Trading Dynamics in the Market for Lemons *

Ayça Kaya† and Kyungmin Kim‡

May 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. By encompassing various patterns of trading dynamics, our model broadens the applicability of dynamic adverse selection theory. We also demonstrate 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 up for sale for a while, what would potential buyers infer about its quality? The answer depends on the perceived source of the delay. First, it could simply be that no potential buyer has shown up yet. 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 may have observed an unfavorable attribute of the asset and, therefore, decided not to purchase it. In this case, 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, any one of them

*We are grateful to Michael Choi, William Fuchs, Martin Gervais, Stephan Lauermann, Benjamin Lester, Qingmin Liu, Santanu Roy, Galina Vereshchagina, and Gabor Virag for many helpful comments and suggestions.

†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 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 Banerjee (1992) and Bikhchandani et al. (1992). However, the existing literature is largely silent on how these sources are linked to one another and to 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. The goal of this paper is to clarify the interplay among various sources of delay, thereby providing 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 or of high quality. 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. This allows us to explicitly study trading dynamics that arises under adverse selection.

We show that there is generically a unique equilibrium in this dynamic trading problem, and equilibrium trading dynamics crucially depends on the asset's initial reputation (i.e., buyers' prior beliefs about the quality of the asset). We find that for an asset enjoying a high reputation, delay typically results from unfavorable inspection outcomes, and the asset's reputation deteriorates over time. In contrast, if the asset starts out with a low reputation, then delay stems mainly from the seller's rejecting low prices, and the asset's reputation improves over time. To understand these patterns, first note that the higher the asset's reputation is, the more likely are buyers to offer a high price. This implies that even a seller of a low-quality asset would be reluctant to accept a low price while enjoying a high reputation. In this case, trade can be delayed only when, despite the asset's high reputation, buyers are unwilling to offer a high price, which is the case when they receive unfavorable inspection outcomes. Since a low-quality asset is more likely to generate such inspection outcomes, the asset is deemed less likely to be of high quality, the longer it stays on the market. In the opposite case when the asset for sale suffers from a low reputation, a seller of a low-quality asset would be willing to settle for a low price, while that of a high-quality asset would still insist on a high price. Since a high-quality asset would stay on the market relatively longer than a low-quality asset, the asset's reputation improves over time. Interestingly, buyers' beliefs converge to a certain level, regardless of whether the asset's initial reputation is high or low.

[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

This is the level at which the two effects, a low-quality asset's generating unfavorable inspection outcomes more often vs. a high-quality seller's insisting on a high price, are exactly balanced.

The role of search frictions in determining equilibrium trading dynamics deserves elaboration. First, since search frictions have the same impact on both types, they are neutral to the *direction* of the evolution of beliefs. 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 contacted the seller yet. This forces buyers' beliefs to change gradually. Second, search frictions are responsible only for a portion of the delay: even if search frictions are arbitrarily small, the expected time to trade remains bounded away from zero. This may be surprising in light of the fact that each buyer generates a constant amount of information by inspecting the asset, and therefore in the frictionless limit an arbitrarily large amount of *exogenous* information is instantaneously generated about the quality of the asset. Yet, strategic interactions between the seller and the buyers preclude efficient use of this information, causing delay to persist.

The richness of our equilibrium dynamics broadens the applicability of the theory of dynamic adverse selection. The literature on dynamic adverse selection 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. A common theoretical insight of this literature is that, since a seller enjoys a higher payoff from holding on to a high-quality asset than a low-quality asset, the latter trades faster than the former and thus an asset's reputation improves over time. This is precisely the main working mechanism in our model when the asset's initial reputation is low. Our contribution is to enrich the model of dynamic adverse selection by allowing for buyer inspection, thereby accommodating the possibility that a high-quality asset trades faster than a low-quality asset, and thus an asset's reputation may deteriorate over time. Indeed, casual observation suggests that in several markets, including housing and labor markets, the longer an asset stays on the market, the lower is its perceived quality. Although it might be controversial whether information asymmetries are indeed present in such markets, the ability of our model to generate such dynamics enlarges the set of empirical patterns that can be accommodated within models of dynamic adverse selection.

Our model, with its unique equilibrium and clean characterization, is particularly suited to evaluating government policies that influence the market structure.[3] Within our framework, probably

[2] A non-exhaustive list includes Camargo and Lester (2011), Chang (2012), Chari et al. (2012), Chi and Koeppl (2011), Choi (2013), Daley and Green (2012), Guerrieri and Shimer (2012), Fuchs and Skrzypacz (2013a,b), Fuchs et al. (2012), Hörner and Vieille (2009), Kurlat (2010), Lauermann and Wolinsky (2013), Moreno and Wooders (2012), Roy (2013), and Zhu (2012).

[3] For instance, the government may mitigate search frictions by centralizing matching and/or trading mechanisms.

the most intriguing exercise is to explore the impact of the informativeness of buyers' signals (i.e., the quality of buyers' inspection technology) on equilibrium outcomes. 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, both of which include provisions for stricter disclosure requirements on the part of 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 in the precision of buyers' signals can increase or decrease the expected time to trade for each type, depending on the asset's initial reputation. More precisely, if an asset has a low initial reputation, 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 in the informativeness influences buyers' inferences from delay as well as their own signal. In particular, when the initial reputation of the asset is high, delay results from unfavorable inspection outcomes and, therefore, becomes an even stronger indication of low quality as the precision of buyers' signals increases. This reduces buyers' incentive to offer a high price. Even though increased precision of *own* signals directly reduces buyers' uncertainty about the quality of the asset and thus speeds up trade, the aforementioned indirect effect works in the opposite direction and may be significant. The interplay of these two forces makes the overall effect ambiguous.

Related Literature

As already discussed, most existing studies on dynamic adverse selection focus on the the implications of the difference in different types' reservation values, and, therefore, feature one form of equilibrium dynamics in which low quality assets trade faster. 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 that 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 a high quality asset in the second period than in the first period (that is, buyers' beliefs decrease over time). The logic

It may provide subsidies to trading parties, so as to alleviate their incentive problems. Our complete and neat characterization makes it fairly straightforward to analyze the effects of those and related policies.

[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.

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 (i.e., one in which 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 as complete as ours.

Two papers consider an environment similar to ours. Lauermann and Wolinsky (2013) investigate the ability of prices to aggregate dispersed information in a setting where, just like in our model, an informed player (buyer in their model) faces an infinite sequence of uninformed players, each of whom receives a noisy signal about the informed player's type. Using a similar model with the additional feature that the informed player can contact only a finite number of uninformed players, Zhu (2012) makes a number of interesting observations regarding trading in opaque over-the-counter markets.[6] In contrast to our model, in both studies, uninformed players have no access to the informed player's trading history. In particular, uninformed players do not observe the informed player's time-on-the-market. This induces buyers' beliefs and strategies to be necessarily stationary.[7] In other words, buyers' beliefs do not evolve over time. To the contrary, the evolution of uninformed players' beliefs and the resulting trading dynamics are the main focus of this paper.

Daley and Green (2012) study the role of the 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. Therefore, buyers do not face any inference problem regarding other buyers' signals. This makes their trading dynamics quite distinct from ours.[8] 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. In particular, the negative effect of increased informativeness stems from buyers' inferences about other buyers' signals in our model, while in Daley and Green (2012), it is due to its impact on the incentive of a high-quality seller to wait for good news.

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' equilibrium strategies and beliefs in Section

[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 induces the high type to always trade faster than the low type.

[6] One novelty in Zhu's model is that the seller can revisit buyers. Naturally, the seller revisits a buyer only after she has contacted all other buyers: there is no search cost in Zhu's model. 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. This new type of inference problem is another important ingredient of Zhu's analysis.

[7] This creates a non-trivial inference problem on the part of uninformed players, as their actual beliefs do not have to coincide with their prior beliefs: Uninformed players must take into account "contact-order uncertainty", that they do not know how many other uninformed players the informed player has met before, as well as the fact that different types of the informed player trade at different rates.

[8] Daley and Green (2012) consider a model without search frictions: the seller continuously receives price quotes from a pool of competitive buyers. Yet, the differences in equilibrium trading dynamics persist even in the limit as search frictions disappear in our model. See Section 5 for our frictionless market outcome.

3. The formal equilibrium characterization and the proof of uniqueness are in Section 4. We then present the limit equilibrium outcome as search frictions vanish in Section 5 and investigate the effects of improving the informativeness of buyers' signals in Section 6. We discuss two technical issues regarding equilibrium characterization in Section 7 and then conclude in Section 8. 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 payoffs at rate $r > 0$.

The asset is either of low quality (L) or of high quality (H). If the asset is of 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 of the seller. It is commonly known that the asset is of high quality with probability $\widehat{q}$ at time 0.

Upon arrival, each buyer receives a private signal about the quality of the asset. A signal s is drawn from the set $S = \{s_1, ..., s_N\}$.[9] 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 $\gamma_H(s_n)/\gamma_L(s_n)$ is strictly increasing in n, so that the higher a signal is, the more likely it is that the asset is of high quality. For later use, denote by $\Gamma_a(s_n)$ (respectively, $\Gamma_a^-(s_n)$) the probability that a buyer receives a signal *weakly* (respectively, *strictly*) below s_n from the type-a seller. Formally, for each $a = L, H$, $\Gamma_a(s_n) \equiv \sum_{n' \leq n} \gamma_a(s_{n'})$, and $\Gamma_a^-(s_n) \equiv \sum_{n' < n} \gamma_a(s_{n'})$.

We assume that buyers observe (only) how long the asset has been up for sale (i.e. time t).[10] This is consistent with the arrangement in the housing market and also seems to be plausible in various 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. This information

[9] We discuss the case with a continuum of signals in Section 7.

[10] It is well-known that the information buyers have about past histories of the game plays a crucial role in this type of game. See Nöldeke and van Damme (1990), Swinkels (1999), Hörner and Vieille (2009), Kim (2012), Kaya and Liu (2013), and Fuchs et al. (2012). We note that in our model, buyers observe neither the number of buyers who have arrived before nor the offers they have made to the seller.

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. This implies that buyers' beliefs at any point in the play of the game can be obtained through Bayesian updating.[11]

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 the type-a seller accepts price p at time t.[12] 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 $v_a - p$. All other buyers obtain zero payoff.[13]

We adopt the perfect Bayesian equilibrium concept. The concept requires 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. In other words, $q(t)$ is the belief held by the buyer who arrives at time t *prior to* his inspection. 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 the type-a seller's continuation payoff at time t, and (3) given σ_S and σ_B, for any t, $q(t)$ is obtained through Bayesian updating.

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

Assumption 1

$$v_L - c_L < \frac{\lambda}{r+\lambda}(c_H - c_L) \;\Leftrightarrow\; r(v_L - c_L) < \lambda(c_H - v_L).$$

The assumption states that the low-type seller has no incentive to accept buyers' maximal willingness-to-pay for the low-quality asset (v_L), should she expect to receive at least the high-type seller's reservation value (c_H) from the next buyer. If this assumption is violated, then the seller types can be easily separated and the resulting trading dynamics would be rather trivial. We note that $c_H > v_L$ is necessary for this assumption to hold, but a stronger assumption commonly adopted in

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

[12]In principle, the seller's strategy could depend on her private history, such as the number of previous buyers and rejected offers. This would enlarge the set of equilibria by allowing the seller to use her private history as a randomization device. However, all such equilibria would be essentially equivalent to the one in which the seller's strategy is independent of her private history, because buyers' strategies depend only on the seller's observable characteristic (time-on-the-market).

[13]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.

the literature, $\widehat{q}v_H + (1-\widehat{q})v_L < c_H$, is not.

3 Preliminary Observations

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

3.1 Seller's Optimal Acceptance Strategy

Since the seller's trading history 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, and in equilibrium, she accepts c_H with probability 1. In what follows, we denote by $p(t)$ the low-type seller's reservation price at time t. Since $r(p(t)-c_L) \leq \lambda(c_H - p(t))$, $p(t)$ 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, by definition, indifferent between accepting and rejecting $p(t)$. However, in equilibrium, if $p(t) < v_L$, then she necessarily accepts $p(t)$ 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 by the low type with probability 1, 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. In contrast, if $p(t) > v_L$, and it is offered in equilibrium, then the low-type seller must reject it with probability 1. This is because, otherwise, the buyer who offers $p(t)$ would receive a strictly negative payoff and, therefore, would strictly prefer offering a price below $p(t)$.

3.2 Buyers' Optimal Offer Strategies

Without loss of generality, we assume that each buyer offers either the reservation price of the low type (i.e., $p(t)$) or that of the high type (i.e., c_H). This assumption incurs no loss of generality for the following reasons. First, it is strictly suboptimal for a buyer to offer a price 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.

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—i.e., his belief upon receiving a signal—is strictly increasing in the signal received.

3.3 Evolution of Beliefs

The previous results have an important implication for 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 trading rate is equal to the arrival rate of buyers, λ. On the other hand, the high-type seller accepts only c_H. Therefore, her trading 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 (i.e., 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 (i.e., those above a cutoff signal) more frequently than the low type. Therefore, the trading 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 in buyers' beliefs are strict unless buyers' strategies are to offer c_H regardless of their signal.*

4 Equilibrium Characterization

Equipped with the observations from 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.

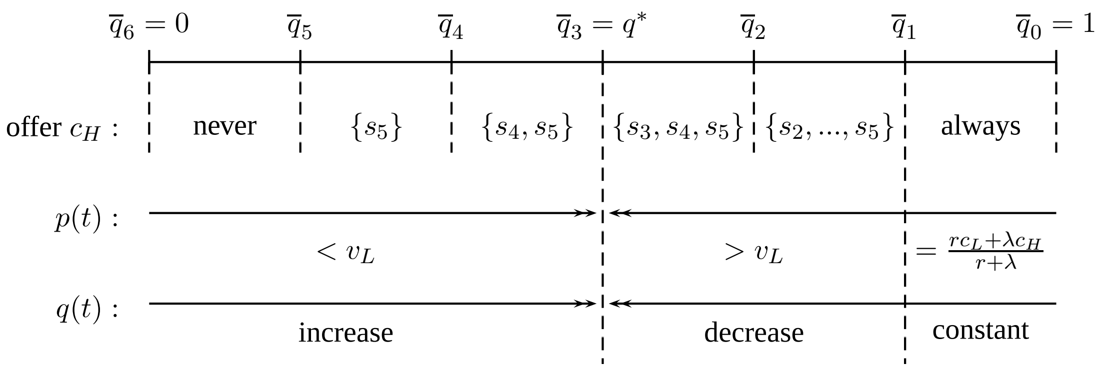

Figure 1: Equilibrium structure when there are 5 signals ($N = 5$) and $q^* = \overline{q}_3$.

(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.[14]

(iii) There exists $q^* \in (0, 1)$ to which buyers' beliefs conditional on no trade converge monotonically and in finite time.

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 show that the constructed equilibrium necessarily satisfies (iii). Our construction makes it clear that for a generic set of parameter values, there is a unique equilibrium that satisfies these conditions. We then show that any equilibrium must satisfy (i) and (ii), and thus the constructed equilibrium is the unique equilibrium in this game.[15]

4.1 Equilibrium Construction and Trading Dynamics

In this section we construct an equilibrium that satisfies properties (i) and (ii) and illustrate the resulting trading dynamics.

[14]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.

[15]Although the properties are fairly intuitive, a priori, it is not clear why an equilibrium that violates these properties cannot exist. 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 only with 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.

4.1.1 Unique Stationary Path

We first establish that there exists a unique *non-trivial* path along which buyers' beliefs can remain constant at some level q^*.[16] We identify this critical level q^* and characterize the unique strategy profile that supports such a stationary path.

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 receives offer c_H when buyers' beliefs remain constant at q^*. Then, ρ_L is uniquely pinned down by

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

to guarantee that indeed $p(q^*) = v_L$.

Suppose all buyers employ an identical pure strategy of offering c_H if and only if their signal is strictly above s_n. Then, the low-type seller receives offer c_H at rate $\lambda(1 - \Gamma_L(s_n))$. Notice that generically, there is no n such that $\rho_L = \lambda(1 - \Gamma_L(s_n))$. In other words, for a generic set of parameter values, the low-type seller's reservation price cannot be equal to v_L if all buyers adopt an identical pure offer strategy. In what follows, we focus on the generic case where $\rho_L \neq \lambda(1 - \Gamma_L(s_n))$ for any $n = 1, ..., N$, relegating the discussion of the non-generic cases to Section 7.

To pin down buyers' equilibrium offer strategies, let n^* be the unique integer such that

$$\lambda(1 - \Gamma_L(s_{n^*})) < \rho_L < \lambda(1 - \Gamma_L(s_{n^*-1})). \tag{1}$$

In other words, if buyers' strategies are to offer c_H if and only if their signal is *weakly* above s_{n^*}, then the low-type seller's reservation price exceeds v_L; while if buyers' strategies are to offer c_H if and only if their signal is *strictly* above s_{n^*}, then the low-type seller's reservation price falls short of v_L. Assumption 1 ensures that n^* is well-defined. In order to make the low-type seller's reservation price exactly equal to v_L, let σ_B^* be the value that satisfies

$$\rho_L = \lambda(\gamma_L(n^*)\sigma_B^* + 1 - \Gamma_L(s_{n^*})). \tag{2}$$

Then, 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^*}. For the generic case we are considering, σ_B^* always lies in $(0, 1)$.

The identification of s_{n^*} allows us to uniquely pin down the value of q^* using the optimality of buyers' strategies. Consider a buyer who has prior belief q^* and receives signal s_{n^*}. By Bayes'

[16] A *trivial* case is when $\widehat{q} > \overline{q}_1$. In this case, trade occurs with probability 1 upon arrival of a buyer, regardless of the seller's type, and thus buyers' beliefs obviously do not change.

rule, his belief updates to

$$\frac{q^*\gamma_H(s_{n^*})}{q^*\gamma_H(s_{n^*}) + (1-q^*)\gamma_L(s_{n^*})}.$$

At this belief, since $\sigma_B^* \in (0,1)$, the buyer must be indifferent between offering c_H and $p(q^*)$. But $p(q^*) = v_L$, and thus, he must receive zero expected payoff, regardless of the low-type seller's acceptance strategy. It follows that q^* is uniquely determined by the following equation:

$$q^*\gamma_H(s_{n^*})(v_H - c_H) + (1-q^*)\gamma_L(s_{n^*})(v_L - c_H) = 0 \Leftrightarrow \frac{1-q^*}{q^*} = \frac{\gamma_H(s_{n^*})}{\gamma_L(s_{n^*})}\frac{v_H - c_H}{c_H - v_L}. \tag{3}$$

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 trades at the same rate as the high type. The high type accepts only c_H. Therefore, given buyers' offer strategies characterized by (n^*, σ_B^*), her trading 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 trading rate is equal to

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

Then, σ_S^* is the unique value that equates the above two rates, that is,

$$\rho_H = \rho_L + (1-\rho_L)\sigma_S^*. \tag{4}$$

It is well-defined in $(0,1)$ because $\Gamma_H(\cdot)$ first-order stochastically dominates $\Gamma_L(\cdot)$.

Proposition 1 summarizes the ongoing discussion.

Proposition 1 (Stationary path) *Let q^*, n^*, σ_B^*, and σ_S^* be the values defined by* (1), (2), (3), *and* (4). *Once buyers' beliefs reach q^*, the following strategy profile constitutes a continuation equilibrium: Each buyer offers c_H with probability 1 if his signal is strictly above s_{n^*} and with probability σ_B^* if his signal is s_{n^*}, and offers v_L otherwise. The low-type seller accepts $p(q^*) = v_L$ with probability σ_S^*. Along the path of this equilibrium, buyers' beliefs stay constant at q^*. Moreover, this is the unique non-trivial stationary path: There is no other belief level below $\overline{q}_1$ at which buyers' beliefs can stay constant. In addition, there is no other strategy profile under which buyers' beliefs remain constant at q^*.*

4.1.2 Evolution of Beliefs

Given n^* and the partition $\{\overline{q}_{N+1}, \overline{q}_N, ..., \overline{q}_1, \overline{q}_0\}$ described in condition (ii), buyers' beliefs evolve deterministically, conditional on no trade. It is convenient to describe the belief evolution separately for three distinct regions of beliefs.

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

Next, consider $q(t) \in (q_{n+1}, q_n)$ with $0 < n < n^*$, so that $q(t) \in (q^*, \overline{q}_1)$. In this case, due to (i), $p(q(t)) > v_L$. This implies that the low-type seller, as well as the high-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 trades at rate $\lambda(1 - \Gamma_L(s_n))$, while the high type trades at rate $\lambda(1 - \Gamma_H(s_n))$. It then follows that $q(\cdot)$ evolves according to the following law of motion:[17]

$$\dot{q}(t) = q(t)(1 - q(t))\lambda(\Gamma_H(s_n) - \Gamma_L(s_n)). \tag{5}$$

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 $\Gamma_H(\cdot) - \Gamma_L(\cdot)$ changes as well.

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

$$\dot{q}(t) = q(t)(1 - q(t))\lambda\Gamma_H(s_n). \tag{6}$$

This expression is obviously positive, and thus $q(t)$ strictly increases in t. The dashed line in Figure 2 exemplifies a typical path that buyers' beliefs follow starting 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 $q(t)$ stays constant. If $0 < n < n^*$, then $q(t)$ evolves according to (5). If $n \geq n^*$, then $q(t)$ evolves according to (6).*

[17] 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.

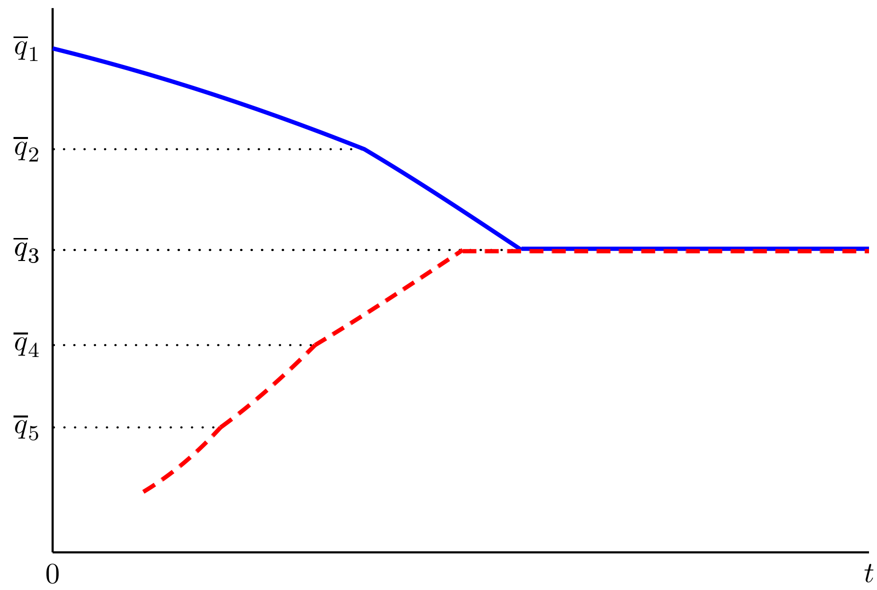

Figure 2: The evolution of buyers' beliefs. The parameter values used for this figure and Figure 3 are as follows: $N = 5$, $\gamma_L(s_n) = (6-n)/15$, $\gamma_H(s_n) = n/15$, $r = 0.1$, and $\lambda = 0.4$.

Two remarks are in order. First, starting from $\widehat{q}$, regardless of whether $\widehat{q} \in (q^*, \overline{q}_1)$ or $\widehat{q} < q^*$, buyers' beliefs $q(\cdot)$ converge to q^* in finite time, because $\dot{q}(t)$ is bounded away from 0.[18] 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 point in time when his belief is exactly 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.

[18]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

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

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 the reservation price schedule $p(q)$ and the 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 has prior belief $\overline{q}_n$ and receives signal s_n. This 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

$$\frac{1-\overline{q}_n}{\overline{q}_n} = \frac{\gamma_H(s_n)}{\gamma_L(s_n)} \frac{v_H - c_H}{c_H - \min\{v_L, p(\overline{q}_n)\}}. \tag{7}$$

The use of $\min\{v_L, p(q)\}$, instead of $p(q)$, reflects the fact that if $p(q) \geq v_L$, then the buyer's expected payoff from 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 $\{\overline{q}_N, ..., \overline{q}_1\}$. Notice that, $p(q)$ is determined only by the rate at which the low type receives offer c_H: each buyer offers either c_H or $p(t)$, and the low type is indifferent between accepting and rejecting $p(t)$. Therefore, for the purpose of calculating the expected payoff, it can be assumed that she accepts only c_H. Then, the low-type seller's reservation price 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,

$$p(q) = \int_0^{T(q,\tilde{q})} ((1-e^{-rt})c_L + e^{-rt}c_H) d\left(1 - e^{-\lambda(1-\Gamma_L(s_n))t}\right) + e^{-(r+\lambda(1-\Gamma_L(s_n)))T(q,\tilde{q})} p(\tilde{q}), \tag{8}$$

where $T(q, \tilde{q})$ is the length of time it takes for buyers' beliefs to move from q to $\tilde{q}$.[19] A closed-form

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

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

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

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

solution can be found by applying the definition of $T(\cdot,\cdot)$ and using the fact that $p(\overline{q}_{n^*}) = v_L$.

Finally, we combine the findings so far 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 (7) reduces to

$$\frac{1-\overline{q}_n}{\overline{q}_n} = \frac{\gamma_H(s_n)}{\gamma_L(s_n)}\frac{v_H - c_H}{c_H - v_L}. \tag{9}$$

Therefore, the cutoffs $\overline{q}_n$ are determined independently of the value of $p(q)$. The uniqueness follows from the fact that the left-hand side is decreasing in $\overline{q}_n$, while the right-hand side is constant. Given $\overline{q}_n$ for each $n < n^*$, $p(q)$ can be calculated using (8) for any $q > q^*$.

The determination of $\overline{q}_n$ for $n > n^*$ is more involved, because these cutoffs 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 $k = 0, ..., N - n^* - 1$,

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

  where $T(\cdot,\cdot)$ is defined in footnote 19.

- $\overline{q}_{n^*+k+1}$ 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 $p(\overline{q}_{n^*}) = v_L$.

To understand what each step accomplishes, consider the case of $k = 0$. In this case, $p(q^*) = p(\overline{q}_{n^*}) = v_L$ is already known. Then it is possible to compute the (hypothetical) reservation price of the low-type seller starting from any belief $q < q^*$, under the supposition that (i) she receives offer c_H at rate $\lambda(1 - \Gamma_L(s_{n^*}))$; (ii) the low type trades at rate λ, while the high type trades at rate $\lambda(1 - \Gamma_H(s_{n^*}))$. This is what is done in the first step.[20] Given $p(\cdot)$, the second step identifies the value of $\overline{q}_{n^*+1}$, using the fact that the buyer with prior belief $\overline{q}_{n^*+1}$ and signal s_{n^*+1} must be

[20]Notice that the equation in the first step is simply a special case of (8).

indifferent between offering c_H and $p(\overline{q}_{n^*+1})$. Note that the fact that the low-type seller accepts $p(\overline{q}_{n^*+1}) < v_L$ with probability 1 is also used to derive the equation in the second step.

The above two steps characterize the cutoff beliefs $\bar{q}_n$ for each $n \geq n^*$. Given these cutoffs, (8) can be used to compute $p(q)$ for each $q < q^*$, as in the previous case.

The following theorem summarizes the results so far.

Theorem 1 *For generic values of the parameters, 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 offers $p(q(t))$ 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^*, the game is played as described in Proposition 1. This is the unique equilibrium 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 (e.g., $\widehat{q}_1$ in the figure), then the low-type seller's reservation price $p(\widehat{q}_1)$ 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 price 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 remains constant at v_L. If buyers' initial beliefs are pessimistic (e.g., $\widehat{q}_2$ in the figure), then the resulting dynamics is exactly the opposite: buyers' beliefs and the low-type seller's reservation price increase over time and converge to q^* and v_L, respectively.

4.2 Uniqueness

We now prove that the equilibrium described 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 other things, 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 $q < \overline{q}_1$. 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).

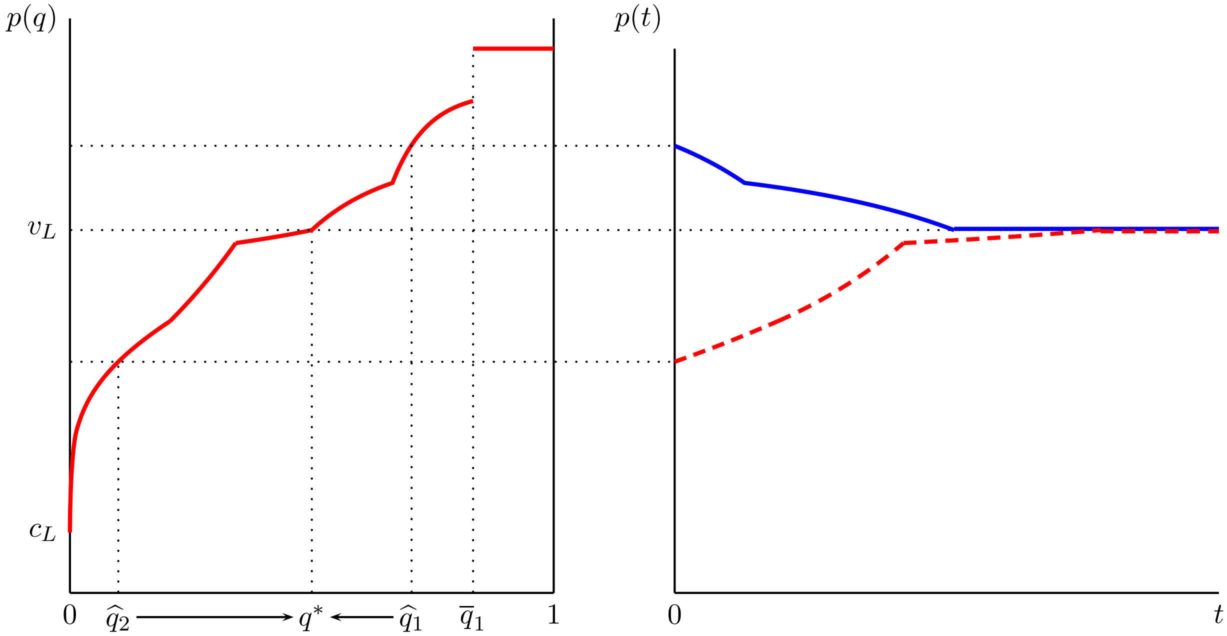

Figure 3: The low-type seller's reservation price as a function of buyers' beliefs q (left) and that of time t (right).

4.2.1 Monotonicity of Beliefs

By Lemma 1, 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 follows 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 $p(\bar{t}) < v_L$, then 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: if t is after $\bar{t}$, but before the convergence occurs, then $p(t) < v_L$, and thus $q(t)$ must be increasing, which implies $q(t) > q(\bar{t}) > q^*$. In the meantime, the low-type seller's reservation price 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 close to v_L. On the other hand, upon receiving signal 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 a function of q. Let $\bar{s}(q)$ denote the cutoff signal that a buyer with prior belief q uses. We now establish that $p(q)$ is increasing, while $\bar{s}(q)$ is non-increasing. We argue this separately for the two ranges of beliefs: $q \in (q^*, \overline{q}_1)$ and $q < q^*$.

First consider $q \in (q^*, \overline{q}_1)$. Then by Lemma 2, $\min\{p(q), v_L\} = v_L$. Inspecting (7) immediately reveals that for this range of beliefs, $\bar{s}(q)$ is non-increasing. This, in turn, implies the monotonicity of $p(q)$ via (8).

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 immediately follows. The next lemma establishes that $p(t)$ is increasing over time whenever $q(t) < q^*$.

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 meantime, 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 both buyers' beliefs and the low-type seller's reservation price are higher after time t' than within the interval (t, t'). This leads to a contradiction.

4.2.3 Proof of Theorem 2

Let us summarize how all the components we have established so far prove the uniqueness of equilibrium.

First, Proposition 1 establishes that *if* buyers' beliefs are to remain constant at some level less than $\overline{q}_1$, this level must be q^* as defined in Section 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 they reach q^*.

If $\widehat{q} > q^*$, then, by Lemma 1, the low-type seller, as well as the high-type seller, accepts only c_H, until buyers' beliefs reach q^*. This uniquely pins down buyers' behavior for $q > q^*$ via (9). Then, the low-type seller's reservation prices for this range of beliefs are determined by (8).

Finally, if $\widehat{q} < q^*$, then the uniqueness argument requires an extra step, since in this case the 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 $p(t)$ must be increasing over time. Given this result, it is clear that any equilibrium must be constructed as in Section 4.1.3. The uniqueness of equilibrium then follows from the fact that the construction yields a unique strategy profile.

5 Equilibrium Outcomes in the Frictionless Limit

In this section we present the limit of equilibrium outcomes as the arrival rate of buyers λ tends to infinity. This limit 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 trade can be facilitated through such policies.[21] 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 has an opportunity to trade at each instant. 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 positive gains from trade regardless of the asset quality, the earlier trade takes place, the larger is the surplus generated in the market. Therefore, the time to trade can be considered a measure of surplus generated. Meanwhile, the low-type seller's expected payoff can be interpreted as a measure of the 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 station-

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

ary cutoff signal s^* is necessarily equal to the highest signal s_N.[22] Intuitively, if buyers arrive frequently, then the low-type seller has a strong incentive to wait for c_H. Expecting this, buyers would be even more cautious and, therefore, would offer c_H only when they receive the highest signal s_N (even then with small probability). This implies that the stationary belief q^* is equal to $\overline{q}_N$, which in turn implies that for any n, $\overline{q}_n$ can be uncovered from the following equation:

$$\frac{1-\overline{q}_n}{\overline{q}_n} = \frac{\gamma_H(s_n)}{\gamma_L(s_n)}\frac{v_H - c_H}{c_H - v_L}. \tag{10}$$

It also implies that the rate at which each type trades along the stationary path is given by

$$\rho_H = \frac{\gamma_H(s_N)}{\gamma_L(s_N)}\frac{v_L - c_L}{c_H - v_L}r, \tag{11}$$

which is independent of λ.[23]

Second, buyers' beliefs immediately jump to q^* in the frictionless limit, as long as their initial beliefs are smaller than $\overline{q}_1$. 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 19). Therefore, the length of time for buyers' beliefs to move from $\widehat{q}$ to q^* also shrinks to 0. Intuitively, at each belief level, the 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 these distributions and, in particular, the probability that each type trades immediately.

Proposition 3 *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 (11) and

[22]The precise condition under which this is true is $\lambda > \frac{r(v_L - c_L)}{\gamma_L(s_N)(c_H - v_L)}$. We consider only the case where $\gamma_L(s_N) > 0$. If $\gamma_L(s_N) = 0$, then the result is trivial: If the asset is of high quality, then each buyer receives the perfectly informative signal s_N with a positive probability. Therefore, in the limit as λ tends to infinity, the high type trades immediately. Given this, the low type also trades immediately.

[23]If $n^* = N$, then the rates at which the two types receive offer c_H are given by $\rho_L = \lambda\gamma_L(s_N)\sigma_B^*$ and $\rho_H = \lambda\gamma_H(s_N)\sigma_B^*$, respectively. The expression in (11) is obtained by solving for σ_B^* with the fact that ρ_L is necessarily equal to $\frac{v_L - c_L}{c_H - v_L}r$.

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

$$F_L(0|\widehat{q}) = 1 - \frac{\widehat{q}}{1-\widehat{q}} \frac{\gamma_H(s_N)}{\gamma_L(s_N)} \frac{v_H - c_H}{c_H - v_L} \quad \textit{and} \quad F_H(0|\widehat{q}) = 0; \tag{12}$$

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

$$F_a(0|\widehat{q}) = 1 - \left(\frac{1}{l_n} \frac{c_H - v_L}{v_H - c_H} \frac{1-\hat{q}}{\hat{q}} \right)^{\psi_n^a} \times \prod_{i=n+1}^{N} \left(\frac{l_{i-1}}{l_i} \right)^{\psi_i^a}, \tag{13}$$

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 (12) 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^*.[24] Similarly, in (13), the first multiplicative term is the probability with which the type-a seller should not trade in order for buyers' beliefs to 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 jump from $\overline{q}_i$ to $\overline{q}_{i-1}$.[25]

We now turn to the expected payoff of the low-type seller, which is necessarily equal to her reservation price with belief $\widehat{q}$.[26] 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^*). Therefore, $p(\widehat{q})$ is simply a weighted average of c_H and v_L, with the weight to c_H given by the probability that she trades immediately at price c_H. Recall that if $\widehat{q} < q^*$, then no buyer offers c_H until buyers' beliefs reach

[24] Precisely, due to (10),

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

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

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

[26] Note that the seller derives flow payoff from owning the asset.

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 3.

Proposition 4 *In the limit as λ tends to infinity, the low-type seller's expected payoff is equal to*

$$p(\widehat{q}) = \begin{cases} v_L, & \textit{if } \widehat{q} \leq q^*, \\ v_L + F_L(0|\widehat{q})(c_H - v_L), & \textit{if } \widehat{q} \in (q^*, \overline{q}_1), \\ c_H, & \textit{if } \widehat{q} \geq \overline{q}_1. \end{cases} \tag{14}$$

It is well-recognized that in dynamic environments, real-time delay is unavoidable with severe adverse selection: for immediate trade, it is necessary that all types of the seller trade at the same price. However, if adverse selection is rather severe, in the sense that the probability that the seller is the low type is sufficiently high, then the high-type seller's reservation value exceeds buyers' expected valuation of the asset, and thus there does not exist a price that yields non-negative payoffs to all types of the seller as well as buyers.

The presence of real-time delay in our model, demonstrated in Proposition 3, does not immediately follow from this well-known wisdom. In the presence of inspection which exogenously reveals type-dependent information, even if the low-type seller exactly mimics the equilibrium behavior of the high-type seller, she cannot generate (statistically) identical outcomes.[27] Since each arriving buyer generates a fixed amount of information about the seller's type (by drawing an inspection outcome), in the frictionless limit, potentially there is sufficient information in the cumulative inspection outcomes to fully reveal the seller's type. If this were the case, the need for delay to distinguish between the two types of the seller would disappear. However, in the absence of a fully informative signal, this mechanism fails.[28] The reason why it fails is subtle and lies in the strategic interaction between the buyers and the seller. As λ increases, each buyer becomes more reluctant to offer c_H, because the low-type seller's incentive to wait for c_H–i.e. mimic the high type– increases. Then, buyers use the most informative signal (the highest signal) as their cutoff, and even with that signal, offer c_H only with a vanishing probability. This is why real time delay persists in our model.

[27] This, in particular, eliminates the above-mentioned requirement that for immediate trade both types must trade at the same price (with probability 1). Indeed, as shown in Section 4.1.1, on the stationary path from which real-time delay results, different types of the seller trade at different prices: the high type trades only at c_H, while the low type trades not only at c_H, but also at v_L.

[28] If $\gamma_L(s_N) = 0$, then delay would disappear, since offering c_H if and only if the signal is s_N would provide a trivial way of screening out the low type.

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, but to obtain some preliminary insights. To this end, we focus on the simplest case where 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.[29]

Specifically, we examine how the market outcomes change as the likelihoods of the signals vary. In particular, we focus on the effects of a marginal increase in $l_2 \equiv \gamma_L(s_1)/\gamma_H(s_1)$.[30] Such changes in the parameter values are consistent with various common notions of more precise signals.[31] In the same spirit as in Section 5, we focus on the effects of these changes on the expected time to trade (as a measure of market liquidity and efficiency) and the low-type seller's expected payoff (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 parsimony, we simply say $\widehat{q} > q^*$, in order to refer to $\widehat{q} \in (q^*, \overline{q}_1)$.

Proposition 3 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. 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 rate is given by ρ_H. Formally, $E[\tau_a(\widehat{q})] = (1 - F_a(0|\widehat{q}))/\rho_H$. The following proposition is then immediate by basic calculus.[32]

Proposition 5 *Suppose l_2 increases marginally.*

- *If $\widehat{q} < q^*$, then $E[\tau_L(\widehat{q})]$ remains constant, while $E[\tau_H(\widehat{q})]$ decreases.*

---

[29]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.

[30]We explain the effects of a marginal decrease in $l_1 \equiv \gamma_L(s_1)/\gamma_H(s_1)$ at the end of this section.

[31]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 all of them.

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

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

In addition,

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

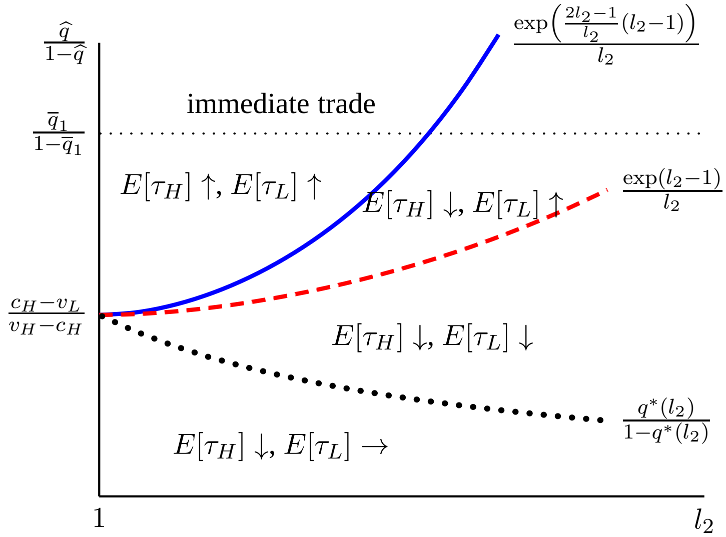

Figure 4: The effects of an increase in l_2 on the expected times to trade.

- *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. Moreover, there are three systematic patterns governing how this effect relates to market conditions. First, more precise signals are less likely to be beneficial when $\widehat{q}$ is high; i.e., the range of l_2 for which a further increase in l_2 speeds up trade becomes smaller 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 $\widehat{q}$ for which an increase in l_2 speeds up trade becomes larger as l_2 increases. Finally, increased precision tends to speed up trade for the high type more than for the low type.

To understand the first two patterns, notice that the increased precision of signals affects the equilibrium dynamics through two channels. First, since each buyer's *own* signal becomes more informative, it directly reduces his risk of paying a high price for a low-quality asset, thereby encouraging him to offer c_H more often. Ceteris paribus, this effect reduces the expected time to trade.[33] Second, since *other buyers'* signals also become more informative, beliefs evolve faster when delay indicates the inability to generate good signals, i.e., when delay carries negative information about the signals of previous buyers. Recall that this is the case when $\widehat{q} > q^*$. In this case, for a given length of 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

[33] 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 $\rho_H = \frac{r(v_L-c_L)}{c_H-v_L}$.

buyers' incentive to offer a high price, thereby slowing down trade.[34] The aforementioned 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 in 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 trades also 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 countervails the effect due to an increase in buyers' willingness to offer c_H: 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 countervailing 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.[35]

Remark 1 (Distribution of time to trade) Although we have focused on how the expected times to trade respond to a change in 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 in the distribution cannot be clearly ranked. Since first-order stochastic dominance implies an increase in the expected value, it follows that the region at which an increase in l_2 speeds up trade in the 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 4 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.[36] The following result is then straightforward to obtain from Proposition 3.

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

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

[36]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 sense.

Proposition 6 *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 stronger indication of low quality, and thus for a buyer with signal s_1 to be willing to offer c_H, 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 Discussion

In this section, we discuss two technical issues that are relevant to our equilibrium characterization.

7.1 Equilibria in Non-generic Cases

We have restricted attention to the (generic) case where there does not exist n^* such that $\rho_L = \frac{r(v_L-c_L)}{\lambda(c_H-v_L)} = \lambda(1-\Gamma_L(s_{n^*}))$. In this section, we show that the equilibrium uniqueness fails in the non-generic cases and illustrate its underlying cause. In addition, we show that, nevertheless, the resulting trading dynamics exhibits the same patterns as in the generic cases.

To illustrate the basic problem of the non-generic cases, suppose $\rho_L = \lambda(1-\Gamma_L(s_{n^*}))$, and 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 s_{n^*}. Then, the low-type seller's reservation price stays equal to v_L, because she receives offer c_H at a constant rate of $\rho_L = \lambda(1-\Gamma_L(s_{n^*}))$. Suppose she accepts v_L with probability σ_S^*, where σ_S^* is the value that satisfies

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

Then, by construction, buyers' beliefs do not change over time: the left-hand side represents the trading rate of the high-type seller, while the right-hand side that of the low-type seller. Furthermore, given that v_L is the low-type seller's reservation price, buyers' offer strategies described above 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 accept v_L with probability 1, until buyers' beliefs reach $\overline{q}_{n^*}$ and stay constant thereafter. Or, she may reject 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 $[\overline{q}_{n^*+1}, \overline{q}_{n^*}]$, 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 this range. 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 when buyers' beliefs are in the interval $[\overline{q}_{n^*+1}, \overline{q}_{n^*}]$,[37] as it varies depending on the low-type seller's acceptance strategy of offer v_L.

7.2 A Continuum of Signals

Even though we have chosen to consider an environment with finite number of signals (possible inspection outcomes), it is easy to see that the qualitative properties of the equilibrium characterized in Section 4 carries over to a model with a continuum of signals: There is a unique level of belief q^* which supports stationary equilibrium play. In addition, buyers' beliefs, whether they are above

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

or below q^*, gradually converge to q^*. The low-type seller's reservation also gradually converges to v_L.

One technical advantage of a model with a continuum of signals is that in such a model, no randomization is necessary on the equilibrium path. In particular, the buyers' randomization between c_H and v_L when they receive the cutoff signal on the stationary path would not be needed. Formally, suppose the signal space $\mathcal{S}$ is given by an interval $[\underline{s}, \overline{s}]$ and $\Gamma_a(\cdot)$ (respectively, $\gamma_a(\cdot)$) now represents the distribution (respectively, density) function of signals conditional on type a. Let s^* be the unique value that satisfies

$$\lambda(1 - \Gamma_L(s^*)) = \rho_L = \frac{r(v_L - c_L)}{c_H - v_L}. \tag{15}$$

Then, for the stationary path to be sustained, it suffices that each buyer offers c_H if and only if his signal is above s^*.

One disadvantage of having a continuum of signals is that in such a model, the equilibrium can be described only by a system of partial differential equations without closed-form solutions, unlike in the finite-signal case. This is because with a continuum of signals, the cutoff signal, as well as buyers' beliefs and the low-type seller's reservation price, varies continuously. This particularly complicates the analysis of the cases where the initial belief is below q^*, because all three equilibrium objects influence one another and, therefore, must be simultaneously determined in those cases.[38]

It is also easy to see how the analysis of the frictionless limit case in Section 5 carries over with a continuum of signals. As λ tends to infinity, the stationary cutoff signal approaches the highest signal $\overline{s}$ (see (15)). This implies that the stationary belief level converges to q^* that satisfies

$$\frac{1 - q^*}{q^*} = \frac{\gamma_H(\overline{s})}{\gamma_L(\overline{s})} \frac{v_H - c_H}{c_H - v_L}.$$

This, in turn, implies that in the limit, buyers' beliefs immediately jump to q^*, whether the initial belief is above or below q^*.

8 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

[38] If $\widehat{q} > q^*$, then buyers' optimal offer strategies (cutoff signals) and the evolution of buyers' beliefs can be determined independently of the low-type seller's reservation price $p(t)$, because $p(t) > v_L$, and thus $p(t)$ is never accepted in equilibrium.

we are able to accommodate different sources of trading delay and provide an understanding of how these sources interact with one another. 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.

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 exist $t' \in (t, \infty)$ 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 1), 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 $\frac{rc_L+\lambda c_H}{r+\lambda}$. This is a contradiction, because $\frac{rc_L+\lambda c_H}{r+\lambda}$ is strictly larger than v_L by Assumption 1.

Let t' be the smallest value such that $p(t') = v_L$. 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 $p(t) = v_L$.

Suppose $q(t) = q^*$ but $p(t) < v_L$. Since buyers' beliefs would be increasing, there would exist t' such that $q(t') > q^*$, while $p(t') < v_L$, which is a contradiction. Symmetric arguments lead to a contradiction for the case where $p(t) > v_L$. ■

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 follows that $p(t) < p(t')$, which is a contradiction. ■

Proof of Proposition 3:. First consider the case where $\widehat{q} < q^*$. In that case, for any fixed λ, the high-type seller does not trade until buyers' beliefs reach q^*. Therefore, $F_H(0|\widehat{q}) = 0$. The low type, on the other hand, trades with probability 1 conditional on the arrival of a buyer during this period. Therefore, the probability that she trades on the convergence path is equal to $1-e^{-\lambda T(\widehat{q},q^*)}$, where $T(\widehat{q},q^*)$ is the value that satisfies

$$q^* = \frac{\widehat{q}}{\widehat{q}+(1-\widehat{q})e^{-\lambda T(\widehat{q},q^*)}}.$$

Combining this with (10),

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

Next, consider the case where $\widehat{q} \in (\overline{q}_n, \overline{q}_{n-1})$ for some $n \leq n^*$. In this case, conditional on the arrival of a buyer, each type trades if and only if the buyer offers c_H, which happens with probability $1-\Gamma_a(s_{n'})$ if $q \in [\overline{q}_{n'+1}, \overline{q}_{n'})$. Therefore, the probability that the type-a seller trades before buyers' beliefs reach q^* is equal to

$$1-e^{-\lambda(1-\Gamma_a(s_{n-1}))T(\widehat{q},\overline{q}_n)}\prod_{i=n}^{N-1} e^{-\lambda(1-\Gamma_a(s_i))T(\overline{q}_i,\overline{q}_{i+1})},$$

where

$$\overline{q}_n = \frac{\widehat{q}e^{-\lambda(1-\Gamma_H(s_{n-1}))}}{\widehat{q}e^{-\lambda(1-\Gamma_H(s_{n-1}))}+(1-\widehat{q})e^{-\lambda(1-\Gamma_L(s_{n-1}))}},$$

and

$$\overline{q}_{i+1} = \frac{\overline{q}_i e^{-\lambda(1-\Gamma_H(s_n))}}{\overline{q}_i e^{-\lambda(1-\Gamma_H(s_n))}+(1-\overline{q}_i)e^{-\lambda(1-\Gamma_L(s_n))}},$$

The result follows by successively applying (10).

■

References

Banerjee, Abhijit V., "A simple model of herd behavior," *Quarterly Journal of Economics*, 1992, *107* (3), 797–817.

Bikhchandani, Sushi, David Hirshleifer, and Ivo Welch, "A theory of fads, fashion, custom, and cultural change as informational cascades," *Journal of Political Economy*, 1992, *100* (5), 992–1026.

Blackwell, David, "Comparison of experiments," *Proceedings, Second Berkeley Symposium on Mathematical Statistics and Probability*, 1951, pp. 93–102.

Camargo, Braz and Benjamin Lester, "Trading dynamics in decentralized markets with adverse selection," *mimeo*, 2011.

Chang, Briana, "Adverse selection and liquidity distortion," *mimeo*, 2012.

Chari, V.V., Ali Shourideh, and Ariel Zetlin-Jones, "Adverse selection, reputation and sudden collapses in secondary loan markets," *mimeo*, 2012.

Chi, Jonathan and Thorsten Koeppl, "Trading dynamics with adverse selection and search: market freeze, intervention and recovery," *mimeo*, 2011.

Choi, Michael, "Asset prices and information transmission," *mimeo*, 2013.

Daley, Brendan and Brett Green, "Waiting for news in the dynamic market for lemons," *Econometrica*, 2012, *80* (4), 1433–1504.

Deneckere, Raymond and Meng-Yu Liang, "Bargaining with interdependent values," *Econometrica*, 2006, *74* (5), 1309–1364.

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

Fuchs, William and Andrzej Skrzypacz, "Bridging the gap: bargaining with interdependent values," *Journal of Economic Theory*, 2013, *forthcoming.*

_ and _ , "Costs and benefits of dynamic trading in a Lemons market," *mimeo*, 2013.

_ , Aniko Öry, and Andrzej Skrzypacz, "Transparency and distressed sales under asymmetric information," *mimeo*, 2012.

Guerrieri, Veronica and Robert Shimer, "Dynamic adverse selection: a theory of illiquidity, fire sales, and flight to quality," *mimeo*, 2012.

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

Janssen, Maarten C. W. and Santanu Roy, "Dynamic trading in a durable good market with asymmetric information," *International Economic Review*, 2002, *43* (1), 257–282.

Kaya, Ayca and Qingmin Liu, "Transparency and price formation," *mimeo*, 2013.

Kim, Kyungmin, "Information about sellers' past behavior in the market for lemons," *mimeo*, 2012.

Kurlat, Pablo, "Lemons, market shutdowns and learning," *mimeo*, 2010.

Lauermann, Stephan and Asher Wolinsky, "Search with adverse selection," *mimeo*, 2013.

Lehmann, Eric L., "Comparing location experiments," *The Annals of Statistics*, 1988, *16* (2), 521–533.

Moreno, Diego and John Wooders, "Decentralized trade mitigates the lemons problem," *International Economic Review*, 2010, *51* (2), 383–399.

_ and _ , "Dynamic markets for lemons: performance, liquidity, and policy evaluation," *mimeo*, 2012.

Nöldeke, Georg and Eric van Damme, "Signalling in a dynamic labor market," *Review of Economic Studies*, 1990, *57* (1), 1–23.

Roy, Santanu, "Dynamic sorting in durable goods markets with buyer heterogeneity," *Canadian Journal of Economics*, 2013, *forthcoming.*

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

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

Swinkels, Jeroen, "Education signalling with preemptive offers," *Review of Economic Studies*, 1999, *66* (4), 949–970.

Taylor, Curtis R., "Time-on-the-market as a sign of quality," *Review of Economic Studies*, 1999, *66* (3), 555–578.

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

_ , "Dynamic auctions," *Review of Economic Studies*, 1990, *57* (1), 49–61.

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