacz (2012). We note that in our model, buyers observe neither the number of buyers who had arrived before nor the offers they had made to the seller.

[9]For alternative approaches, see Zhu (2012) and Lauermann and Wolinsky (2013).

updating.[10]

The offer strategies of buyers are represented by a function $\sigma_B : \mathbb{R}_+ \times S \times \mathbb{R}_+ \to [0,1]$, where $\sigma_B(t,s,p)$ denotes the probability that the buyer who arrives at time t and receives a signal s offers a price p to the seller. The offer acceptance strategy of the seller is represented by a function $\sigma_S : \{L,H\} \times \mathbb{R}_+ \times \mathbb{R}_+ \to [0,1]$ where $\sigma_S(a,t,p)$ denotes the probability that a type a seller accepts price p in period t. An outcome of the game is a tuple (t,p) where t denotes the time of trade and p is the accepted price. All agents are risk neutral. If the seller accepts price p at time t and her type is $a \in \{L,H\}$, then her payoff is $(1-e^{-rt})c_a + e^{-rt}p$, while the payoff of the buyer who offered the price is $p - v_a$. All other buyers obtain zero payoff.[11]

We adopt the perfect Bayesian equilibrium concept. The concept necessitates the specification of agents' beliefs at each information set. Let $q(t)$ represent buyers' beliefs that the seller who has not traded until t is the high type. Then, a tuple (σ_S, σ_B, q) is a perfect Bayesian equilibrium if (1) given σ_S and q, for any t and s, $\sigma_B(t,s,p) > 0$ only when p maximizes the expected payoff of the buyer with signal s at time t, (2) given σ_B, for any a and t, $\sigma_S(a,t,p) > 0$ only when p is weakly greater than type a seller's continuation payoff at time t, and (3) given σ_S and σ_B, for any t, $q(t)$ is derived through Bayesian updating.

In order to avoid trivial cases, we focus on the environment that satisfies the following assumption.

Assumption 1

$$r(v_L - c_L) < \lambda(c_H - v_L).$$

The left-hand side is buyers' willingness-to-pay for the low-quality asset, while the right-hand side is the low-type seller's reservation price when all subsequent buyers offer c_H regardless of their signal. If this assumption is violated, then the seller types can be easily separated and the resulting trading dynamics is trivial. We note that $c_H > v_L$ is necessary for this assumption to hold, but a stronger assumption commonly adopted in the literature, $\widehat{q}v_H + (1-\widehat{q})v_L < c_L$, is not.

3 Preliminary Observations

In this section, we provide some useful observations regarding the players' equilibrium strategies and beliefs.

[10]This is certainly not the case, for example, if the number of previous buyers or past offers are observable.

[11]This means that buyers have no outside options. The accommodation of positive outside options is fairly straightforward and does not alter any result in a significant way.

3.1 Seller's Optimal Acceptance Strategy

Since the seller's action of rejecting a price is not observable to future buyers, it is clear that each type of the seller adopts a "reservation price" strategy, accepting all prices above her reservation price and rejecting all prices below it. In addition, for the same reasoning as in the Diamond paradox, buyers never offer a price strictly above c_H. This implies that the high-type seller's reservation price is always equal to her reservation value of the asset c_H. In what follows, we denote by $p(t)$ the low-type seller's reservation price at time t. Since $c_L < c_H$, $p(t)$ is always strictly smaller than c_H. We also drop the seller's intrinsic type from the arguments of her acceptance strategy and use $\sigma_S(t, p)$, instead of $\sigma_S(L, t, p)$, to denote the low-type seller's acceptance strategy.

To fully characterize an equilibrium, it is necessary to determine the low-type seller's acceptance strategy when a buyer's offer is exactly equal to $p(t)$. She is, of course, indifferent between accepting and rejecting $p(t)$. However, equilibrium requires that if $p(t) < v_L$, then she accept $p(t)$ with probability 1. This results from the optimality of buyers' offer strategies. If the low type does not accept $p(t)(< v_L)$ with probability 1, then the buyer can offer a price slightly above $p(t)$. Since the alternative offer would be accepted with probability 1 by the low type, it would strictly increase the buyer's expected payoff. This holds for any price above $p(t)$, and thus in equilibrium the low type must accept $p(t)$ with probability 1.

3.2 Buyers' Optimal Offer Strategies

Without loss of generality, we assume that each buyer offers either the reservation price of the low type $p(t)$ or that of the high type c_H. This assumption incurs no loss of generality for the following reasons. First, it is strictly suboptimal for a buyer to offer strictly above c_H or between $p(t)$ and c_H. Second, if in equilibrium a buyer makes a losing offer (a price strictly below $p(t)$), then it suffices to set his offer to be equal to $p(t)$ and the low type's acceptance strategy $\sigma_S(t, p(t))$ to reflect her rejection of the buyer's losing offer. For example, if the buyer at time t with signal s is designated to offer a price below $p(t)$ with probability 1, then his offer can be set to be equal to $p(t)$, and the low-type seller can be assumed to reject the price with probability 1.

Given the low-type seller's reservation price $p(t)$, the buyer's optimal strategy is a "cutoff signal" strategy, offering c_H if his signal is above the cutoff signal, while offering $p(t)$ if his signal is below it. This follows immediately from the fact that the information structure exhibits the monotone likelihood ratio property, and thus the buyer's interim belief is strictly increasing in his signal.

3.3 Evolution of Beliefs

The previous results have an important implication on the way buyers' beliefs evolve over time. Suppose $p(t) < v_L$. In this case, the low-type seller accepts both $p(t)$ and c_H with probability 1. Therefore, her exit rate is equal to the arrival rate of buyers λ. On the other hand, the high-type seller accepts only c_H. Therefore, her exit rate is equal to the arrival rate of buyers λ multiplied by the probability that the buyer's signal exceeds a certain cutoff. Since the latter is always smaller than the former (the low type trades faster than the high type), buyers' beliefs about the seller's type necessarily *increase* over time. If $p(t) > v_L$, then both types accept only c_H. But, the high type generates good signals (above a cutoff signal) with a higher probability than the low type. Therefore, the exit rate of the high type exceeds that of the low type. This induces buyers' beliefs to *decrease* over time.

The following lemma summarizes the most useful observations of this section.

Lemma 1 *If $p(t) < v_L$ (respectively, $p(t) > v_L$), then the low-type seller accepts (respectively, rejects) $p(t)$ with probability 1, and thus buyers' beliefs $q(t)$ increase (respectively, decrease) over time. The changes of buyers' beliefs are strict unless buyers offer c_H regardless of their signal.*

4 Equilibrium Characterization

Equipped with the observations in the previous section, we now turn to the equilibrium characterization. The unique equilibrium of our model exhibits the following intuitive, yet not obvious, properties:

(i) The low-type seller's reservation price $p(t)$ depends only on buyers' beliefs $q(t)$ and is increasing in $q(t)$. In a slight abuse of notation, we denote by $p(q)$ the low-type seller's reservation price when buyers' beliefs are equal to q.

(ii) There is a finite partition of the belief space, $\{\overline{q}_{N+1} = 0, \overline{q}_N, ..., \overline{q}_1, \overline{q}_0 = 1\}$, which informs the cutoff signal above which each buyer offers c_H. Specifically, if a buyer's prior belief $q(t)$ belongs to the interval $(\overline{q}_{n+1}, \overline{q}_n)$, then he offers c_H if and only if his signal is strictly above s_n.[12]

(iii) Buyers' beliefs conditional on no trade converge monotonically to a stationary value q^*.

Figure 1 illustrates these properties in an example with 5 signals.

In this section, we first construct an equilibrium that satisfies (i) and (ii), and illustrate the resulting trading dynamics. Our construction makes it clear that for a generic set of parameter

[12] If $q(t) > \overline{q}_1$, then he offers c_H regardless of his signal, while if $q(t) < \overline{q}_N$, then he never offers c_H.

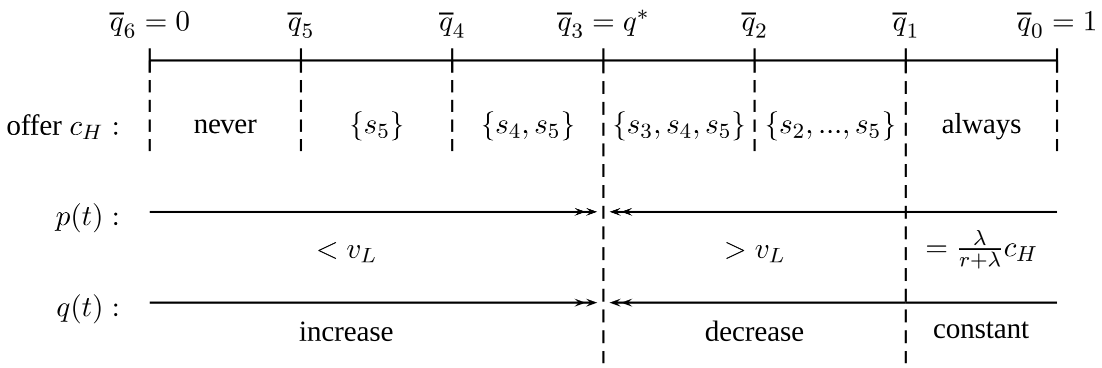

Figure 1: Equilibrium structure when there are 5 signals ($N = 5$) and $q^* = \overline{q}_3$.

values, there is a unique equilibrium that satisfies the conditions. Moreover, the constructed equilibrium necessarily satisfies (iii). We then show that any equilibrium must satisfy (i) and (ii), and thus the constructed equilibrium is the unique equilibrium in this game.[13] Finally, we discuss the non-generic cases where the uniqueness of equilibrium fails.

4.1 Equilibrium construction and trading dynamics

In this section we construct an equilibrium that satisfies properties (i) and (ii), and discuss the trading dynamics that arises in the equilibrium.

4.1.1 Unique Stationary Path

We first consider the behavior along the stationary path where buyers' beliefs remain constant at q^*. We identify the critical belief level q^* and characterize the equilibrium when buyers' beliefs are equal to q^*.

Notice that Lemma 1 implies that $p(q^*) = v_L$: otherwise, $q(t)$ either increases or decreases. In order to utilize this fact, denote by ρ_L (ρ_H) the rate at which the low-type (high-type) seller

[13]The equilibrium properties are fairly intuitive. A buyer is more willing to offer c_H when his interim belief is higher, and for a given signal s_n, his interim belief is strictly increasing in his prior belief. Therefore, he should be *less reluctant* to offer c_H; i.e. his cutoff signal must decrease as his belief increases (property (ii)). Then, from the perspective of the low-type seller, the higher buyers' beliefs are, the more frequently would he receive offer c_H. Therefore, the low-type seller's reservation price must increase in buyers' beliefs (property (i)). However, a priori, it is not clear whether there cannot exist an equilibrium that violates these properties. In particular, if the low-type seller's reservation price were to decline in buyers' beliefs, buyers may be *more reluctant* to offer c_H at higher beliefs, simply because a lower reservation price makes the option of trading with only the low type more attractive. This, in turn, could be consistent with the seller's reservation price falling. The main thrust of our uniqueness proof is to rule out this possibility.

receives offer c_H when buyers' beliefs are equal to q^*. Then, ρ_L must satisfy

$$r(v_L - c_L) = \rho_L(c_H - v_L).$$

Generically, there does not exist n such that $\rho_L = \lambda(1-\Gamma_L(s_n))$. This implies that buyers typically play a mixed offer strategy once their beliefs reach q^*. In what follows, we focus on the generic case where $\rho_L \neq \lambda(1 - \Gamma_L(s_n))$ for any $n = 1, ..., N$. For the sake of completeness, we discuss the non-generic case at the end of this section.

To pin down buyers' equilibrium offer strategies, let n^* be the largest value of n such that $\rho(q^*) > \lambda(1-\Gamma_L(s_n))$. By Assumption 1, n^* is well-defined. In addition, let σ_B^* be the value such that $\rho(q^*) = \lambda(\gamma_L(n^*)\sigma_B^* + 1 - \Gamma_L(s_{n^*}))$. For the generic case we are considering, σ_B^* always lies in $(0, 1)$. By construction, $p(q^*)$ is equal to v_L if all subsequent buyers offer c_H with probability 1 when their signals are strictly above s_{n^*}, with probability σ_B^* when their signals are s_{n^*}, and with probability 0 when their signals are strictly below s_{n^*}.

The identification of s_{n^*} allows us to find the value of q^*. Consider a buyer who had prior belief q^* and received signal s_{n^*}. By Bayes' rule, his belief updates to $\frac{q^*\gamma_H(s_{n^*})}{q^*\gamma_H(s_{n^*})+(1-q^*)\gamma_L(s_{n^*})}$. The buyer must be indifferent between offering c_H and $p(q^*)$. But, since $p(q^*) = v_L$, he must receive zero expected payoff, regardless of the low-type seller's acceptance strategy. This implies that q^* must satisfy

$$q^*\gamma_H(s_{n^*})(v_H - c_H) + (1 - q^*)\gamma_L(s_{n^*})(v_L - c_H) = 0 \Leftrightarrow \frac{1 - q^*}{q^*} = \frac{\gamma_H(s_{n^*})}{\gamma_L(s_{n^*})}\frac{v_H - c_H}{c_H - v_L}.$$

It remains to determine the probability that the low-type seller accepts $p(q^*) = v_L$, which we denote by σ_S^* for notational simplicity. To identify σ_S^*, notice that $q(t)$ is time-invariant if and only if the low type has the same exit rate as the high type. The high type accepts only c_H. Therefore, given buyers' offer strategies characterized by (n^*, σ_B^*), her exit rate is equal to

$$\rho_H = \lambda(\gamma_H(s_{n^*}) + 1 - \Gamma_H(s_{n^*})).$$

If the low type accepts $p(q^*)$ with probability σ_S^*, then her exit rate is equal to

$$\rho_L + (1 - \rho_L)\sigma_S^* = \lambda(\gamma_L(n^*)\sigma_B^* + 1 - \Gamma_L(s_{n^*})) + \lambda(\Gamma_L^-(s_{n^*}) + \gamma_L(n^*)(1 - \sigma_B^*))\sigma_S^*.$$

σ_S^* is the unique value that equates the above two rates, ρ_H and $\rho_L + (1 - \rho_L)\sigma_S^*$. It is well-defined in $(0, 1)$ because $\Gamma_H(\cdot)$ first-order stochastically dominates $\Gamma_L(\cdot)$.

We summarize all the findings on the stationary path in the following proposition.

Proposition 1 (Stationary path) *Let q^*, n^*, σ_B^*, and σ_S^* be the values defined as above. When* *buyers' beliefs are equal to q^*, the following strategy profile constitutes an equilibrium: Buyers' beliefs stay constant. Each buyer offers c_H with probability 1 if his signal is strictly above s_{n^*}, with probability σ_B^* if his signal is s_{n^*}, and with probability 0 otherwise. The low-type seller accepts $p(q^*) = v_L$ with probability σ_S^*. Moreover, this is the unique stationary path; i.e. if buyers' beliefs are to remain constant, then the equilibrium must be as described.*

4.1.2 Evolution of Beliefs

Given n^* and the partition $\{\overline{q}_{N+1}, \overline{q}_N, ..., \overline{q}_1, \overline{q}_0\}$, buyers' beliefs evolve deterministically, conditional on no trade. There are three distinct regions, depending on the location of $q(t)$.

If $q(t) > \overline{q}_1$, then the outcome is trivial. Each buyer offers c_H regardless of his signal. Therefore, trade takes place as soon as a buyer arrives. Since both types exit at the same rate, buyers' beliefs do not change over time.

Suppose $q(t) \in (q_{n+1}, q_n)$ and $0 < n < n^*$ (so that $q(t) \in (q^*, \overline{q}_1)$). In this case, $p(q(t)) > v_L$, and thus in equilibrium the low-type seller accepts only c_H (see Lemma 1). Since the buyer offers c_H if and only if his signal is strictly above s_n, the low type exits the market at rate $\lambda(1 - \Gamma_L(s_n))$, while the high type at rate $\lambda(1 - \Gamma_H(s_n))$. It then follows that $q(\cdot)$ evolves according to the following law of motion:[14]

$$\dot{q}(t) = q(t)(1 - q(t))\lambda(\Gamma_H(s_n) - \Gamma_L(s_n)). \tag{1}$$

Since $\Gamma_H(\cdot)$ first-order stochastically dominates $\Gamma_L(\cdot)$, $\Gamma_H(s_n) - \Gamma_L(s_n)$ is always negative, and thus $q(t)$ strictly decreases in t. In Figure 2, a typical path for buyers' beliefs starting from $\widehat{q} > q^*$ is illustrated by the solid line. The kink at $\overline{q}_2$ is due to the fact that the cutoff signal changes from s_1 to s_2 at that point, and thus $\Gamma_H(\cdot) - \Gamma_L(\cdot)$ as well.

Now suppose $q(t) \in (q_{n+1}, q_n)$ and $n \geq n^*$ (so that $q(t) < q^*$). In this case, $p(q(t)) < v_L$, and thus in equilibrium the low type accepts $p(q(t))$ with probability 1 (see Lemma 1). Therefore, her exit rate is equal to λ. The high type still accepts only c_H, and thus her exit rate is equal to $\lambda(1 - \Gamma_H(s_n))$. As above, it follows that the law of motion of $q(t)$ in this case is given by

$$\dot{q}(t) = q(t)(1 - q(t))\lambda\Gamma_H(s_n). \tag{2}$$

This expression is obviously positive, and thus $q(t)$ strictly increases in t. The dashed line in Figure

[14]Heuristically, by Bayes' rule,

$$q(t + dt) = \frac{q(t)e^{-\lambda(1-\Gamma_H(s_n))dt}}{q(t)e^{-\lambda(1-\Gamma_H(s_n))dt} + (1 - q(t))e^{-\lambda(1-\Gamma_L(s_n))dt}}.$$

The equation can be derived by subtracting $q(t)$ from both sides and dividing by dt.

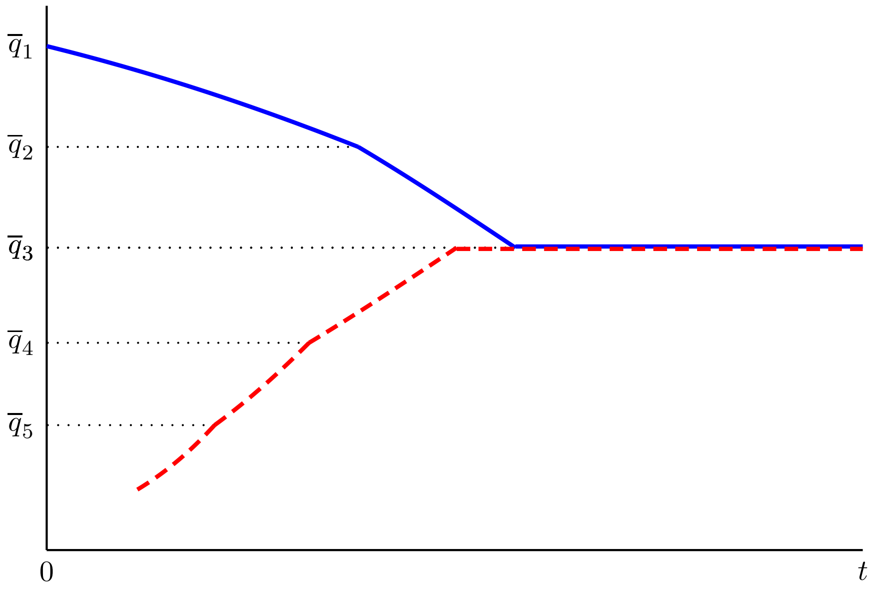

Figure 2: The evolution of buyers' beliefs. The parameter values used for this figure and Figure 3 are as follows: $N = 5$, $\gamma_L(s_n) = (6 - n)/15$, $\gamma_H(s_n) = n/15$, $r = 0.1$, and $\lambda = 0.4$.

2 exemplifies a typical path buyers' beliefs follow from below q^*.

Proposition 2 (Evolution of Beliefs) *Consider an equilibrium that satisfies (i) and (ii), and suppose $q(t) \in (q_{n+1}, q_n)$. If $n = 0$, then $q(t)$ does not change. If $0 < n < n^*$, then $q(t)$ evolves according to (1). If $n \geq n^*$, then $q(t)$ evolves according to (2).*

Two remarks are in order. First, starting from $\widehat{q}$, whether $\widehat{q} \in (q^*, \overline{q}_1)$ or $\widehat{q} < q^*$, $q(\cdot)$ converges to q^* in finite time, because $\dot{q}(t)$ is bounded away from 0.[15] Notice that this implies that property (iii) follows from (i) and (ii). Second, the offer strategies of the buyers with prior belief $\overline{q}_n$ and signal s_n for some $n \neq n^*$ are indeterminate, since such buyers are exactly indifferent between offering c_H or $p(\overline{q}_n)$. However, the offer strategies at such decision nodes do not affect the equilibrium play, in particular, the evolution of buyers' beliefs. This is because the beliefs are strictly monotone, and the arrival rate of buyers is finite so that the probability that a buyer arrives at a

[15] Formally, if $\widehat{q} \in (q^*, \overline{q}_1)$, then $\dot{q}(t)$ is bounded above by

$$\left(\min_{q' \in [q^*, \widehat{q}]} q'(1-q') \right) \lambda \min_{n<n^*} (\Gamma_H(s_n) - \Gamma_L(s_n)) < 0.$$

If $\widehat{q} < q^*$, then $\dot{q}(t)$ is bounded below by

$$\left(\min_{q' \in [\widehat{q}, q^*]} q'(1-q') \right) \lambda \Gamma_H(s_{n^*}) > 0.$$

point at which his belief is equal to one of the cutoffs is 0. In fact, this is the reason why the behavior at these nodes cannot be determined from other equilibrium requirements, unlike in the case of $\overline{q}_{n^*}$. In what follows, without loss of generality, we assume that for each $n \neq n^*$, the buyer with prior belief $\overline{q}_n$ and signal s_n offers c_H with probability 1.

4.1.3 Reservation Prices and Equilibrium Belief Cutoffs

We conclude the equilibrium construction by jointly identifying $p(\cdot)$ and $\{\overline{q}_N, ..., \overline{q}_1\}$. We also establish that given (i) and (ii), both reservation price schedule $p(q)$ and partition $\{\overline{q}_N, ..., \overline{q}_1\}$ are uniquely determined.

First fix $p(q)$. Given the previous characterization of the low-type seller's equilibrium acceptance strategy, each buyer's optimal offer strategy depends only on his belief q and the low-type seller's reservation price $p(q)$. Consider a buyer who had prior belief $\overline{q}_n$ and received signal s_n. The buyer must be indifferent between offering c_H and $p(q)$. Using the fact that his interim belief is equal to $\frac{\overline{q}_n \gamma_H(s_n)}{\overline{q}_n \gamma_H(s_n) + (1-\overline{q})\gamma_L(s_n)}$, it is straightforward to show that $\overline{q}_n$ is the unique value that satisfies

$$\frac{c_H - \min\{v_L, p(\overline{q}_n)\}}{v_H - c_H} = \frac{\overline{q}_n}{1 - \overline{q}_n} \frac{\gamma_H(s_n)}{\gamma_L(s_n)}. \tag{3}$$

The use of $\min\{v_L, p(t)\}$, instead of $p(t)$, reflects the fact that if $p(q) \geq v_L$, then the buyer's expected payoff by offering $p(q)$ is equal to 0, either because the offer itself is v_L, which is accepted only by the low type, or because it is greater than v_L, in which case in equilibrium it is rejected with probability 1.

Now fix the partition $\{\overline{q}_N, ..., \overline{q}_1\}$. Intuitively, $p(t)$ is determined only by the rate at which the low type receives offer c_H: each buyer offers either c_H or her reservation price $p(t)$, which she is indifferent between accepting and rejecting. Therefore, for the purpose of calculating the expected payoff, she can be assumed only to accept c_H. Then, the low-type seller's reservation price for the belief level can be calculated recursively as follows: Suppose $q \in [\overline{q}_{n+1}, \overline{q}_n)$. While buyers' beliefs are in this interval, the low-type seller receives offer c_H at rate $\lambda(1 - \Gamma_L(s_n))$. If $n < n^*$, then $q(t)$ decreases to $\overline{q}_{n+1}$, while if $n \geq n^*$, then it increases to $\overline{q}_n$. Define $\tilde{q}$ so that $\tilde{q} = \overline{q}_{n+1}$ if $n < n^*$, while $\tilde{q} = \overline{q}_n$ if $n \geq n^*$. Then,

$$p(q) = \int_0^{T(q,\tilde{q})} ((1 - e^{-rt})c_L + e^{-rt}c_H) d\left(1 - e^{-\lambda(1-\Gamma_L(s_n))t}\right) + e^{-(r+\lambda(1-\Gamma_L(s_n)))T(q,\tilde{q})} p(\tilde{q}), \tag{4}$$

where $T(q, \tilde{q})$ is the length of time it takes for buyers' beliefs to move from q to $\tilde{q}$.[16] A closed-form solution can be found by applying the definition of $T(\cdot, \cdot)$ and using the fact that $p(q_{n^*}) = v_L$.

Finally, we combine the characterizations and compute the cutoffs $\overline{q}_N, ..., \overline{q}_1$ as well as the reservation price schedule $p(\cdot)$. Their uniqueness follows from the explicit computation.

Consider first $\overline{q}_n$ for $n < n^*$. In this case, $p(q) > v_L$, and thus (3) reduces to

$$\frac{1 - \overline{q}_n}{\overline{q}_n} = \frac{\gamma_H(s_n)}{\gamma_L(s_n)} \frac{v_H - c_H}{c_H - v_L}. \tag{5}$$

Therefore, the cutoffs $\overline{q}_n$ are determined independently of the true value of $p(q)$. The uniqueness follows from the fact that the left-hand side is increasing, while the right-hand side is constant. Given $\overline{q}_n$ for each $n < n^*$, $p(q)$ can be calculated using (4) for any $q > q^*$.

The determination of $\overline{q}_n$ for $n > n^*$ is more involved, because they cannot be identified separately from $p(\cdot)$. Nonetheless, as shown in the previous section, $q(t)$ evolves deterministically from $\overline{q}_{n+1}$ to $\overline{q}_n$ for any $n \geq n^*$. Therefore, $\overline{q}_{n^*+k}$ for each $k = 0, 1, ..., N - n^* - 1$ can be found recursively as follows:

- For each $k = 0, ..., N - n^* - 1$,

$$\begin{aligned} p(\overline{q}_{n^*+k+1}) &= \int_0^{T(q_{n^*+k+1}, q_{n^*+k})} ((1 - e^{-rt})c_L + e^{-rt}c_H) d\left(1 - e^{-\lambda(1-\Gamma_L(s_{n^*+k}))}\right) \\ &\quad + e^{-(r+\lambda(1-\Gamma_L(s_{n^*+k})))T(q_{n^*+k+1}, q_{n^*+k})} p(\overline{q}_{n^*+k}). \end{aligned}$$

- $\overline{q}_{n^*+k+1}$ is the value that satisfies

$$\frac{1 - \overline{q}_{n^*+k+1}}{\overline{q}_{n^*+k+1}} = \frac{\gamma_H(s_{n^*+k+1})}{\gamma_L(s_{n^*+k+1})} \frac{v_H - c_H}{c_H - p(\overline{q}_{n^*+k+1})}.$$

  Such $\overline{q}_{n^*+k+1}$ uniquely exists because $p(\cdot)$ is strictly increasing and it is known that $p(\overline{q}_{n^*}) = v_L$.

As in the previous case, given $\overline{q}_n$ for each $n \geq n^*$, (4) can be used to compute $p(q)$ for each $q < q^*$.

The results are summarized in the following proposition.

[16] Formally, if $n < n^*$, then $T(q, q')$ is defined to be the value that satisfies

$$q' = \frac{q e^{-\lambda(1-\Gamma_H(s_n))T(q,q')}}{q e^{-\lambda(1-\Gamma_H(s_n))T(q,q')} + (1-q) e^{-\lambda(1-\Gamma_L(s_n))T(q,q')}} \Leftrightarrow \frac{1-q'}{q'} = \frac{1-q}{q} e^{-\lambda(\Gamma_H(s_n) - \Gamma_L(s_n))T(q,q')},$$

while if $n \geq n^*$, then

$$q' = \frac{q e^{-\lambda(1-\Gamma_H(s_n))T(q,q')}}{q e^{-\lambda(1-\Gamma_H(s_n))T(q,q')} + (1-q) e^{-\lambda T(q,q')}} \Leftrightarrow \frac{1-q'}{q'} = \frac{1-q}{q} e^{-\lambda(\Gamma_H(s_n) - 1)T(q,q')}.$$

Proposition 3 *There is a unique equilibrium of* $p(\cdot)$ *and* $\{\overline{q}_{N+1} = 0, \overline{q}_N, ..., \overline{q}_1, \overline{q}_0 = 1\}$ *that satisfies properties (i) and (ii).*

Figure 3 illustrates the trading dynamics that emerges from our model. If buyers' beliefs about the asset's quality are rather optimistic ($\widehat{q}_1$), then the low-type seller's reservation price ($p(\widehat{q}_1) = p(0)$) exceeds v_L. In this case, trade takes place only at price c_H. Since the high type generates good signals more often than the low type, buyers' beliefs $q(\cdot)$ decline over time (see the solid line in Figure 2). As $q(\cdot)$ decreases, buyers offer c_H less frequently, and thus the low-type seller's reservation also declines (the solid line in the right panel). Once $q(t)$ becomes equal to q^*, buyers' beliefs do not change thereafter, and the low-type seller's reservation price stays equal to v_L. If buyers' initial beliefs are pessimistic ($\widehat{q}_2$), then the resulting dynamics is exactly the opposite: buyers' beliefs and the low-type seller's reservation price decline over time and converge to q^* and v_L, respectively.

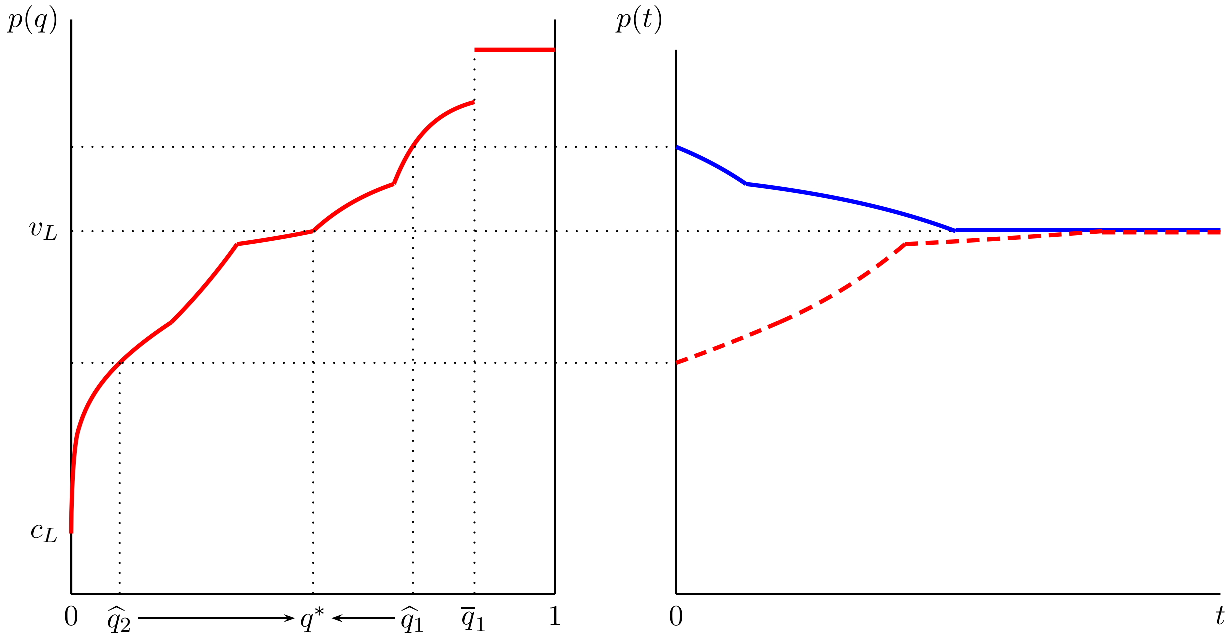

Figure 3: The low-type seller's reservation price as a function of buyers' beliefs q (left) and that of time t (right).

We conclude this section by summarizing all the results regarding the equilibrium construction.

Theorem 1 *For a generic set of parameter values, there exists an equilibrium in which buyers' beliefs evolve as described in Proposition 2. Until buyers' beliefs reach* q^**, buyers' equilibrium offer strategies are characterized by a partition* $\{\overline{q}_N, ..., \overline{q}_1\}$*: If* $q(t) \in (\overline{q}_{n+1}, \overline{q}_n]$*, then the buyer offers* c_H *if his signal is strictly above* s_n *and* $p(q(t))$ *otherwise. The low-type seller accepts* $p(q(t))$

with probability 1 *if* $q(t) < q^*$ *and with probability* 0 *if* $q(t) > q^*$. *Once buyers' beliefs reach* q^*, *then the game is played as described in Proposition 1. This is the unique equilibrium that satisfies properties (i) and (ii).*

4.2 Uniqueness

We now prove that the equilibrium presented in Theorem 1 is the unique equilibrium of our model.

Theorem 2 *For a generic set of parameter values, there is a unique equilibrium.*

In light of Theorem 1, it suffices to show that any equilibrium satisfies properties (i) and (ii). We obtain this result in two steps. First, we show that in any equilibrium, even without reference to properties (i) and (ii), buyers' beliefs evolve monotonically, regardless of their starting point. This implies, among others, that that the low-type seller's reservation price can be regarded as a function of buyers' beliefs q. Second, we show that the reservation price $p(q)$ must be strictly increasing in q whenever $q < \overline{q}^*$. From this, we deduce that buyers' cutoff signals must be non-increasing in q. The latter, in particular, implies that there exists a partition $\{\overline{q}_N, ..., \overline{q}_1\}$ that describes buyers' equilibrium offer strategies as in property (ii).

4.2.1 Monotonicity of beliefs

Lemma 1 already argued that if $p(t) < v_L$, then $q(t)$ is increasing, while if $p(t) > v_L$, then it is decreasing. The next lemma links the ranking of $p(t)$ relative to v_L to the ranking of $q(t)$ relative to q^*. The monotonicity of beliefs immediately follow by combining Lemmas 1 and 2.

Lemma 2 *In any equilibrium,* $q(t) \leq q^*$ *if, and only if,* $p(t) \leq v_L$.

Proof. See the Appendix. ■

This intuitive lemma would immediately follow if it were already established that the low-type seller's reservation price is increasing in buyers' beliefs. In the absence of that result, establishing Lemma 2 requires subtler arguments. The crucial step is to observe that, if, for instance, it were the case that for some $\bar{t}$, $q(\bar{t}) > q^*$ while $p(\bar{t}) < v_L$, the reservation price must eventually converge to v_L from below. This observation can be used to reach a contradiction as follows: just before the convergence occurs, the beliefs must be bounded away from q^*. This follows from Lemma 1: for $t > \bar{t}$ and before the convergence occurs, $p(t) < v_L$, and thus $q(t)$ must be increasing, which implies $q(t) > q(\bar{t}) > q^*$. In the meantime, the low-type seller's reservation value is *almost* v_L. Consider a buyer arriving at such an instant. If he makes an offer of $p(t)$, his payoff is almost 0, since $p(t)$ is almost v_L. On the other hand, upon receiving a signal of s_{n^*}, his payoff from offering

c_H is bounded away from 0: the payoff would be exactly zero if his belief were exactly q^*, yet his belief is strictly above (and bounded away from) q^*. It follows that such a buyer would offer c_H *at least as often as on the stationary path*. Clearly, this provides the low-type seller with a higher payoff than his stationary payoff v_L, leading to a contradiction. We formalize this argument in the Appendix.

4.2.2 Monotonicity of $p(\cdot)$ and buyers' cutoff signals

The monotonicity of beliefs discussed above implies that, in any equilibrium, there is a one-to-one correspondence between the time-on-the-market, which is the only relevant history in this game, and buyers' beliefs q. Therefore, the low-type seller's reservation price, which is formally a function of his time-on-the-market, can be expressed as a function of buyers' beliefs. For the same reason, buyers' offer strategies can also be expressed as functions of q. Let $\bar{s}(q)$ denote the cutoff signal that a buyer with prior belief q uses. Next we establish that $p(q)$ is strictly increasing and $\bar{s}(q)$ is non-increasing.

First consider $q > q^*$. Then by Lemma 2, $\min\{p(t), v_L\} = v_L$. Inspecting (3) immediately reveals that for this range of beliefs, $\bar{s}(q)$ is non-increasing. This, in turn, implies the monotonicity of $p(q)$.

Next consider $q < q^*$. Now, it is not a priori clear that a higher prior belief would be associated with a lower cutoff signal, because if $p(\cdot)$ is decreasing, then offering $p(q)$ could be relatively more attractive when q is higher, and thus the buyer could be more reluctant to offer c_H. Yet, if $p(q)$ is monotone in q, then the monotonicity of the cutoff signals follows. The next lemma establishes that $p(t)$ is increasing over time whenever $q(t) < q^*$. Combining this with the monotonicity of beliefs, the desired result follows.

Lemma 3 *In any equilibrium, if $q(t) < q^*$, then $p(\cdot)$ is strictly increasing in t.*

The argument for this result uses the idea that if $p(\cdot)$ is not increasing, then there exists an interval, say (t, t'), such that $p(\cdot)$ is $\cup$-shaped over the interval and the values at the two end points are identical. In the mean time, since $p(t) < v_L$, buyers' beliefs must increase over the interval. But then, within the interval, the buyers cannot be offering c_H more frequently than *after* time t', since at that point both buyers' beliefs and the low-type seller's reservation price are higher. This leads to a contradiction.

4.2.3 Proof of Theorem 2

Let us summarize how all the components we have established so far lead to the proof of the uniqueness.

Firstly, Proposition 1 establishes that *if* buyers' beliefs are to remain constant at some level, this level must be q^* as defined in 4.1.1. Lemmas 1 and 2 establish that in any equilibrium buyers' beliefs must be decreasing if $q(t) > q^*$, while increasing if $q(t) < q^*$. Since $q(t)$ is continuous in t, buyers' beliefs cannot "jump over" q^* and must remain constant once buyers' beliefs reach q^*.

Next, if $\widehat{q} > q^*$, then, by Lemma 1, the low-type seller also trades only when an offer of c_H is made, until buyers' beliefs reach q^*. This uniquely pins down buyer behavior via (5) for $q > q^*$. Then, the low-type seller's behavior for this range of beliefs is determined by (4).

Finally, if $\widehat{q} < q^*$, then the uniqueness argument requires an extra step, since in this case buyers' equilibrium offer strategies cannot be pinned down independently of the low-type seller's reservation price. In this case, the crucial step is Lemma 3 which establishes that $p(t)$ cannot be decreasing over time. Proposition 3 then proves the uniqueness of equilibrium strategies.

4.3 Non-Generic Cases

We conclude this section by illustrating the equilibrium outcome for the non-generic case, where there exists n^* such that $\rho_L = \frac{r(v_L - c_L)}{\lambda(c_H - v_L} = \lambda(1 - \Gamma_L(s_{n^*}))$.

To illustrate the basic problem of these cases, let $\overline{q}_{n^*}$ and $\overline{q}_{n^*+1}$ be the values such that

$$\frac{1 - \overline{q}_{n^*}}{\overline{q}_{n^*}} = \frac{\gamma_H(s_{n^*})}{\gamma_L(s_{n^*})} \frac{v_H - c_L}{c_L - v_L},$$

and

$$\frac{1 - \overline{q}_{n^*+1}}{\overline{q}_{n^*+1}} = \frac{\gamma_H(s_{n^*+1})}{\gamma_L(s_{n^*+1})} \frac{v_H - c_L}{c_L - v_L}.$$

Suppose $q(t) \in [\overline{q}_{n^*+1}, \overline{q}_{n^*}]$ and all subsequent buyers offer c_H if and only if their signal is strictly above c_H. Then, the low-type seller's expected payoff stays equal to v_L. Suppose she accepts v_L with probability σ_S^*, where σ_S^* satisfies

$$\lambda(1 - \Gamma_H(s_{n^*})) = \lambda(1 - \Gamma_L(s_{n^*}) + \Gamma_L(s_{n^*})\sigma_S^*).$$

Then, buyers' beliefs do not change over time. Furthermore, given that v_L is the low-type seller's reservation price, buyers' offer strategies are also optimal. This implies that any belief level between $[\overline{q}_{n^*+1}, \overline{q}_{n^*}]$ can serve as the critical stationary belief in our model.

In fact, buyers' beliefs do not even need to converge to a certain level. Once buyers' beliefs fall between $\overline{q}_{n^*+1}$ and $\overline{q}_{n^*}$, any belief path that stays within the interval can be supported as an equilibrium: Although buyers' offer strategies are fixed, buyers' beliefs can decrease or increase, depending on the low-type seller's acceptance strategy of offer v_L. For instance, the low-type seller may reject v_L with probability 1, until it reaches $\overline{q}_{n^*}$ and stays constant thereafter. Or, she may accept v_L with probability 1, until buyers' beliefs hit $\overline{q}_{n^*+1}$. Buyers' beliefs may even keep oscillating between (or any two levels between) $\overline{q}_{n^*}$ and $\overline{q}_{n^*+1}$.

Nevertheless, all these equilibria have crucial properties in common. First, within the range of beliefs, the buyers play the same offer strategy across all equilibria, offering c_H if and only if their signal is strictly above s_{n^*}. Therefore, in any equilibrium the low-type seller's reservation price is equal to v_L, once buyers' beliefs fall into the interval $[\overline{q}_{n+1}, \overline{q}_n]$. Second, outside $[\overline{q}_{n+1}, \overline{q}_n]$, buyers' beliefs gradually converge to the interval $[\overline{q}_{n^*+1}, \overline{q}_{n^*}]$, just as they converge to q^* in the generic case. Furthermore, the unique convergence path can be fully characterized as in the generic case: If $q(t) \in [\overline{q}_{n+1}, \overline{q}_n]$ for some $n < n^*$, then the low type trades at rate $\lambda(1 - \Gamma_L(s_n))$, while the high type at rate $\lambda(1 - \Gamma_H(s_n))$. If $n \geq n^* + 1$, then the low type trades at rate λ, while the high type at rate $\lambda(1 - \Gamma_H(s_n))$. Finally, given the first two properties, it follows that all the equilibria are payoff-equivalent, whether the initial belief $\widehat{q}$ belongs to the interval $[\overline{q}_{n+1}, \overline{q}_n]$ or not. The only difference among the equilibria is the low-type seller's trading rate,[17] as it varies depending on the low-type seller's acceptance strategy of offer v_L.

5 Equilibrium Outcomes in the Frictionless Limit

In this section we present equilibrium outcomes for the limit case where λ is arbitrarily large. The limit case is of interest for at least three reasons. First, the market outcome characterized in the previous section is influenced by the level of search frictions as well as information asymmetry. The analysis of the limit case allows us to separate the effects due to the latter from those due to the former. Second, while search frictions are physically inherent in various markets, such as labor markets and over-the-counter markets, they can be mitigated by government policies. For example, the government can increase λ by introducing a more efficient job-matching mechanism or promoting electronic trading, as opposed to over-the-counter trading. The limit case informs us of the extent to which the government can facilitate trade through such policies.[18] Finally, it permits a direct comparison of our model to the existing ones that assume away search frictions (that is, the models in which the seller can trade any time she wants). In this section, we provide the result and intuition, while relegating the formal derivation to the Appendix.

We focus on the time to trade for each type, denoted by $\tau_a(\widehat{q})$ for each $a = L, H$, and the low-type seller's expected payoff from the game $p(\widehat{q})$. Since there are gains from trade, whether the asset is of high or low quality, the surplus generated in the market is larger, the earlier does trade takes place. Therefore, the time to trade can be considered as a measure of surplus generated.

[17] The high-type seller's trading rate is identical across all equilibria, because buyers' offer strategies are identical.

[18] It is fairly straightforward to show that a marginal increase of λ always increases the low-type seller's expected payoff and speeds up trade of both types.

Meanwhile, the low-type seller's expected payoff can be interpreted as a measure of division of surplus: recall that the high-type seller's expected payoff is always equal to 0.

We start with two convenient observations. First, if λ is sufficiently large, then the stationary cutoff signal s^* is necessarily equal to the highest signal s_N.[19] Intuitively, if buyers arrive frequently, then the low-type seller has a strong incentive to wait for c_H. For her reservation price to stay equal to v_L, each buyer must offer c_H with a sufficiently small probability and, therefore, only when he receives the highest signal s_N (even then with probability less than 1). This implies that the stationary cutoff belief q^* is equal to $\overline{q}_N$, which in turn implies that for any n, the corresponding cutoff signal $\overline{q}_n$ is determined to be the value that satisfies

$$\frac{1-\overline{q}_n}{\overline{q}_n} = \frac{\gamma_H(s_n)}{\gamma_L(s_n)}\frac{v_H - c_H}{c_H - v_L}. \tag{6}$$

It also implies that the rate at which each type trades after $q(t)$ reaches q^* is given by

$$\rho_H = \frac{\gamma_H(s_N)}{\gamma_L(s_N)}\frac{v_L - c_L}{c_H - v_L} r, \tag{7}$$

which is also independent of λ.[20]

Second, buyers' beliefs immediately jump to q^* in the frictionless limit, as long as their initial beliefs are smaller than $\bar{q}_1$. The length of time it takes for buyers' beliefs to move from $\overline{q}_n$ to $\overline{q}_{n+1}$ (in the case of $n < n^*$) or $\overline{q}_{n-1}$ (in the case of $n \geq n^*$) shrinks to 0 as λ tends to infinity (see footnote 16). Therefore, the length of time for buyers' beliefs to move from $\widehat{q}$ to q^* also shrinks to 0. Intuitively, at each belief level, (expected) arrival of each buyer moves buyers' beliefs at a constant rate. Therefore, buyers' beliefs move arbitrarily fast as λ tends to infinity.

Let $F_a^*(\cdot|\widehat{q})$, $a = L, H$, represent the limit distributions of random variables $\tau_a(\widehat{q})$ as λ approaches infinity. The following proposition characterizes the distributions, in particular, the probability that each type trades immediately.

Proposition 4 *For each $\widehat{q} \in (0, \overline{q}_1)$ and $a \in \{L, H\}$, as λ tends to infinity, the probability that the type a seller trades by time t converges to*

$$F_a(t|\widehat{q}) = 1 - (1 - F_a(0|\widehat{q}))e^{-\rho_H t},$$

where ρ_H is given in (7) and

[19] The precise condition under which this is true is $\lambda > \frac{r(v_L - c_L)}{\gamma_L(s_N)(c_H - v_L)}$.

[20] If $n^* = N$, then the rate at which each type receives offer c_H is equal to $\rho_L = \lambda\gamma_L(s_N)\sigma_B^*$ and $\rho_H = \lambda\gamma_L(s_N)\sigma_B^*$, respectively. The result follows from the fact that ρ_L is necessarily equal to $\frac{v_L - c_L}{c_H - v_L} r$, and both types must have the same exit rate.

- *if* $\widehat{q} < q^*$*, then*

$$F_L(0|\widehat{q}) = 1 - \frac{\widehat{q}}{1-\widehat{q}} \frac{\gamma_H(s_N)}{\gamma_L(s_N)} \frac{v_H - c_H}{c_H - v_L} \quad \textit{and} \quad F_H(0|\widehat{q}) = 0; \tag{8}$$

- *if* $q^* < \widehat{q} < \bar{q}_1$*, then for each* $a = L, H$*,*

$$F_a(0|\widehat{q}) = 1 - \left(\frac{1}{l_n} \frac{c_H - v_L}{v_H - c_H} \frac{1-\hat{q}}{\hat{q}} \right)^{\psi_n^a} \times \prod_{i=n+1}^{N} \left(\frac{l_{i-1}}{l_i} \right)^{\psi_i^a}, \tag{9}$$

with $l_i = \frac{\gamma_H(s_i)}{\gamma_L(s_i)}$ *and* $\psi_i^a = \frac{1-\Gamma_a(s_i)}{\Gamma_H(s_i)-\Gamma_L(s_i)}$.

Proof. See the Appendix. ■

Intuitively, the atoms at time 0, $F_a(0|\widehat{q})$, reflect the probabilities with which each type trades before $q(t)$ reaches q^*. These probabilities, and thus their limits as well, are straightforward to calculate from the characterization in the previous section. To understand how the expressions in (8) correspond to these probabilities, first recall that if $\widehat{q} < q^*$, then, since $n^* = N$, the high-type seller trades with zero probability until buyers' beliefs reach q^*, which implies $F_H(0|\widehat{q}) = 0$. For this range of initial beliefs, the low type trades with probability 1 conditional on the arrival of a buyer. Given this, $1 - F_L(0|\widehat{q})$ is precisely the probability with which the low type should *not* trade so that buyers' beliefs jump from $\widehat{q}$ to q^*.[21] Similarly, in (9), the first multiplicative term is the probability with which the type-a seller does not trade before the buyers' beliefs move to $\overline{q}_n$, and $\left(\frac{l_{i-1}}{l_i}\right)^{\psi_i^a}$ is the corresponding probability for buyers' beliefs to evolve from $\overline{q}_i$ to $\overline{q}_{i-1}$.[22] It is important to note that all these probabilities are independent of λ; i.e. as λ varies, the equilibrium strategies adjust so that these probabilities remain constant.

We now turn to the reservation price schedule of the low type seller. In the limit, since buyers' beliefs immediately jump from $\widehat{q}$ to q^*, the low-type seller either trades immediately at price c_H (while buyers' beliefs converge to q^*) or receives expected payoff v_L (once buyers' beliefs become equal to q^*).[23] Therefore, $p(\widehat{q})$ is simply a weighted average of c_H and v_L, with the weight to c_H

[21]Precisely, due to (6),

$$\frac{\widehat{q}}{\widehat{q} + (1-\widehat{q})(1-F_L(0|\widehat{q}))} = \frac{1}{1 + \frac{\gamma_H(s_N)}{\gamma_L(s_N)} \frac{v_H - c_H}{c_H - v_L}} = q^*.$$

[22]Notice that $\psi_i^L - \psi_i^H = -1$. Therefore,

$$\frac{1-q_{i+1}}{q_{i+1}} = l_i \frac{v_H - c_H}{c_H - v_L} = \frac{1-q_i}{q_i} \left(\frac{l_i}{l_{i+1}} \right)^{\psi_i^L - \psi_i^H} = l_i \frac{v_H - c_H}{c_H - v_L} \frac{l_{i+1}}{l_i}.$$

[23]For the latter, we again invoke the fact that $p(q)$ can be calculated by assuming that the low type would accept only c_H.

being equal to the probability that she immediately trades at price c_H. Recall that if $\widehat{q} < q^*$, then no buyer offers c_H until buyers' beliefs reach q^* (recall that $n^* = N$), while if $\widehat{q} \in (q^*, \overline{q}_1)$, then the low type receives offer c_H with probability $F_L(0|\widehat{q})$ before buyers' beliefs reach q^*. The following result is then immediate from Proposition 4.

Proposition 5 *In the limit as λ tends to infinity, the low-type seller's expected payoff is equal to*

$$p(\widehat{q}) = \begin{cases} v_L, & \text{if } \widehat{q} \leq q^*, \\ v_L + F_L(0|\widehat{q})(c_H - v_L), & \text{if } \widehat{q} \in (q^*, \overline{q}_1), \\ c_H, & \text{if } \widehat{q} \geq \overline{q}_1. \end{cases} \quad (10)$$

6 Effects of Improving the Precision of Buyers' Signals

In this section, we study the effects of improving the precision of buyers' signals. Our main goal is not to perform comprehensive comparative statics analysis regarding information structures, but to obtain some insights about the effects of improving the quality of buyers' signals. To this end, we focus on the simplest case in which there are only two signals, that is, $N = 2$. In order to highlight the purely informational effects, we further restrict attention to the limit case where λ is arbitrarily large.[24]

Specifically, we examine how the market outcome varies as the likelihoods of the signals change. In particular, we focus on the effects of a marginal increase in l_2.[25] Such changes in the parameter values are consistent with various common notions of more precise signals.[26] In the same spirit as in Section 5, we focus on the effects of those changes on the expected time to trade (as a measure of market liquidity and efficiency) and the low-type seller's reservation price (as a measure of surplus division). Since the result is obvious if $\widehat{q} > \overline{q}_1$, we consider only the case where $\widehat{q} < \overline{q}_1$. In what follows, for notational simplicity, we simply say $\widehat{q} > q^*$, in order to refer to $\widehat{q} \in (q^*, \overline{q}_1)$.

Proposition 4 informs us of how to calculate the expected time to trade for each type. In the limit as λ tends to infinity, each type trades either immediately or at a constant rate of ρ_H thereafter. Therefore, the expected time to trade for the type-a seller is equal to the probability that the seller does not trade immediately ($1 - F_a(0|\widehat{q})$) times the expected duration when the hazard ratio is

[24] Although we present the results only for the limit case, all the qualitative results established in this section hold as long as λ is sufficiently large.

[25] We explain the effects of a marginal decrease in l_1 at the end of this section.

[26] There are several ways to rank information structures. The most common criteria are Blackwell's *garbling* (Blackwell (1951)), Lehmann's *accuracy* (Lehmann (1988)), and Shannon's entropy (Shannon (1948)). Even though it is not clear which concept is appropriate in strategic environments in general and in our model in particular, the variations we consider are consistent with each of them.

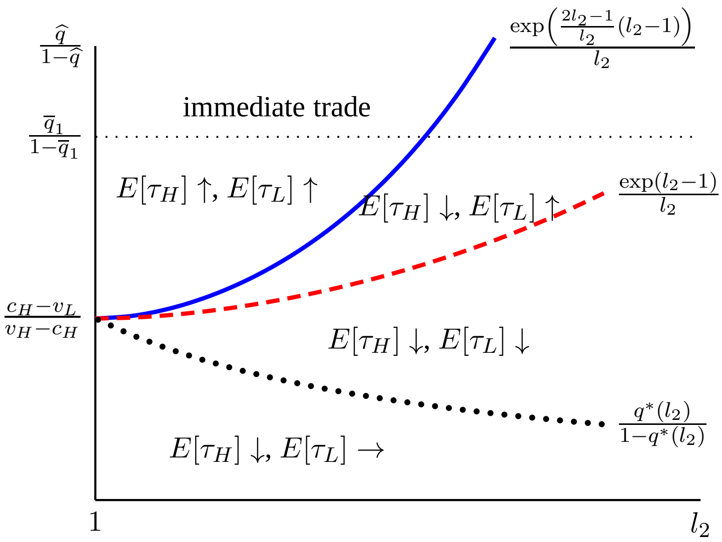

Figure 4: The effects of an increase in l_2 on the expected times to trade.

given by ρ_H. Formally, $E[\tau_a(\widehat{q})] = \frac{1-F_a(0|\widehat{q})}{\rho_H}$. The following proposition is then immediate by basic calculus.[27]

Proposition 6 *Suppose l_2 increases marginally.*

- *If $\widehat{q} < q^*$, then $E[\tau_L(\widehat{q})]$ remains constant, while $E[\tau_H(\widehat{q})]$ decreases.*
- *If $\widehat{q} > q^*$, then $E[\tau_L(\widehat{q})]$ decreases if and only if $l_2 > \frac{\exp(l_2-1)}{l_2}\frac{c_H-v_L}{v_H-c_H}$, while $E[\tau_H(\widehat{q})]$ decreases if and only if $l_2 > \frac{exp\left(\frac{2l_2-1}{l_2}(l_2-1)\right)}{l_2}\frac{c_H-v_L}{v_H-c_H}$.*

Figure 4 visualizes Proposition 1. It is clear that more transparency (increased precision of signals) may or may not contribute to market liquidity and efficiency. Nonetheless, there are three systematic patterns. First, more precise signals are less likely to be beneficial when $\widehat{q}$ is high; i.e. the range of l_2 at which an increase in l_2 speeds up trade decreases as $\widehat{q}$ increases. Second, more precise signals are more likely to be beneficial when l_2 is already high; i.e. the range of $\widehat{q}$ at which an increase in l_2 speeds up trade increases as l_2 increases. Finally, increased precision tends to speed up trade for the high type than for the low type.

[27] For the case of $N = 2$,

$$\rho_H = l_2\frac{v_L - c_L}{c_H - v_L}r.$$

In addition,

$$F_L(0|\widehat{q}) = \begin{cases} 1 - l_2\frac{v_H-c_H}{c_H-v_L}\frac{\widehat{q}}{1-\widehat{q}}, & \text{if } \widehat{q} \le q^*, \\ 1 - \left(\frac{1}{l_2}\frac{c_H-v_L}{v_H-c_H}\frac{1-\widehat{q}}{\widehat{q}}\right)^{\frac{1}{l_2-1}}, & \text{if } \widehat{q} > q^* \end{cases} \quad F_H(0|\widehat{q}) = \begin{cases} 0 & \text{if } \widehat{q} \le q^* \\ 1 - \left(\frac{1}{l_2}\frac{c_H-v_L}{v_H-c_H}\frac{1-\widehat{q}}{\widehat{q}}\right)^{\frac{l_2}{l_2-1}} & \text{if } \widehat{q} \in q^* \end{cases}.$$

To understand the first two patterns, notice that increased precision of signals affects the equilibrium dynamics through two channels. First, it directly reduces the risk of each buyer's paying a high price for a low-quality asset, thereby encouraging buyers to offer c_H more often. Ceteris paribus, this effect reduces the expected time to trade.[28] Second, since other buyers also receive more precise signals, it speeds up the evolution of buyers' beliefs. In particular, if $\widehat{q} > q^*$, for a fixed time-on-the-market, the higher l_2 is, the more pessimistic are buyers about the quality of the asset; i.e. for a fixed t, an increase in l_2 decreases $q(t)$. This indirectly reduces buyers' incentive to offer a high price, thereby slowing down trade.[29] The patterns emerge because the magnitude of the former effect is essentially constant in $\widehat{q}$ and l_2, while that of the latter increases in $\widehat{q}$ and decreases in l_2. Buyers' beliefs travel from $\widehat{q}$ to q^*. Therefore, the latter effect amplifies as $\widehat{q}$ increases. On the other hand, when l_2 is already high, an additional increase of l_2 has a small effect on the speed of the belief convergence, and thus the latter effect is relatively small.

To understand why increased precision tends to speed up trade for the high type more than for the low type, recall that, while the high type trades only at c_H, the low type also trades at her reservation price $p(t)$. Therefore, the low-type seller's expected time to trade also depends on her incentive to accept $p(t)$, which in turn depends on the rate at which buyers offer c_H. This means that there is a countervailing effect to the first one in the previous paragraph: increased precision increases the low type's chance to trade at c_H. But, this decreases the low type's incentive to accept $p(t)$ and, therefore, slows down trade. Clearly, this additional effect operates only for the low type. It follows that the low type's expected time to trade decreases only when the high type's also decreases, while the high type's can decrease even when the low type's increases. The countervailing effect is particularly strong when $\widehat{q}$ is smaller than q^*. In that case, it fully offsets the direct effect, and thus the expected time to trade stays constant.[30]

Remark 1 (Distribution of time to trade) Although we have focused on how the expected times to trade respond to the change of l_2, the effects on the entire distributions of time to trade are straightforward to obtain. Recall that the distribution for each type consists of two components: the probability that trade takes place immediately ($F_a(0|\widehat{q})$), and the stationary rate of trade conditional on no trade at time 0 (ρ_H). The latter always increases in l_2, while the former may or may not increase. Therefore, if the former also increases, then the distribution decreases in l_2 in the first order stochastic dominance sense, while if the former decreases, then the change of the distribution cannot be clearly ranked. Since first-order stochastic dominance implies an increase of the expected value, it follows that the region at which an increase in l_2 speeds up trade in the

[28]In the formal expression of $E[\tau_a(\widehat{q})]$, this effect manifests itself as a decrease in $q^* = l_2 \frac{v_H - c_H}{c_H - v_L}$ and an increase in $\rho_H = \frac{r(v_L - c_L)}{c_H - v_L}$.

[29]In the formal expression, this effect is present in the power terms, $\frac{1}{l_2 - 1}$ and $\frac{l_2}{l_2 - 1}$.

[30]When $\widehat{q} < q^*$, buyers' beliefs increase over time. Therefore, the second effect in the previous paragraph is absent.

sense of first-order stochastic dominance is smaller than the region at which the expected time to trade decreases in l_2.

Our next result concerns how the low-type seller's expected payoff is affected by an increase in l_2. Proposition 5 implies that if $\widehat{q} < q^*$, then it is constant in l_2, while if $\widehat{q} > q^*$, then it depends on how $F_L(0|\widehat{q})$ responds to an increase in l_2.[31] The following result is then straightforward to obtain from Proposition 4.

Proposition 7 *Suppose l_2 increases marginally.*

- *If $\widehat{q} < q^*$, then $p(\widehat{q})$ stays constant.*
- *If $\widehat{q} \geq q^*$, then $p(\widehat{q})$ increases if and only if $\frac{\widehat{q}}{1-\widehat{q}} < \frac{\exp\left(\frac{l_2-1}{l_2}\right)}{l_2} \frac{c_H - v_L}{v_H - c_H}$.*

Remark 2 (Impact of a marginal decrease in l_1) So far we have considered only an increase in l_2, since ρ_H and $F_a(t;\widehat{q})$ are independent of the other likelihood ratio l_1. Yet, decreasing l_1 (making signal s_1 more informative) is another way to improve the quality of buyers' signals. The only role that l_1 plays is to change the lower bound $\overline{q}_1$ on $\widehat{q}$, above which trade is immediate. It is easy to see that a decrease in l_1 leads to an increase in $\overline{q}_1$. Intuitively, this is because a decrease in l_1 means that signal s_1 becomes an even worse signal, and thus for a buyer to be willing to offer c_H with signal s_1, his prior must be even higher. It is immediate that if $\widehat{q} = \overline{q}_1$, then a marginal decrease in l_1 sharply slows down trade of both types: Before the change, both types trade upon arrival of the first buyer, while after the change, there is substantial delay.

7 Conclusion

The main contribution of our paper is to provide a simple and intuitive framework, which, nevertheless, leads to a rich set of predictions for equilibrium trading dynamics. Within this framework we are able to identify different sources of trading delay and provide an understanding of how these sources interact. The simple comparative statics exercise we present in Section 6 demonstrates how this model can be used to identify the role of "asset transparency" which has recently been the target of market regulations. The simple yet rich structure of the equilibrium of our model easily lends itself to such policy analysis. Moreover, we believe that its further modifications may help shed light on other issues of interest such as increased transparency of market transactions and various market regulations.

[31] Notice that the condition under which $p(\widehat{q})$ increases coincides with the condition for the distribution $F_L(\cdot;\widehat{q})$ to decrease in the first-order stochastic dominance.

Appendix: Omitted Proofs

Proof of Lemma 2:. We establish the result in three steps.

(1) If $q(t) < q^*$, then $p(t) < v_L$.

Suppose $q(t) < q^*$, but $p(t) > v_L$. Then, there must exists $t' \in (t, \infty)$ such that $p(t') = v_L$. Suppose not, that is, $p(t') > v_L$ for any $t' > t$. Lemma 1 implies that $q(\cdot)$ then keeps decreasing. This, in turn, implies that there must exist $q_\infty \in [0, q(t))$ such that $q(\cdot)$ converges to q_∞. Recall that, since $p(t') > v_L$ for any t, both types trade only when the buyer offers c_H. Therefore, in the long run, both types must trade at the same rate, which can be the case only when either every buyer always offers c_H or every buyer never offers c_H. Since $q(t) < q^*$, the former obviously cannot be true. The latter also cannot be the case, because if so, the low-type seller's reservation price would be close to c_L, which is strictly smaller than v_L.

Let t' be the smallest value such that $p(t') = v_L$. Then, for any $x \in (t, t')$, $p(x) > v_L$. Therefore, by Lemma 1, $q(x) \leq q(t) < q^*$, which implies that the probability that the buyer at $x \in (t, t')$ offers c_H is strictly less than $\gamma_L(s^*)\sigma_B^* + 1 - \Gamma_L(s^*)$. Combining this with $p(t') = v_L$, it follows that $p(t) < v_L$, which is a contradiction (recall that if all the buyers between t and t' offer c_H with probability $\gamma_L(s^*)\sigma_B^* + 1 - \Gamma_L(s^*)$ and $p(t') = v_L$, then $p(t) = v_L$).

Now suppose $q(t) < q^*$, but $p(t) = v_L$. Together, they imply that the buyer at t offers c_H with a strictly lower probability than $\gamma_L(s^*)\sigma_B^* + 1 - \Gamma_L(s^*)$. If $\dot{p}(t) \leq 0$, then clearly $p(t) < v_L$, which is a contradiction. If $\dot{p}(t) > 0$, there exists t' such that $q(t') < q^*$, but $p(t') > v_L$. We showed above that this can never be the case.

(2) If $q(t) > q^*$, then $p(t) > v_L$.

Suppose $q(t) > q^*$, but $p(t) < v_L$. We first show that there exists $t' \in (t, \infty)$ such that $p(t') = v_L$. Suppose not, that is, $p(t') < v_L$ for any $t' \geq t$. Lemma 1 implies that $q(\cdot)$ keeps increasing. Since $q(t) \in [0, 1]$ for any t, this means that there exists $q^\infty \in (q(t', 1]$ such that $q(\cdot)$ converges to q^∞. Since the low type trades whenever a buyer arrives (see Lemma **??**), the convergence can occur only when the high type trades with almost probability 1. This, in turn, implies that in the long run, each buyer offers c_H with probability 1, regardless of his signal. But, then the low-type seller's reservation price becomes arbitrarily close to c_H. This is a contradiction, because c_H is strictly larger than v_L by Assumption 1.

Let t' be the smallest value at which $p(t') = v_L$. Since $p(x) < v_L$ for any $x \in (t, t')$, $q(\cdot)$ cannot decrease on (t, t'). Therefore, $q(x) > q^*$ for any $x \in (t, t')$. Let $t'' \equiv t' - \epsilon$ for ϵ positive, but sufficiently small. Then, for any $x \in (t'', t')$, the buyer must offer c_H with probability 1 whenever his signal is weakly above s^*: Since x is close to t', $p(x)$ is close to v_L. Therefore, when the buyer's signal is s^*, his expected payoff by offering $p(x)$ is also close to 0. To the contrary, his expected payoff by offering c_H is bounded away from 0, because $q(x) \geq q(t) > q^*$ (recall that the payoff is equal to 0 if $q(x) = q^*$). But, then $p(x) > v_L$, because the buyers on (x, t') offer c_H at least with probability $1 - \Gamma_L^-(s^*)$, while the low-type seller's reservation price at t' is equal to v_L (recall that the low-type seller's reservation price is equal to 0 if every buyer offers c_H with probability $\gamma_L(s^*)\sigma_B^* + 1 - \Gamma_L(s^*)$). This is, of course, a contradiction.

Now suppose $q(t) > q^*$, but $p(t) = v_L$. In this case, the low type does not necessarily accept $p(t)$ with probability 1. Therefore, $q(\cdot)$ is not necessarily increasing. However, we do know that the buyer would offer c_H with probability 1 whenever his signal is weakly above s^*: Since $p(t) = v_L$, the buyer with signal s^* obtains zero expected payoff by offering $p(t)$, while his expected payoff by offering c_H is strictly positive, because $q(t) > q^*$. If $\dot{p}(t) \geq 0$, then it is clear that $p(t) > v_L$. If $\dot{p}(t) < 0$, then there exists $t' > t$ such that $q(t') > q^*$, but $p(t') < v_L$. We showed above that this can never be the case.

(3) If $q(t) = q^*$, then $p(t) = v_L$.

That $p(t) = v_L$ implies $q(t) = q^*$ is immediate. Now suppose $q(t) = q^*$ but $p(t) < v_L$. Then the belief is increasing so that for $t' > t$ sufficiently close to t, $q(x) > q^*$ and $p(x) < v_L$ whenever $x \in (t, t')$, since both $p(\cdot)$ and $q(\cdot)$ are continuous, a contradiction. Symmetric arguments lead to a contradiction for the case of $p(t) > v_L$. ■

Proof of Lemma 3:. Suppose there exists t such that $q(t) < q^*$, but $\dot{p}(t) \leq 0$. Since $p(\cdot)$ is continuous and eventually converges to v_L, there exists t' such that $t' > t$ and $p(t') = p(t)$. Without loss of generality, assume that $p(x) \leq p(t)$ for any $x \in (t, t')$ and $\dot{p}(x) > 0$ for any $x > t'$ such that $q(x) < q^*$ (if $p(\cdot)$ is not strictly increasing until it reaches v_L, there always exist t and t' that satisfy these properties). For $x \in (t, t')$, $p(x) \leq p(t')$, while $q(x) < q(t')$. This implies that the probability that the buyer at $x \in (t, t')$ offers c_H cannot be larger than the corresponding probability for the buyer at t'. To the contrary, whenever $x > t'$, $p(x) > p(t')$ and $q(x) > q(t')$. Therefore, the probability that the buyer at $x > t'$ offers c_H is strictly larger than the corresponding probability for the buyer at t'. Since the low-type seller's reservation price $p(\cdot)$ is determined by the rate at which buyers offer c_H, it immediately follows that $p(t) < p(t')$, which is a contradiction. ■

Proof of Proposition 4:. First consider the case where $\widehat{q} < q^*$. In that case, for any fixed λ, the H-type seller does not trade until the belief reaches q^*. Therefore, $F_H(0|\widehat{q}) = 0$. The L-type, on the other hand, trades with probability 1 conditional on the arrival of a buyer during this period. Then, the probability he trades before the stationary path is reached is given by $1 - e^{-\lambda T(\widehat{q}, q^*)}$ where

$$T(\widehat{q}, q^*) = -\frac{1}{\lambda} \log\left(\frac{1 - q^*}{q^*} \frac{\widehat{q}}{1 - \widehat{q}}\right).$$

Combined with (6), this implies that for $\widehat{q} < q^*$,

$$F_L(0|\widehat{q}) = 1 - \frac{\widehat{q}}{1-\widehat{q}} \frac{\gamma_H(s_N)}{\gamma_L(s_N)} \frac{v_H - c_H}{c_H - v_L}.$$

Next, consider the case where $\widehat{q} \in (q^*, \bar{q}_1)$. In this case, for a given λ, conditional on the arrival of a buyer, each type trades if and only if the offer is c_H, which happens with probability $1 - \Gamma_a^-(s_n) = 1 - \Gamma_a(s_{n-1})$ if $q \in [\bar{q}_n, \bar{q}_{n-1})$. Therefore, for type a, the probability that the trade takes place before the belief reaches q^* is given by

$$1 - \exp\left\{ -(1-\Gamma_a(s_{n-1}))T(\hat{q}, \overline{q}_n) - \sum_{i=n+1}^{N} (1-\Gamma_a(s_{i-1}))T(\overline{q}_{i-1}, \overline{q}_i) \right\},$$

so that

$$-(1-\Gamma_a(s_{n-1}))T(\hat{q}, \overline{q}_n) = \log \left(\frac{1-\bar{q}_n}{\bar{q}_n} \frac{\hat{q}}{1-\hat{q}} \right)^{-\frac{1-\Gamma_a(s_{n-1})}{\Gamma_H(s_{n-1})-\Gamma_L(s_{n-1})}},$$

and

$$(1-\Gamma_a(s_{i-1}))T(\overline{q}_{i-1}, \overline{q}_i) = \log \left(\frac{1-\bar{q}_i}{\bar{q}_i} \frac{\overline{q}_{i-1}}{1-\overline{q}_{i-1}} \right)^{-\frac{1-\Gamma_a(s_i)}{\Gamma_H(s_i)-\Gamma_L(s_i)}}.$$

Then, the result follows immediately using (6). ■

References

Akerlof, G.A.(1970), ‘The market for “Lemons”: quality uncertainty and the market mechanism,’ *Quarterly Journal of Economics* **84**(3), 488–500.

Benerjee, A. (1992), ‘A simple model of heard behavior,’ *Quarterly Journal of Economics* **107**(3), 797–817.

Bikhchandani, S, Hirshleifer, D. and Welch, I.(1992), ‘A theory of fads, fashion, custom, and cultural change as informational cascades,’ *Journal of Political Economy* **100**(5), 992–1026.

Blackwell, D. (1951), ‘The comparison of experiments,’ in *Proceedings, Second Berkeley Symposium on Mathematical Statistics and Probability*, Berkeley, CA: University of California Press, 93-102.

Camargo, B. and B. Lester (2011), ‘Trading dynamics in decentralized markets with adverse selection,’ mimeo.

Chang, B. (2010), ‘Adverse selection and liquidity distortion in decentralized markets,’ mimeo.

Chari, V.V., A. Shourideh, and A. Zetlin-Jones (2010), ‘Adverse selection, reputation and sudden collapses in secondary loan markets,’ mimeo.

Chiu, J. and T.V. Koeppl (2011), ‘Market freeze and recovery: trading dynamics under optimal intervention by a market-maker-of-last-resort,’ mimeo.

Choi, M. (2013), ‘Asset prices and information transmission,’ mimeo.

Daley, B. and B. Green (2011), ‘Waiting for news in the dynamic market for lemons,’ *Econometrica*, **80**(4), 1433–1504.

Deneckere, R. and M.-Y. Liang (2006), ‘Bargaining with interdependent values,’ *Econometrica* **74**(5), 1309–1364.

Diamond, P.A. (1971), ‘A model of price adjustment,’ *Journal of Economic Theory* **3**, 156–168.

Evans, R. (1989), ‘Sequential bargaining with correlated values,’ *Review of Economic Studies* **56**(4), 499–510.

Fuchs, W., Öry, A. and Skrzypacz, A. (2012), ‘Transparency and distressed sales under asymmetric information,’ *mimeo*.

Guerrieri, V. and R. Shimer (2012), 'Dynamic adverse selection: a theory of illiquidity, fire sales, and flight to quality,' mimeo.

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

Janssen, M.C.W. and S. Roy (2002), 'Dynamic trading in a durable good market with asymmetric information,' *International Economic Review* **43**(1), 257-282.

Janssen, M.C.W. and S. Roy (2004), 'On durable goods markets with entry and adverse selection,' *Canadian Journal of Economics* **37**(3), 552-589.

Kaya, A. and Q. Liu (2012), 'Transparency and price formation,' mimeo.

Kim, K. (2012), 'Information about sellers' past behavior in the market for lemons,' mimeo.

Kurlat, P. (2010), 'Lemons, market shutdowns and learning,' mimeo.

Lauermann, S. and A. Wolinsky (2006), 'Search with adverse selection,' mimeo.

Lehmann, E. (1988), 'Comparing location experiments,' *The Annals of Statistics* **16**(2), 521–533.

Moreno, D. and J. Wooders (2010), 'Decentralized trade mitigates the lemons problem,' *International Economic Review* **51**(2), 383–399.

Moreno, D. and J. Wooders (2012), 'Dynamic markets for lemons: performance, liquidity, and policy evaluation,' mimeo.

Nöldeke, G. and E. Van Damme (1990), 'Signalling in a dynamic labor market,' *Review of Economic Studies* **57**, 1-23.

Roy, S. (2012), 'Dynamic sorting in durable goods markets with buyer heterogeneity,' mimeo.

Shannon, C.E. (1948), 'A mathematical theory of communication,' *Bell System Technical Journal* **27**(3), 379–423.

G.J. Stigler (1961), 'The economics of information,' *The Journal of Political Economy* **69**(3), 213–225.

Swinkels, J.M. (1999), 'Education signalling with preemptive offers,' *Review of Economic Studies* **66**, 949-970.

Taylor, C. R. (1999), 'Time-on-the-market as a sign of quality,' *Review of Economic Studies* **66**, 555-578.

Vincent, D. R. (1989), 'Bargaining with common values,' *Journal of Economic Theory* **48**(1), 47–62.

Vincent, D. R. (1990), 'Dynamic auctions,' *Review of Economic Studies* **57**(1), 49–61.

Zhu, H. (2012), 'Finding a good price in opaque over-the-Counter markets,' *Review of Financial Studies* **25**(4), 1255–1285.