Waiting for News in the Dynamic Market for Lemons*

Brendan Daley[†] Brett Green[‡§]

December 2008

Abstract

Trade breaks down in the market for lemons because high-type sellers have a reservation value greater than expected market value. Unraveling occurs and only the lowest types trade. Two related questions arise: What happens the next day? And, from where does the reservation value come? We model these considerations in a dynamic setting with gradual arrival of noisy information. We characterize the unique equilibrium in a continuous-time framework. The equilibrium involves a region of no trade or *market failure*. The no-trade region ends in one of two ways: either enough good news arrives restoring confidence and markets re-open or bad news arrives making buyers more pessimistic forcing market *capitulation* i.e., a sell-off of low value assets. Reservation values arise endogenously from the *option* to sell in the future. Our model also encompasses dynamic signaling environments. In a dynamic setting with sufficiently informative news, Spence's *Job Market Signaling* and Akerlof's *Market for Lemons* have the same unique equilibrium. The predictions help explain "irrational" trading patterns in financial markets.

*The authors are greatly indebted to Jeremy Bulow, Michael Harrison and Andrzej Skrzypacz for useful conversations, comments, and encouragement. Special thanks also to Peter DeMarzo, Ilan Kremer, Jon Levin, Paul Milgrom, Yuliy Sannikov, Bob Wilson and the seminar participants at Stanford, Duke and UNC for their useful comments and conversations.

†Fuqua School of Business, Duke University **Email**:bd2@duke.edu **Web**: http://faculty.fuqua.duke.edu/∼bd28

‡Stanford GSB (Job Market Paper 2008-09) **Email**:greenb@stanford.edu **Web**: www.stanford.edu/∼greenb

§Brett gratefully acknowledges the support of the SIEPR Dissertation Fellowship.

Introduction

Market breakdown occurs in the market for lemons because high-value sellers have a reservation value greater than expected market value (Akerlof, 1970). Unraveling occurs, and only the lowest types trade. Two related questions arise. What happens the next day? If all the lowest types sell on day one, a buyer should be willing to pay more on day two, destroying the equilibrium. Also, from where does the reservation value come? The option to sell at a later date surely contributes in addition to any consumption value derived in the meantime. We model these considerations in a dynamic market where, as time passes:

- A privately informed seller of an asset receives offers from a market of buyers.
- Prior to trade, the seller accrues a flow payoff that depends on the asset's value
- News about the asset is gradually revealed by an exogenous stochastic process.

For example, consider an entrepreneur who is interested in selling her company due to liquidity constraints. Naturally, she is better informed about the company's fundamentals than is the market. She would like to sell, yet there is no reason she is forced to sell on any given day—if she does not sell today, she can sell tomorrow or any day after. With every passing day that she does not sell, she gains (or loses) the day's profit, and the market learns more about her company through observation of sales data, revenue, customer base, etc.

There are several intricacies. What offers should buyers make at date t conditional on what they have learned? Clearly, they have exposure to an adverse selection problem. How long should the entrepreneur with a high-value company wait for favorable information to be revealed? How long should the entrepreneur with a low-value company hold out and risk too much negative information being revealed about it?

Our first main result is the characterization of equilibrium trading behavior. The equilibrium involves three distinct regions. When beliefs about the seller are favorable, the market is fully efficient—trade is immediate with price set at the expected market value of the asset. When beliefs are unfavorable, there is a partial sell-off in which only low-type assets are traded with positive probability. When beliefs are intermediate, the market dries up—no trade occurs, and the market waits for more information to be revealed about the seller's asset before making serious offers.

As a second example, consider a student's use of education to signal his ability to potential employers. On any given day, he can leave school and take up employment. With every passing day that he remains in school he incurs the cost (or joy) of education and completes more graded coursework, which reveals information about his ability to potential employers.

It is clear that the two examples share a similar structure. However, most economists consider the first to be a market for *lemons*, while the second is considered a model of *signaling* (Spence, 1973). Our model encompasses both environments. We demonstrate that their only difference is a single inequality on parameters. Namely, the high-type seller would never accept the market value for a low type under the lemons condition. Under the signaling condition, she would.

Our second main result is that market behavior is characterized by the same three regions under the signaling condition as long as the news arrives fast enough.[1] An *endogenous* lemons problem develops because the *option to sell* in the future drives the seller's continuation value above the market price. Trade is delayed despite the fact that buyers are always willing to offer more than the high types' *outside option*.[2]

Unlike under the lemons condition, the results for signaling depend on the rate at which information arrives. The lemons condition implies that the high type is never tempted to sell when the market belief about her asset is unfavorable—foregoing offers carries no opportunity cost until the market belief improves. In contrast, under the signaling condition, the high type's opportunity cost is positive regardless of the market belief. Only when news arrives quickly does the expected benefit from higher prices in the future make the high type willing to forego immediate trade.

The game is played in continuous time where news arrives according to a Brownian diffusion process. This setting makes the analysis tractable; the equilibrium is precisely characterized through a system of differential equations. It also allows us to solve examples and analyze market efficiency numerically. In Appendix B we consider the following discrete-time analog: in every period of length Δ, a market of buyers make offers to a privately informed seller. Then good or bad news is revealed about the asset via a random walk process before the next period begins. The unique equilibrium in the discrete-time version exhibits the same three regions: immediate trade, no trade and partial sell-off. As $\Delta \to 0$, the equilibrium converges to the equilibrium of the continuous-time game.

In equilibrium, both type sellers trade eventually. Informative news mitigates the well-known inefficiency associated with market breakdown in pessimistic markets. However, a new inefficiency develops for intermediate beliefs due to delay. When news is completely uninformative, the welfare to each player is the same as in Akerlof's original static model. The welfare of the high-type seller improves with the quality of the news, while the low type does worse. We demonstrate how the quality of the news affects the severity of the *dry market*: the expected inefficiency caused by delay is single-peaked in the news quality. When the quality is low, neither seller nor buyer has any incentive to wait for arrival of meaningless information. The inefficiency from delay is small. On the opposite extreme, when the quality is very high, a low type has virtually nothing to gain by holding out because the market will almost surely correctly identify her. In equilibrium, the sell-off region grows, leading to faster trade. It is only for intermediate qualities that both type sellers have incentive to wait: the low type takes his chances hoping to get lucky while the high type knows that good news will eventually fetch her a higher price.

In the leading example, a company that has been sold will never be re-traded among homogeneous agents. However, in most liquid financial markets, assets are traded repeatedly. The value

[1] We measure news by the signal-to-noise ratio, which can be interpreted either as the rate at which information arrives (i.e., how long does it take to learn a certain amount) or the quality of the arriving information (i.e., how much can be learned in a certain amount of time).

[2] A seller's *outside option* is her payoff from never trading, whereas her *continuation value* is her expected payoff from rejecting *in equilibrium*.

to a trader is determined primarily from the price at which the asset can be sold in the future rather than from dividends or coupon payments. To more appropriately describe such settings, we introduce random arrivals of liquidity shocks in Part II. The shock increases the rate at which an owner of the asset discounts future payoffs. The owner of an asset is *not* forced to sell upon arrival, but she is more eager to do so. A buyer's value for the asset depends not only on her beliefs about the asset type but also on her expectations about future liquidity in the market.

We characterize the equilibrium in the model with shocks through a system of differential equations and describe an algorithm for solving it. Using numerical techniques, we solve for the equilibrium and its dependence on the underlying parameters. The prices at which assets are traded are lower than in the model without shocks. However, the equilibrium structure is similar. It involves three analogous regions: (1) the market is liquid when conditions are favorable, sellers are able to trade quickly at high prices; (2) a sell-off region when the market is very pessimistic, an owner hit by a shock in this region is forced to either sell at rock-bottom prices or hold out; (3) a no-trade region where both sides of the market wait for news until either good news restores confidence to (1) or bad news forces (2). Asset prices decrease with the arrival rate of shocks because investors face costly liquidation more frequently. The equilibrium asset prices converge to the prices in the model without shocks as the arrival rate goes to zero.

The predictions of the model may help to explain some of the seemingly irrational trading patterns observed in financial markets. For example, our model predicts a sell-off of assets at low prices can often help stabilize a shaky market. Wall Street traders and analysts refer to this as *market capitulation*.[3] We also find that a small amount of bad news can cause drastic changes in prices, volume and liquidity.[4] The recent collapse of the mortgage-backed securities market is a particularly relevant example. Until recently, trade and issuance of mortgage-backed securities has occurred in a liquid and well functioning market. News of the declining housing market created uncertainty in the value of the underlying collateral and led to a catastrophic drop in liquidity. Investors have been unwilling to buy these securities from banks (even at a substantial discount) for fear of being stuck with the most "toxic" assets. Rightly so. A bank willing to sell mortgages for 25 cents on the dollar is likely holding the worst of the worst. As a result, these assets continue to remain on the balance sheet of the firms most in need of capital.

In the next subsection we discuss the literature as it relates to our model. The rest of the paper is broken into two parts. In Part I we study the model without liquidity shocks. We first present a stylized two-period example that illustrates several key points and provides intuition. Section 2 describes the model and relates lemons and signaling in a dynamic framework. The main results under the lemons condition are given in Section 3. Welfare and efficiency are discussed in Section 4. The results under the signaling condition are given in Section 5. We discuss alternative specifications and robustness in Section 6. Liquidity shocks are introduced in Part II. Section 7

[3] See for example Zweig (2008) or Cox (2008)

[4] See for example Reuters (2008) or Smith (2008)

describes the model with shocks. Section 8 characterizes equilibrium asset prices. We describe an algorithm to solve for the equilibrium and provide a numerical example in Section 9. Proofs are located in Appendix A. Appendix B contains results for the discrete-time analog of the model.

Related Literature

There is a fundamental strategic difference between Spence's market signaling and Akerlof's market for lemons. In Spence's model, sellers can commit to delay trade (and incur the costs from doing so) by attending school. In Akerlof's model, the choice is whether to trade *now or never*. In many markets, costly investment or signaling takes time to materialize; indeed the signaling variable in Spence's primary application is time spent in school. Moreover, rejecting an offer today rarely prevents a seller from trading in the future. In a dynamic environment (without commitment) these two strategic settings are virtually identical.

We are not the first to recognize some of the deficiencies of the aforementioned static models. A well-known objection raised by Weiss (1983) and Admati and Perry (1987) goes as follows: if only the high-ability students attend school, then attending for even a single day is sufficient to demonstrate high ability. Employers should hire students who show up on the first day rather than waiting for them to finish schooling. Of course, this is not an equilibrium either.

Nöldeke and Damme (1990) developed a convenient framework for analyzing a dynamic signaling model by splitting the game into many small periods, each of length Δ. They show that when firms offers are publicly observable, the unique equilibrium satisfying the never-a-weak-best-response refinement (Kohlberg and Mertens, 1986) is payoff equivalent to the least-cost-separating outcome. Swinkels (1999) argues that their result hinges on the combination of public offers and the refinement. He shows that when offers are kept private, the unique sequential equilibrium involves immediate trade with all types. Under the signaling condition and when news is completely uninformative, our model is the continuous-time analog to Swinkels. We confirm that trade is immediate and demonstrate how this result changes with the quality of the news.

Kremer and Skrzypacz (2007) introduce exogenous information into dynamic signaling model with private offers. In their model, a grade is revealed at some fixed time, provided that trade has not already occurred. In contrast to Swinkels, trade is always delayed with positive probability. A key insight of their work is that noisy information causes an *endogenous* lemons market to develop. In equilibrium, trade breaks down completely just prior to revelation of the grade. A similar result obtains in our model. However, in an infinite horizon model with gradual information arrival, the region of breakdown depends on market beliefs rather than time.

Hörner and Vieille (2008) address the issue of public versus private offers in a dynamic adverse selection model. They demonstrate that trade always [eventually] occurs when offers are private, while bargaining often ends at an impasse when offers are made publicly. Janssen and Roy (2002) take a Walrasian approach to dynamic adverse selection. They show that every equilibrium involves a sequence of increasing prices and qualities traded over time. Trade does not (necessarily) occur in every period. However, all goods are traded in finite time. In contrast to our model, the inefficiency does not disappear as the discount rate goes to one.

We provide a framework that accommodates both adverse selection and signaling and clarify their single point of differentiation in a dynamic environment. We then fully characterize the interaction between trade dynamics and information revelation in a continuous-time setting. This framework enables precise predictions and is ripe for extensions.

Another strand of literature has focused on extending the models of Akerlof and Spence to understand institutions of specific markets as well as the interactions between the markets for new and used goods. See, for example, Hendel and Lizzeri (1999,2002), Hendel et al. (2005), Johnson and Waldman (2003) and Taylor (1999). It is [by now] well understood that markets can and have overcome inefficiencies through a variety of innovations such as quality inspection, certification intermediaries and rental contracts.

The equilibrium in our model bears resemblance to that of two recent works on reputation. Bar-Isaac (2003) investigates learning and reputation in a dynamic signaling model where a privately informed monopolist decides whether to sell in each period. Lee and Liu (2008) consider a reputation game where in each period a short-lived plaintiff makes a take-it-or-leave-it settlement demand to a long-run defendant who is privately informed of her liability. The common thread is that the high type never succumbs when beliefs are low. This perseverance serves as an imperfect signal that boosts reputation because the low type mimics only with some probability.

Daley and Green (2008) study a static signaling model where the sender chooses a perfectly observable signal that also generates a stochastic grade. Under the signaling condition, the present paper is a dynamic analog of this model. One interpretation is that the dynamic model relaxes the assumption of a seller's ability to commit to delay trade until a fixed date t. From this standpoint, one could investigate the impact of the commitment assumption on welfare and on trading patterns. Despite their differences, a first pass shows a few similarities in the equilibria of the two models. In each there is cutoff prior below which the low type mixes between delaying trade and not, while the high type always chooses to delay. The payoffs to each type are constant over priors below the cutoff, with the low type's equal to her market value. Finally, delay is decreasing, and the expected value to both types is increasing, in the prior once it is above the cutoff. We would like to pursue this investigation further in future work.

Part I

1 Two-Period Example

Suppose that an entrepreneur's start-up company has just launched a beta version of a new technology. The new technology will make the company either a *hit* (high value) or a *letdown* (low value). The entrepreneur knows the type. The market of [potential] buyers does not. The company produces an i.i.d. cashflow X_t for $t = 1, 2, \ldots$ A *hit* generates $X_t = \$1$ with probability $q = 3/4$ and \$0 with probability 1/4. A *letdown* has probability $1 - q$ of generating \$1 and q of generating \$0. The market discount rate is $r_M = 10\%$. The entrepreneur discounts future payoffs at $r_E = 25\%$. It is efficient for the entrepreneur to sell the firm immediately, but an adverse selection problem exists that may prevent efficient trade.

There are two periods in which a sale can take place. In period $t = 0$, two buyers make *private* offers to the entrepreneur who chooses which offer to accept (if any). If the entrepreneur rejects both offers, the buyers exit the game and first-period cashflows are *publicly* realized. In period $t = 1$, two new buyers arrive, offers are made, and again the entrepreneur chooses to accept or reject. No future buyers arrive after period 1. The market value for a *hit* (*letdown*) is

$$\begin{aligned} V_H &= \frac{q}{r_M} = 7.5 \\ V_L &= \frac{(1-q)}{r_M} = 2.5 \end{aligned}$$

The entrepreneur has the outside *outside option* to never sell. The expected value from doing so is

$$\begin{aligned} K_H &= \frac{q}{r_E} = 3 \\ K_L &= \frac{(1-q)}{r_E} = 1 \end{aligned}$$

Two important relations are of note: (1) $V_\theta > K_\theta$, there is common knowledge of gains from trade, and (2) $K_H > V_L$, an entrepreneur with a *hit* would prefer her outside option rather than sell at the market value of a *letdown*. The structure of the equilibrium crucially depends on the market beliefs about the seller's type. Let p_0 denote the probability that buyers initially place on the seller having a *hit*, and let p_1 denote the belief at time $t = 1$.

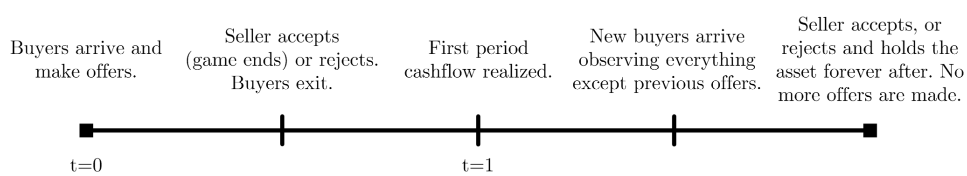

Figure 1: Timing of the game in the two-period example

Using backward induction, a classic market for lemons develops in period 1. If both types trade in period 1, then the market clearing price must be the expected value given p_1, $E[V_\theta|p_1] = p_1 V_H + (1-p_1)V_L$. A high-type seller is willing to trade if and only if the offer is greater than K_H. Thus if $p_1 \geq \underline{p} \equiv \frac{K_H - V_L}{V_H - V_L} = 0.1$, buyers offer $E[V_\theta|p_1]$, and both seller types accept. Otherwise only the low type trades at a price of V_L.

Moving back to period 0, three regions develop:

Market for Everything: When $p_0 > b = 0.548$ both seller types trade immediately at a price equal to expected market value.

Capitulation: When $p_0 < a = 0.125$ buyers offer V_L. A low-type seller mixes over accepting, and a high-type seller rejects.

No Trade: For intermediate $p_0 \in (a, b)$ there is complete market breakdown. No trade occurs at $t = 0$.

When p_0 is close to one, the high-type entrepreneur can get a price close to V_H by trading immediately. Realizing positive cashflows will have only a small effect on the price in the next period. The cost of delay outweighs the expected benefits. As p_0 decreases the expected benefit from good news increases while the cost of rejecting $E[V_\theta|p_0]$ decreases. At $p_0 = b$, the high type is just indifferent. The cutoff between the market for everything and the no-trade region, b, is defined implicitly by

$$bV_H + (1-b)V_L = (1+r_E)^{-1}\big[q(1 + E_b[V_\theta|X_1 = 1]) + (1-q)\max\{E_b[V_\theta|X_1 = 0], K_H\}\big]$$

where $E_q[\cdot]$ denotes the expectation starting from a prior of $p_0 = q$. The term on the left hand side represents the payoff from selling immediately at expected market value. The right hand side is the high type's discounted expected payoff from rejecting.[5] If $p_0 < b$, a high-type seller is unwilling to sell at average market value because the discounted expected future cashflows combined with the option to sell in period 1 is worth more. A low type would always happily accept expected market value. However, for all $p_0 > a$ her expected continuation value is larger than V_L. When $p_0 \in (a, b)$, any offer that is acceptable to a seller loses the buyer money on average. When $p_0 = a$ the low type is just indifferent between accepting V_L immediately versus taking her chances by rejecting. The lower cutoff is defined by the low types' indifference:

$$V_L = (1+r_E)^{-1}\big[(1-q)\cdot(1 + E_a[V_\theta|X_1 = 1]) + qV_L\big]$$

Below a, a low type accepts V_L with a probability such that the market updates its beliefs to a conditional on no trade. The equilibrium play is illustrated in Figure 2.

In period 1 (or a game with only one opportunity to trade), both types trade immediately for all $p > 0.1$, and the low type trades when $p < 0.1$ (see right panel of Figure 2). Adding another

[5] There is a unique $b \in (\underline{p}, 1)$ satisfying this equation. The left hand side is less than the right hand side at $\underline{p}$ but increases faster and is strictly larger at $b = 1$.

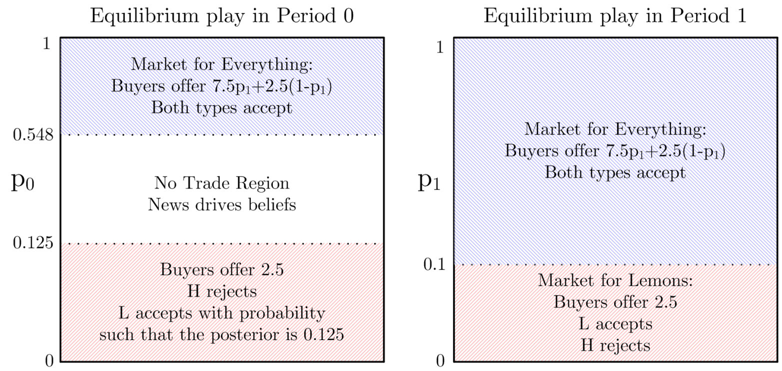

Figure 2: Characterization of equilibrium play as it depends on beliefs.

opportunity to trade and a noisy cashflow creates a region where the adverse selection problem is so severe that *neither* type of entrepreneur sells despite the fact that buyers are willing to pay a price in excess of K_H. Why? The high type's expected payoff from cashflows and higher prices in the next period exceeds the maximum price buyers are willing to pay, and the low type can do better than V_L by mimicking the high type. The option to sell in the future at higher prices drives both type sellers to reject offers even when they exceed their respective outside options.

Of course, this example is somewhat rudimentary and relies on several arbitrary assumptions. Restricting trade from occurring after period 1 is hardly different from doing so after period 0 (the static model).[6] However, the intuition for the no-trade region and the equilibrium dependence on the market belief will be key to understanding the fully dynamic model.

2 The Model

There is one seller holding an asset of type $\theta \in \{L, H\}$ and a market of [potential] buyers.[7] The seller knows her type while buyers do not. The game is played in continuous time, starting at $t = 0$ with an infinite horizon. Whoever is in possession of the asset receives a private flow payoff. The flow value is v_θ for buyers and k_θ for the seller. Any player not holding the asset receives a flow payoff of zero. All players discount future payoff streams at a common discount rate $r > 0$. Flow payoffs satisfy the following three conditions:

1. $v_\theta > k_\theta$: there is common knowledge of gains from trade.

2. $v_H > v_L$: buyers strictly prefer the high-value asset.

3. $k_H \geq k_L$: the seller weakly prefers the high-value asset.[8]

[6] The equilibrium also requires a condition on parameters to ensure that the high type is willing to trade at $E_a[V_\theta | X_1 = 1]$.

[7] We use "the seller holds an asset of type θ" interchangeably with "the seller is of type θ." Similarly, any reference to "buyers" or "the market" is equivalent.

[8] $k_H \geq k_L$ is not crucial for the main results provided that news is sufficiently informative, however it is implied by 1 and 2 under the lemons condition (see § 2.3).

Remark 2.1. *The results from our model remain unchanged under the following alternative specifications: (1) the flow payoff from an asset of type θ is the same for both buyers and the seller but the seller discounts future payoffs at a higher rate or (2) the flow payoffs are stochastic with mean k_θ to the seller and v_θ to buyers—the distinction is irrelevant for risk-neutral agents.*

Define the discounted present value of holding the asset *ad infinitum* as $V_\theta \equiv \int_0^\infty v_\theta e^{-rs} ds$ for a buyer and $K_\theta \equiv \int_0^\infty k_\theta e^{-rs} ds$ for the seller. If a type θ asset trades at time τ for a price w, the payoff (in time 0 dollars) to the seller is

$$\int_0^\tau e^{-rt} k_\theta dt + e^{-r\tau} w$$

Buyers begin discounting as soon as they arrive in the market. This is equivalent to saying that the seller's asset does not depreciate in value. This makes the game stationary and allows us to use a single state variable, the market belief about the seller's type, as a sufficient statistic for the history of the game. After the initial sale, the asset will not be re-traded because the winning buyer immediately learns the asset type following a purchase (Milgrom and Stokey, 1982). The payoff (in time τ dollars) to the buyer is

$$\int_0^\infty e^{-rt} v_\theta dt - w$$

At every time t, two or more short-lived buyers arrive and simultaneously make private offers. The seller either accepts one of the offers or rejects all. If she accepts, trade occurs at the accepted price, and the rights to all future flow payoffs are transferred to the buyer. If she rejects all offers, the seller retains the asset, consumes the flow payoff and can entertain offers from future buyers. A buyer whose offer is rejected exits the game permanently. In addition, news about the seller's type is gradually revealed—as we discuss in immediately detail. A heuristic description of the timing is depicted in Figure 3.

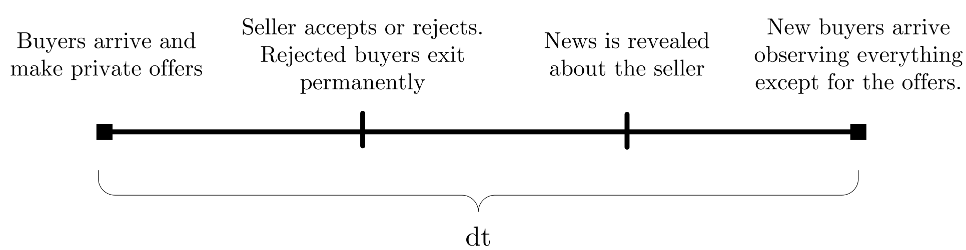

Figure 3: Heuristic timeline of a single "period"

2.1 Information Environment and News Arrival

Having specified what players can do and their payoff functions, we need to specify what they know at each time t. The seller knows θ as well as the entire past history of the game. Buyers do not know θ, but start with a common prior $P_0 = \Pr_{t=0}(\theta = H)$. Further, they do not observe the past (rejected) offers made by previous buyers.[9] What they do know is whether the seller has traded and the realized path from an exogenous stochastic process.

News about the seller's asset is revealed via a Brownian diffusion process. Both types start with the same initial score X_0. The score process then evolves according to the following stochastic differential equation:

$$dX_t^\theta = \mu_\theta dt + \sigma dB_t \tag{1}$$

where $B = \{B_t, \mathcal{F}_t, 0 \leq t \leq \infty\}$ is standard Brownian motion on the canonical probability space $\{\Omega, \mathcal{F}, \mathcal{Q}\}$. Without loss of generality, $\mu_H \geq \mu_L$, meaning that the high type expects to have a higher score than the low type at any point in time. Define the signal-to-noise ratio $s \equiv (\mu_H - \mu_L)/\sigma$. When $s = 0$, the news is completely uninformative. For the remainder we assume that $s > 0$, unless otherwise stated. The larger is s, the faster the news arrives.[10] At each time τ, all players observe the entire history of news, $\{X_t : 0 \leq t \leq \tau\}$. At any time τ the market belief about the seller is conditioned on the revealed news about her $\{X_t : 0 \leq t \leq \tau\}$ and on the fact that the seller did not trade at any time $t < \tau$.

Remark 2.2. *The results of our model remain unchanged under the following alternative specification: the flow payoff to the asset holder* ***is*** *the publicly observable news process, but the seller has a higher discount rate, in which case $\mu_H > \mu_L$.*[11]

Let $\hat{P}$ denote the belief process for a Bayesian who updates *only based on news.* Given an initial prior $\hat{P}_0 = \Pr_{t=0}\{\theta = H\}$

$$\hat{P}_t = g(t, X_t) \equiv \frac{\hat{P}_0 e^{-(X_t - \mu_H t)^2/2\sigma^2 t}}{\hat{P}_0 e^{-(X_t - \mu_H t)^2/2\sigma^2 t} + (1 - \hat{P}_0) e^{-(X_t - \mu_L t)^2/2\sigma^2 t}} \tag{2}$$

Applying Ito's formula to g gives the following stochastic differential equation

$$d\hat{P}_t = \frac{\partial}{\partial t} g(t, X_t) dt + \frac{\partial}{\partial x} g(t, X_t) dX_t + \frac{1}{2} \frac{\partial^2}{\partial x^2} g(t, X_t)(dX_t)^2 \tag{3}$$

Inserting the law of motion from (1) gives a probabilistic representation of how beliefs evolve from the perspective of each type based solely on the news, which we denote by $\hat{P}^\theta$. Because X^θ is Markovian, and Bayesian updating depends only on the prior, $\hat{P}$ is also a Markov process.[12] $\hat{P}$ is a martingale because in expectation beliefs must be correct and the expected belief tomorrow is

[9] This eliminates the possibility of signaling through rejection of high offers (Nöldeke and Damme (1990), Hörner and Vieille (2008)). We discuss the implications of relaxing this assumption in § 6.

[10] We demonstrate that s is a sufficient statistic for the rate of information arrival. That is, if two triples $(\mu_H, \mu_L, \sigma) \neq (\mu'_H, \mu'_L, \sigma')$ but $\frac{\mu_H - \mu_L}{\sigma} = \frac{\mu'_H - \mu'_L}{\sigma'}$ then the two models are equivalent up to the low type's flow rate of acceptance at the lower boundary (see § 3.1).

[11] In this case, the No-Trade Theorem (Milgrom and Stokey, 1982) does not apply after the initial sale because the new owner has the same information about the asset as all arriving buyers. However, it is common knowledge there are *no gains* from trade.

[12] In Appendix A, we make a change of variables in order to get a process that is one-to-one with beliefs but follows arithmetic brownian motion and is more convenient to work with.

the same as the belief today. However, the high type expects to receive good news, hence $\hat{P}^H$ is a submartingale. The low type is expectant of bad news; $\hat{P}^L$ is a supermartingale.

2.2 Strategies and Equilibrium Concept

We restrict attention to equilibria in which players' strategies are stationary in the market's belief about the seller.[13] Let $P = \{P_t, 0 \leq t < \infty\}$ denote the process representing buyers' beliefs in equilibrium and denote an outcome of this process by $P_t(\omega)$, where ω denotes a realization of a state in Ω. In equilibrium, P must be consistent with the players' strategies as well as the news arrival. P differs from $\hat{P}$ because it accounts for the possibility of trade by one or more types before time t. We use p when referring to the state variable as opposed to the stochastic process P. The reader should interpret p as any pair (t, ω) such that $P_t(\omega) = p$ given no trade before time t.

A stationary strategy for a buyer is a function mapping a belief to an offer. We restrict buyers to play pure strategies only to simplify exposition. Trade dynamics remain unchanged if we allow buyers to play mixed strategies.[14] Let $w(p)$ denote the maximum of all offers made at p. A [general] strategy for the type θ seller is a stochastic process $S^\theta = \{S_t^\theta, 0 \leq t \leq \infty\}$ adapted to the filtration $\mathcal{H}_t \equiv \sigma(\{w_s, P_s, X_s\}, 0 \leq s \leq t)$, where $\sigma(y)$ is the σ-algebra generated by the random variable y.[15] The process must satisfy: (1) $S_0 = 0$, (2) $S_t \leq S_{t'} \leq 1$ for all $t \leq t'$ and (3) S_t is right-continuous. In words, the strategy process keeps track of how much probability mass the seller has "used up" by accepting past offers. An upward jump in S corresponds to the seller accepting with an atom of mass. S increasing continuously corresponds to the seller accepting at a flow rate. For any sample path of $\{w_t, P_t, X_t\}$, S^θ is merely a CDF over the seller's acceptance time (unless she *never* accepts).

Stationarity requires that $S_{t'} - S_t$ is independent of $\mathcal{H}_t$ for all $t < t'$.[16] If the type θ seller accepts with an atom of probability when the market belief is p and the current best offer is w, we will denote the probability of acceptance as $\rho_\theta(p, w)$ (that is, ρ is the proportion of the seller's

[13] Beliefs are the natural choice for a state variable since the prior is a sufficient statistic for Bayesian updating. In other words, if buyers believe the seller holds a high-type asset with probability p at time t, how buyers will update those beliefs depends only on p and their belief about the seller's strategy.

[14] See footnote 23.

[15] We have ruled out any dependence on non-maximal offers. Because offers are private, and buyers' offers depend only on p, any conditioning on non-maximal offers done by the seller only substitutes for private randomization.

[16] Recall that in discrete time a stationary strategy is defined simply as a mapping from (θ, w, p) into a probability of acceptance (refer to this as *D-T stationarity* and the restriction in the body above as *C-T stationarity*). First, notice that in a discrete-time game the two notions of stationarity are equivalent. However, in continuous time, D-T stationarity is more restrictive than C-T stationarity. One could restrict the strategy of the seller in our continuous-time game by imposing D-T stationarity. However, we argue D-T stationarity is an overly restrictive notion of stationarity in continuous time. At a conceptual level, we interpret stationarity to restrict behavior to depend only on "today's" information—regardless of how it was reached. C-T stationarity is the precisely formulated statement of this concept. The difference in continuous time is that strategies that depend on the increment $d\hat{P}_t$ are measurable with respect to $\mathcal{H}_t$ (today's information). There is no analog of this information available to the player in a discrete-time game. Therefore, a strategy that depends only on information deemed permissable by D-T stationarity may be inconsistent with the seller's use of any information that *does not pertain to past history.* We will also see that the unique C-T stationary equilibrium does not satisfy D-T stationarity, despite it being the limit of the unique D-T stationary equilibrium in the discrete-time game in Appendix B. Because C-T stationarity is weaker than D-T stationarity, there does not exist an equilibrium to our game that satisfies D-T stationarity.

remaining mass that the atom uses—not the absolute size of the atom). Finally, the news process and flow payoffs are not verifiable and cannot be contracted upon.[17]

Given w and P, the seller faces an optimal stopping problem. We use F_θ to denote the value function for a seller of type θ. The Bellman equation for the seller's problem along the equilibrium path is

$$F_\theta(p) = \max\left\{w(p), k_\theta dt + e^{-rdt}E^\theta\left[F_\theta(p + dP_t)\right]\right\} \tag{SP}$$

The seller chooses between accepting the current offer today or taking her flow payoff and waiting in hopes of a higher offer in the future. Because offers are private, the level of offer rejected will have no effect on future payoffs. This observation implies that the seller will follow a reservation strategy: for a given p, a type θ seller will accept all offers greater than $R_\theta(p)$ and reject all offers less. We now formally define the equilibrium concept.

Definition 2.1. *An equilibrium of the game is a quadruple (S^L, S^H, w, P), such that*

1. *Given w and P, S^θ solves the type θ seller's problem.*
2. *Given S^L, S^H and P, w is consistent with buyers playing best responses.*
3. *P satisfies Bayes' rule whenever possible and is consistent with (S^L, S^H, w) in the sense of sequential equilibrium.*[18]

In addition, we restrict any off-equilibrium-path belief processes to be stationary and satisfy *belief monotonicity*, defined as follows. Fix a candidate equilibrium and denote its *on-path* belief process as P^*. Define $\widetilde{P}$ to be the belief process that "ignores" off path events: in any state p where the probability of trade is less than one $\widetilde{P}$ evolves according to $d\widetilde{P} = dP^*$, and in any state p such that trade was supposed to commence with probability one $\widetilde{P}$ evolves according to $d\widetilde{P} = d\hat{P}$.

Definition 2.2. *Fix a candidate equilibrium and an arbitrary (ω, t) such that $S^H_t(\omega) = S^L_t(\omega) = 1$. The equilibrium satisfies **belief monotonicity** if and only if $P_{t'}(\omega) - P_t(\omega) \geq \widetilde{P}_{t'}(\omega) - \widetilde{P}_t(\omega)$ for all $t' > t$ if trade has not occurred by time t'.*

The sequential equilibrium notion of belief consistency ensures that all buyers share the same off-path beliefs. Without restriction on the off-path belief process many unappealing equilibria can be sustained by "threat beliefs" (e.g. a seller who does not accept immediately is considered to be a low type with probability one). Belief monotonicity restricts off-equilibrium-path beliefs in the following way. Fix a sample path ω and a candidate equilibrium that calls for trade to occur with probability one by time t (given ω). The equilibrium satisfies belief monotonicity if and only if beliefs do not put more weight on the low type solely based on the seller's failure to trade.

[17]The most obvious reasons for this assumption are the "standard" issues with verification and contractibility. There are others. In many applications, the news ends after trade occurs. For example, when a student leaves school to enter the workforce she no longer receives grades. Even if news is contractible, there are reasons why both parties may agree not to do so. For example, if the seller can manipulate the news creating a moral hazard problem.

[18]Technically, the consistency requirement of Kreps and Wilson (1982) applies only to finite-horizon games with finite action sets. We make an extension of this concept in order to apply it to our game.

A behavioral justification can be offered. The high-type seller has a higher expected future payoff from rejecting because she has higher flow payoffs and is more likely to get good news. The high type has more to gain by rejecting. Therefore, buyers should not revise their belief downward upon observing an unexpected rejection. Belief monotonicity is intimately related to the *Divinity* refinement of Banks and Sobel (1987).[19]

2.3 Dynamic Lemons and Dynamic Signaling

In Spence's canonical static signaling story a student is privately informed about her ability and may acquire costly education in an attempt to signal this ability to potential employers. Under the usual single-crossing assumption, the only stable outcome is the least-cost-separating equilibrium.[20] To address the objection raised by Weiss (1983) and Admati and Perry (1987), we have reformulated the game in a fully dynamic context; students incur a flow cost while in school, but are free to leave for employment at any instant in time.

Now consider Akerlof's canonical market for lemons. A used-car owner is privately informed about the quality of her car and faces a market of potential buyers. If the market belief about her car is sufficiently pessimistic, the owner of a high-value car prefers to retain it for use rather than sell at the market price. When beliefs are unfavorable only lemons are traded in the market. A very similar objection can be raised. One can interpret the model as, "There's a market for used cars. It occurs on Monday." The reader should be compelled to ask, "What happens on Tuesday?" If all the low-value cars are sold on Monday only high-value ones remain. The price would sky-rocket on Tuesday. Again, this cannot be an equilibrium. We have reformulated the game in a fully dynamic context; potential sellers drive their cars until they sell them, which they are free to do at any instant in time.

The static models and predictions of Spence and Akerlof are quite different. The privately informed player's action in Spence's model is a time, $t \in \mathbb{R}_+$. Each time is associated with a type-dependent cost. While in Akerlof's, the decision is binary—sell now or never. However, both models are susceptible to the same criticism: failing to incorporate dynamics. Once the dynamics of the environment are formalized the two models are virtually identical. Each fits within our general model and is separated by a single inequality on parameters.

Definition 2.3. *From this point on, define*

- *The* ***lemons condition****:* $v_L < k_H$*, and*
- *The* ***signaling condition****:* $v_L \geq k_H$

[19]One way to see this is to decompose P into two processes: $\hat{P}$ and $\hat{Q}$. $\hat{P}$ tracks the updating of beliefs based solely on the revealed news, while $\hat{Q}$ tracks the updating of beliefs passed on the equilibrium inference that the seller has not yet traded. *Belief monotonicity* then requires that $\hat{Q}$ is weakly greater off the equilibrium path than on it. This is a natural extension of *Divinity* to continuous time.

[20]In the two-type case, only the high type attends school, and does so for long enough to discourage the low type from imitating.

In words, if the buyers were convinced that the seller was the low type, would the high-type seller be willing to sell? Under the signaling condition he would. Under the lemons condition he would not.[21] This is the extent of the difference in a dynamic framework. Not only do both environments fit within our model, their equilibria have the same structure when the news is sufficiently informative.

2.4 An Optimal Stopping Problem: The Ψ-game

Here we present a stopping problem that will play an important role in understanding the equilibrium of the game. Suppose that a type θ seller faces a "naive" market that offers $\Psi(p) \equiv E[V_\theta|p]$ at every (t,ω) such that $\hat{P}_t(\omega) = p$. Define $h_\theta(t,p) \equiv \int_0^t k_\theta e^{-rs}ds + e^{-rt}\Psi(p)$ to be the payoff to a type θ seller from stopping at time t given the state p. The problem facing the seller is to find an optimal policy (a stopping time) τ_θ for $\hat{P}_t$ to maximize her expected payoff given any initial state.

$$E_p^\theta\left[h_\theta(\tau_\theta, \hat{P}_{\tau_\theta})\right] = \sup_\tau E_p^\theta\left[h_\theta(\tau, \hat{P}_\tau)\right] \qquad \text{for all } p \in [0,1]$$

We refer to this problem as the Ψ-game. The optimal policy for a low-type seller is to accept Ψ immediately because $\hat{P}^L$ trends downward (implying a lower expected offer tomorrow), and in the meantime she forgoes a flow payoff of $r\Psi(p) > k_L$.

On the other hand, a high-type seller can benefit from waiting. The form of the optimal policy for a high-type seller depends on whether the lemons (or signaling) condition holds. Under the lemons condition, it can never be optimal to stop when $\Psi(p) < K_H$. The marginal benefit to the high type from delaying trade for an instant is single peaked in p. It is always positive at low beliefs since the high type prefers her outside option. However, as p increases so too does the cost of foregoing $\Psi(p)$ and this eventually outweighs the high type's expected increase in price. As s increases, the peak converges to the belief where good news has its largest effect.

Lemma 2.1. *Under the **lemons condition**, there exists a unique $p_H^* \in (0,1)$ such that the optimal policy for the high type in the Ψ-game is to continue for all $p < p_H^*$ and to stop for all $p > p_H^*$. That is $\tau_H = \inf\{t : \hat{P}_t > p_H^*\}$.*

Under the signaling condition, $\Psi(p) > K_H$ for all p. In this case, the optimal policy stops for p below a lower threshold as well as above an upper threshold.

Lemma 2.2. *Under the **signaling condition**, there exists a unique pair $(\underline{p}_H, \overline{p}_H)$ such that the optimal policy for the high type in the Ψ-game is to continue for all $p \in (\underline{p}_H, \overline{p}_H)$ and stop at the first time such that $\hat{P}_t \notin (\underline{p}_H, \overline{p}_H)$. That is $\tau_H = \inf\{t : \hat{P}_t \notin (\underline{p}_H, \overline{p}_H)\}$.*

The benefit from delaying trade for an instant of time is shown in Figure 4 for both cases. The only significant difference is that under the signaling condition, the optimal policy is to accept when p is close to zero.

[21] Given a prior p_0, trade breaks down in the static model when $k_H > E[v_\theta|p_0]$. The lemons condition implies that there exists a p_0 such that trade breaks down. We define it in this way because we will characterize the equilibrium for all possible beliefs.

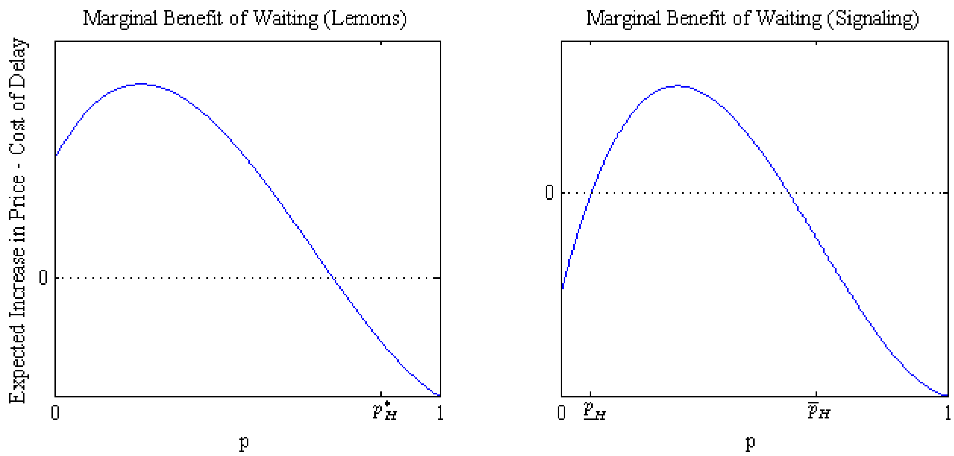

Figure 4: Marginal benefit to the high type of delaying trade for an instant of time in the Ψ-game under the lemons condition (left-panel) and the signaling condition (right-panel).

3 The Dynamic Market for Lemons

The following theorem characterizes the equilibrium trade dynamics under the lemons condition.

Theorem 3.1. *There exists an essentially unique equilibrium under the lemons condition. It is characterized by a unique pair of belief levels* $0 < a < b < 1$ *such that:*

- *For all* $p \geq b$, $w(p) = E[V_\theta|p]$ *and both types trade with probability one.*
- *For all* $p \leq a$, $w(p) = V_L$, *the high type rejects and the low type accepts with probability* $\rho_L(p, V_L) \equiv \frac{a-p}{a(1-p)} \in [0, 1)$.
- *For all beliefs* $p \in (a, b)$ *the buyers make non-serious offers, which both types reject.*

The uniqueness is qualified with "essentially" because it is unique up to strategies over a measure-zero set of beliefs below a and above b. However, all equilibria are payoff equivalent state-by-state; the expected payoff to a player given the belief p is identical in all equilibria. A complete characterization requires specifying equilibrium beliefs. On-path beliefs are discussed in § 3.1. The only off-path behavior is a rejection at some $p \geq b$. In this case, we specify that following an unexpected rejection at time t, P follows $\tilde{P}$.[22]

When beliefs are very favorable the seller has little to gain and a high cost from waiting. Trade occurs immediately at expected market value. As beliefs become less favorable the market shuts down; buyers wait for more news before making serious offers. In this region, the high type will not accept expected market value because the combination of her flow payoff and the option value of trading in the future is more attractive. On the other hand, the low type would be happy to accept $E[V_\theta|p]$; however the combination of *her* flow payoff and the option to trade in the future is more attractive than an offer of V_L. Any offer that would be accepted loses the buyer money. When

[22] If an offer is unexpectedly accepted, than any off-path beliefs that do not jump to $p = 1$ are permissable.

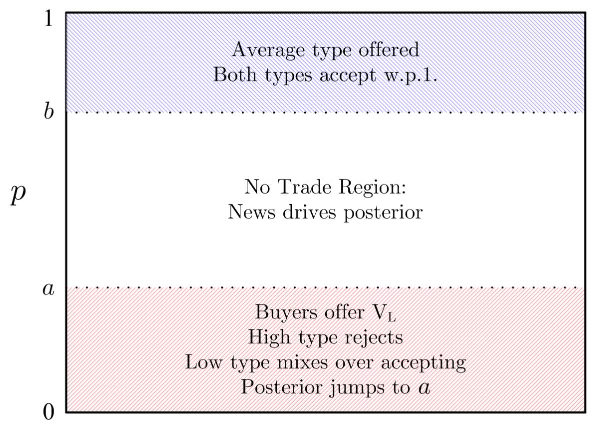

Figure 5: Illustration of the Equilibrium

the belief is low enough, the low type's option value decreases to the point where she is willing to accept V_L instead of delaying trade. The belief where she is just indifferent occurs at a. For $p < a$ the low-type seller mixes between accepting and rejecting V_L in a way such that, conditional on not observing trade, the market belief jumps instantaneously to a.[23] In economic terms, not selling when the market is pessimistic is an imperfect signal of high value.

The value functions are depicted in Figure 6. Both are constant for all $p \leq a$ due to the instantaneous jump in beliefs. F_H is greater than expected market value, and F_L is greater than V_L in the no-trade region. Both value functions are exactly equal to expected market value above b when trade occurs immediately.

3.1 Equilibrium Beliefs

The characterization of the equilibrium in Theorem 3.1 implies that a low-type seller trades with probability zero at $p = a$. This is true; the low type cannot trade with an atom at a because then beliefs would instantaneously jump upward, in which case the low type would prefer to reject. On the other hand, if the low-type seller never traded at the lower boundary then beliefs would sometimes drift below a, which takes strictly positive time. This imposes a cost to a low type, causing F_L to drop below V_L and creating a profitable opportunity for buyers. The following proposition precisely demonstrates that the low type trades at a rate proportional to the arrival of news at the lower boundary.

[23] It is unimportant whether the mixing comes from the seller or buyers. It is only important that the mixing results in an instantaneous jump to a conditional on rejection. We state it in this way because we have innocuously assumed that buyers play pure strategies. In addition, consider a belief p_0 such that the low trades at a price of V_L with probability between 0 and 1. Conditional on rejection, P_t jumps from p_0 to $j(p_0) \equiv \frac{p_0}{p_0+(1-p_0)\rho_L(p_0,V_L)}$. It is possible that the low type trades with positive probability at $j(p_0)$ as well, causing P_t to jump discontinuously again, and then again, etc, all happening in an infinitesimally small amount of time. To deal with this technicality, we impose the following restriction. *Fix a price y and let $\mathcal{P}_y$ be the set of p-values such that trade occurs at a price of y with positive probability when beliefs are p. For an arbitrary $p' \in \mathcal{P}_y$, let $j(p')$ be the updated market belief conditional on no trade at p'. We restrict that $j(p') \notin \mathcal{P}_y$.* Simply put, we restrict jumps in P_t to occur in "one shot." This restriction is not necessary in a discrete-time setting.

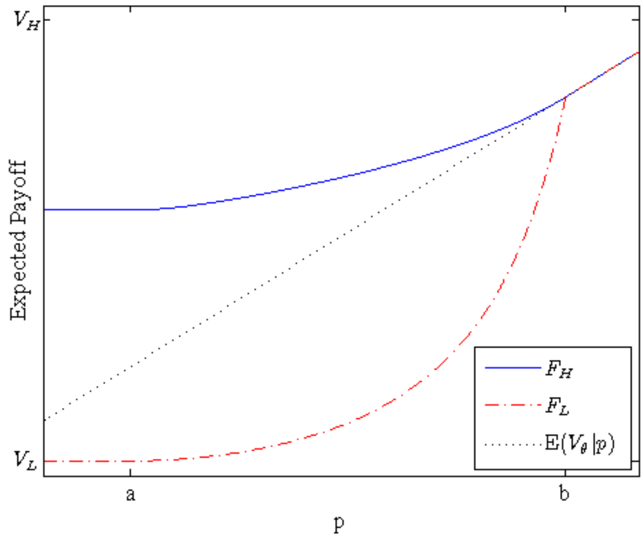

Figure 6: Equilibrium Value Functions

Lemma 3.1. *The low-type seller trades with a flow probability $\kappa = s/\sigma$ proportional to dX_t at the lower boundary. More specifically, starting from any $P_0 \in [a, b)$, the probability that the low type accepts V_L before time t is $1-e^{-\hat{Y}_t}$ where $\hat{Y}_t \equiv \max(\alpha - \hat{M}_t, 0)$, $\hat{M}_t \equiv \inf_{0 \le s \le t} \hat{Z}_t$, $\hat{Z}_t \equiv \ln(\hat{P}_t/(1-\hat{P}_t))$ and $\alpha \equiv \ln(a/(1-a))$.*

Higher quality information makes it more likely that the low type trades at the lower boundary in two ways; the low type will hit the lower boundary with higher probability, and upon doing so she trades at a faster rate.

Corollary 3.1. *Starting at $P_t = a$, the probability that the low type trades at a price of V_L before time $t + \Delta$ is approximately $s\sqrt{\Delta}$ for Δ small. Conditional on $P_t = x \in (a, b)$, the probability that the low type eventually reaches b and sells at a price of $\Psi(b)$ is $\frac{x(1-b)}{b(1-x)}$.*

When trade does not occur at a, the low type's acceptance strategy prevents the equilibrium beliefs from dropping below a. The belief process *reflects* off the lower boundary. By rejecting offers when beliefs are unfavorable, the high-type seller signals her value to the market.

Proposition 3.1. *Prior to trade, the equilibrium belief process has a reflecting boundary at $p = a$. Formally, $P_t = \hat{P}_t + \hat{Q}_t$ where $\hat{Q}(t) = \max(a - \inf_{0 \le s \le t} \hat{P}_t, 0)$.*

Figure 7 depicts a discrete-time analog based on a random walk news process. At time zero, beliefs are $p = a$. In the next short period of time (Δ), either good news or bad news will be revealed about the seller. If good news arrives, Bayesian updating leads to $P_\Delta = a^+ > a$. If bad news arrives, updating leads to $P_\Delta = a^- < a$. At $P_\Delta = a^-$ the low type accepts with probability $\frac{a-a^-}{a(1-a^-)}$. Conditional on rejection, beliefs jump back to a. As $\Delta \to 0$, a^- and a^+ converge toward a, which becomes a reflecting boundary for the continuous-time process P_t until trade occurs.[24]

[24]Reflecting boundaries are often constructed by having the process *reflect* off a if it tries to move below it. That is,

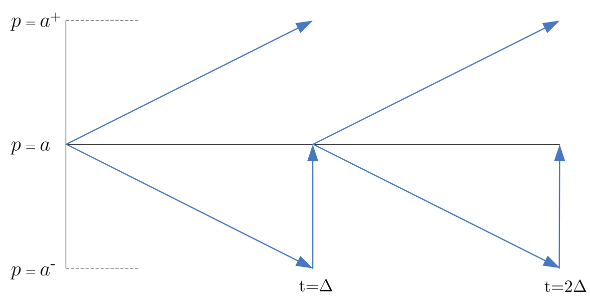

Figure 7: The Reflecting Boundary

Since the high type never accepts at the lower boundary, P^H is purely reflective at a; a high type will eventually reach the upper boundary with probability one. The low type will accept V_L at a with positive probability, in which case the market immediately learns the asset is low value. The process is then absorbed; equilibrium beliefs immediately jump to $p = 0$ where they remain forever after. Figure 8 shows a sample path of equilibrium behavior. In the left panel the asset value is low. The seller rejects V_L at $p = a$ until time τ at which point she accepts. The right panel shows the same path up to τ at which point the high type keeps rejecting at a until eventually reaching b.

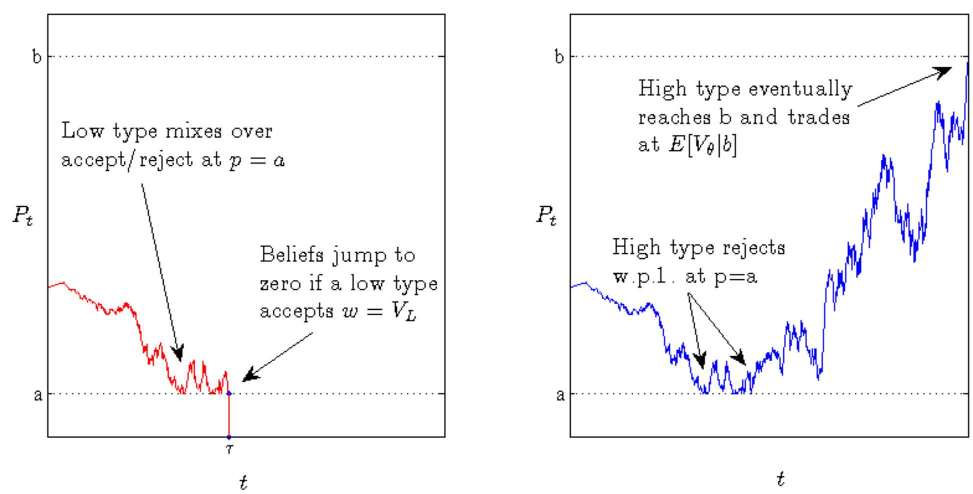

Figure 8: Low type may eventually trade at $p = a$ (left panel), a high type never does (right panel).

3.2 Proof of Main Theorem

The proof of Theorem 3.1 consists of three parts. We begin with a few simple properties that must hold in any equilibrium for any s. Next, we argue that, under the lemons condition with $s > 0$,

if $P_0 = a$ then $P_\Delta = a^+$ regardless of the news arrival. The distribution of these two discrete-time process converges to the same continuous-time limit. However, only the one described in Figure 7 comports with the economics of the situation. See Appendix B for more on the discrete-time model.

the equilibrium must be of the form in Theorem 3.1 for some (a, b). Then we show there is a unique candidate pair (a, b) by solving a pair of differential equations and verify it constitutes an equilibrium of the game.

Lemma 3.2. *In any equilibrium:*

1. *Buyers make zero profit.*

2. $F_L(p) \geq V_L$ *and* $F_H(p) \geq \Psi(p)$ *for all* p.

3. $F_L(p) \leq \Psi(p)$ *and* $F_H(p) \leq V_H$ *for all* p.

4. *For any* p*, there are only two prices at which trade can occur:* V_L *and* $\Psi(p) \equiv E[V_\theta|p]$.

Proof. The zero-profit condition is immediate from the Bertrand nature of the buyers' competition for the seller's asset (similar to Swinkels (1999)). If (2) fails then a buyer can make a profitable offer that will be accepted, violating (1). The total potential value in the game given the state p is $\Psi(p)$, and the amount going to the seller is $pF_H(p)+(1-p)F_L(p)$. The zero profit condition implies that $pF_H(p) + (1 - p)F_L(p) \leq \Psi(p)$, otherwise buyers would be subsidizing the seller and losing money. Combining this with (2) implies (3). For (4), a high type will never accept $w(p) < \Psi(p)$, and therefore (1) implies that any trade that occurs at a price strictly less than $\Psi(p)$ must take place at V_L. From (3), if $w(p) > \Psi(p)$ is offered then the low type will accept with probability one, and the offer loses money regardless of the high type's strategy. ■

Lemma 3.3. *The low-type seller does not trade with a flow probability over any interval in* $[0, 1]$.

Proof. Suppose that the low type accepts with a positive flow rate for almost all p in (p_1, p_2). Then $F_L(p) = V_L$ for all $p \in (p_1, p_2)$. Consider the low type's decision at $p = \frac{1}{2}(p_1 + p_2)$. If she rejects, P will evolve continuously. Therefore, over the next small interval of time she will receive flow payoff $k_L < v_L$, and beliefs will reach some p' such that $F_L(p) = V_L$. It follows that she does strictly better by accepting V_L at p. ■

Lemma 3.4. *Suppose there exists a no-trade interval with endpoints* (a, b)*. There are only three possible values for* $\{F_L(a), F_L(b)\}$*: (1)* $\{V_L, \Psi(b)\}$*, (2)* $\{\Psi(a), \Psi(b)\}$*, or (3)* $\{\Psi(a), V_L\}$.

Proof. Because a is the endpoint of a no-trade region either trade occurs at a or everywhere on an interval $(a-\varepsilon, a)$. By Lemma 3.2, the trading price must be either V_L or $\Psi(a)$. The same holds for b. However, if $F_L(a) = F_L(b) = V_L$ then there exists $p \in (a, b)$ such that $F_L(p) < V_L$, contradicting Lemma 3.2. The lemma follows. ■

Proof of Uniqueness of Form. The proof proceeds in several steps.

Step 1. *There exists a belief* $\underline{a} > 0$ *such that, for almost all* $p < \underline{a}$*:*

1. $w(p) = V_L$

2. $\rho_L(p, V_L) \in (0, 1)$

3. $\rho_H(p, V_L) = 0$

Define $\underline{p}$ implicitly by $\Psi(\underline{p}) = K_H$. The lemons condition is that $\underline{p} > 0$. Therefore, by Lemma 3.2 the high type will never trade at beliefs $p < \underline{p}$. Suppose that the low type never trades at $p < \underline{p}$. Then in equilibrium, P_t evolves in accordance with $\hat{P}_t$ when $P_t < \underline{p}$. From the buyers' individual rationality, $F_L(\underline{p}) \leq V_H$ (this holds for all p). Fixing any $F_L(\underline{p}) \leq V_H$, the limit of $F_L(p)$ as $p \to 0$ is $K_L < V_L$ violating Lemma 3.2. Hence, there exists a $\widetilde{p} < \underline{p}$ such that when beliefs are $\widetilde{p}$ the low type trades with a positive probability at a price of V_L. $F_L(\widetilde{p}) = V_L$.

The continuity of $\hat{P}$ implies that $F_L(p) \leq F_L(\widetilde{p}) = V_L$ for all $p < \widetilde{p}$. Lemma 3.2 then implies that $F_L(p) = V_L$ for all $p < \widetilde{p}$. Define $\underline{a}$ as $\sup\{p : F_L(p') = V_L$ for almost all $p' \leq p\}$. The only equilibrium behavior consistent with this value function is that the low type trades at a price of V_L with positive probability for almost all $p \leq \underline{a}$, and that the high type does not trade for any $p < \underline{a}$.

Consider a generic $p < \underline{a}$ where the low type trades at a price of V_L with positive probability. To see that this probability must be less than one, suppose it were not. Then P_t jumps from p to 1 instantaneously and trade will occur at a price of $V_H > V_L$. The low type should therefore reject V_L at p, producing a contradiction.

Step 2. *There exists an interval* (a, b)*,* $0 < a < b < 1$*, such that trade never occurs when* $p \in (a, b)$*.*

Consider a p_0 less than $\underline{a}$. In zero time, the low type will trade at a price of V_L with interior probability. Hence, conditional on no trade, P_t will discontinuously jump to some $j(p_0) \equiv \frac{p_0}{p_0 + (1-p_0)\rho_L(p_0, V_L)} > p_0$. $j(p_0)$ must belong to an interval $[p_1, p_2]$ such that for almost all $p \in [p_1, p_2]$ either no trade occurs at p or $\Psi(p)$ is offered (see footnote 23). Because the low type is indifferent at p_0, $F_L(j(p_0)) = V_L$. $F_L(p) \geq V_L$ for all p then implies that there must be zero probability that the low type trades at a price of $\Psi(p)$ for any $p > \underline{p}$ in a small positive interval of time. This establishes the claim.

Step 3. $\underline{a}$ *is the lower bound of a no-trade region.*

We demonstrate first that F_L is continuous at $\underline{a}$, then that this is sufficient for the claim. Suppose F_L were discontinuous from the left at $\underline{a}$. That is, suppose $F_L(\underline{a}) > V_L$. The continuity of the $\hat{P}$ process immediately rules that this discontinuity is due to the existence of a no-trade region beginning at $\underline{a}$. Hence, when beliefs are $\underline{a}$, $\Psi(\underline{a})$ is offered and accepted by both types. As we next show, this is not optimal for the high type, yielding a contradiction.

For convenience, we use dP to denote $\int_0^t dP_t$ for arbitrarily small t. There are two cases: (1) there is positive probability that $\underline{a} + dP < \underline{a}$, or (2) there is not. Take the first case first and let F_H^o be the high type's value function off the equilibrium path resulting from the proposed deviation. If the next dP is positive, $F_H^o(\underline{a} + dP) \geq \Psi(\underline{a} + dP) > \Psi(\underline{a})$.[25] If the next dP is negative, then P will drift below $\underline{a}$. *On path* the low type would have traded with positive probability at $\underline{a} + dP$

[25] Sequential rationality of the buyers implies that $F_H^o(p) \geq \Psi(p)$ for all p by the logic given in the proof of Lemma 3.2 for the analogous on-path result.

causing $\widetilde{P}$, and therefore P by *belief monotonicity*, to discontinuously increase. We argued above that on-path beliefs must jump from any p_0 below $\underline{a}$ to a belief $j(p_0)$ bounded away from any belief p where $\Psi(p)$ is offered. Hence, $j(p_0) \gg \underline{a}$, and $F_H^o(p_0) = F_H^o(j(p_0)) \geq \Psi(j(p_0)) \gg \Psi(\underline{a})$. This implies that by rejecting, the high type has a positive probability of a discretely higher payoff in an arbitrarily small amount of time, and zero probability of a lower payoff. It follows that she should reject.

For case (2) it matters whether there exists and $\varepsilon > 0$ such that $F_L(p) = V_L$ for all points in $(\underline{a}, \underline{a} + \varepsilon)$ or not. If so, then the argument given in case (1) applies; the only difference is that the discrete jump in payoffs comes from dP being positive—which, by assumption occurs with probability one under (2). If there does not exist such an ε, then (2) implies that $\underline{a}$ is a reflecting boundary of P^H off the equilibrium path. Hence, rejection by the high type lands him in an optimal stopping problem similar to the Ψ-game, but with a reflecting lower boundary, values $F_H^o(p) \geq \Psi(p)$ and a belief process always at least as favorable as $\hat{P}$. It is immediate that the solution to this problem must involve rejection by the high type at the reflecting boundary $\underline{a}$ (Harrison, 1985). Having covered both cases, we have established that F_L is continuous from the left at $\underline{a}$. That F_L is continuous from the right at $\underline{a}$ follows the same argument with $\underline{a}$ replaced by a belief $\underline{a} + \varepsilon$ for ε arbitrarily small.

By definition of $\underline{a}$, $F_L(p) > V_L$ for almost all $p \in (\underline{a}, b)$, for some $b > \underline{a}$. The only equilibrium behavior consistent with F_L continuous and increasing at $\underline{a}$ and Lemma 3.2 is a no-trade region between $(\underline{a}, b)$.

Step 4. *At $p = b$, the high type is indifferent between accepting and rejecting $\Psi(b)$ given that (1) $dP = d\hat{P}$ conditional on rejection, and (2) $F_H(p) = \Psi(p)$ for all $p > b$.*

Notice first that we are *not* claiming here that (1) and/or (2) are true in equilibrium. The statement merely pertains to a property that helps us identify b.

We have established that there exists a no-trade interval (a, b), such that $F_L(a) = V_L$. Clearly, $b \neq 1$, otherwise the high type would never trade. By Lemma 3.4, $F_L(b) = \Psi(b)$.

We proceed by contradiction. Suppose that the high type strictly prefers to reject at b given (1) and (2). Belief monotonicity implies that $dP \geq d\widetilde{P} \geq d\hat{P}$. It also implies that the off-path value function F_H^o endowed by this deviation satisfies $F_H^o \geq F_H \geq \Psi$ (Lemma 3.2). Hence, rejection is at least as good in any candidate equilibrium as it is under (1) and (2). The high type will therefore strictly prefer to reject at b in any equilibrium, contradicting its definition as the endpoint of the no-trade region.

Suppose instead that the high type strictly prefers to accept at b given (1) and (2). If $d\hat{P} > 0$ her payoff is $\Psi(b + d\hat{P})$ and if $d\hat{P} < 0$ her payoff is $F_H(b + d\hat{P})$. If the high type strictly prefers to accept at b, then it must be that there exists $\varepsilon > 0$ such that $F_H(b - \varepsilon) < \Psi(b - \varepsilon)$. This violates Lemma 3.2. Again, because (1) and (2) are lower bounds on equilibrium properties, *conditional on* $d\hat{P} > 0$, rejecting at b is at least as profitable in any candidate equilibrium. This only strengthens the argument.

Step 5. *For almost all $p > b$, $w(p) = \Psi(p)$ and both types of sellers accept with probability one.*

Proposition 3.2 establishes that there is a unique solution satisfying the necessary conditions we have established for a no-trade region of this form. We are denoting that solution (a, b).

Recall that in the Ψ-game the high type strictly prefers to accept for all $p > p^*_H$. As shown in Lemma 2.1, $b > p^*_H$ because the reflecting lower boundary makes the high type more willing to wait.

Now, suppose there exists a $p > b$ such that the high type does not strictly prefer to accept $\Psi(p)$. $F_L(b) = \Psi(b)$ and $p > p^*_H$ then imply that there must exist a $p' > b$ such that conditional on rejection at p', $dP > d\hat{P}$. For such a belief to be consistent with Bayes' rule and the proposed equilibrium behavior, it must be that the there is positive probability that the low type trades at a price of V_L when beliefs are greater than b.

It is therefore sufficient to show that the low type does not trade at a price of V_L for any $p > b$. For the purpose of contradiction, suppose there exits a $\widetilde{p} > b$ such that $F_L(\widetilde{p}) = V_L$. The same argument given in Step 3 establishes that F_L is continuous at $\widetilde{p}$. It is clear that $F_L(p) = V_L$ for all $p > \widetilde{p}$ cannot hold in equilibrium, because then the high type never trades for any $p > \widetilde{p}$ forcing $F_H(p) < \Psi(\widetilde{p}) < \Psi(p)$. Hence, F_L must increase continuously from V_L to $\Psi(b')$ for some $b' > \widetilde{p}$. The only equilibrium behavior that can support this property of F_L is a second no-trade region where $F_L(b') = \Psi(b')$. However, this violates Proposition 3.2.

Hence, for all $p > b$, the high type strictly prefers to accept Ψ; therefore it must be offered, which completes the argument.

Step 6. $\rho_L(p, V_L) = \frac{a-p}{a(1-p)}$ *for all $p < a$.*

The final detail to collect is the exact probability that the low type accepts when beliefs are $p < a$. There is a unique no-trade region and a unique belief a such that (1) $F_L(a) = V_L$, and (2) a is within the closure of the no-trade region. Therefore, P_t must jump from any $p < a$ to a following a rejection at p. The formula for $\rho_L(p, V_L)$ given in the theorem is necessary and sufficient for this result. ■

Having pinned down the unique form of the equilibrium, we now identify the unique candidate (a, b) and verify that it constitutes an equilibrium. Appendix A provides a more rigorous analysis including the proofs of Propositions 3.2 and 3.3.

In the no-trade region the seller rejects w and takes her continuation payoff. Applying Ito's formula to the right hand side of (SP) gives

$$F_\theta(P) = k_\theta dt + e^{-rdt} E^\theta \left[F_\theta(P) + F'_\theta(P) dP + \frac{1}{2} F''_\theta(P)(dP)^2 \right]$$

Substituting the law of motion (dP_t) yields a second order differential equation for each type seller's value function. There are six necessary boundary conditions that must be satisfied in any equilibrium.[26] The low type's value at a is V_L

$$F_L(a) = V_L \tag{4}$$

and she is indifferent between accepting or taking her continuation payoff at that point.

$$F_L'(a) = 0 \tag{5}$$

Both types accept expected market value at $p = b$

$$F_L(b) = \Psi(b) \tag{6}$$

$$F_H(b) = \Psi(b) \tag{7}$$

The high type is indifferent at this point

$$F_H'(b) = \Psi'(b) \tag{8}$$

and the belief process is reflecting at $p = a$ for the high type

$$F_H'(a) = 0 \tag{9}$$

The slope conditions on F_H at b and F_L at a are also known as *smooth pasting* conditions and are required for seller indifference in a continuous-time setting (Dixit, 1993). The high type must be indifferent between accepting at b and waiting an instant of time. Heuristically, if good news comes in the next instant then the high type will accept Ψ, and if bad news arrives then she continues to wait. To see that (8) is necessary, suppose that $F_H'(b) > \Psi'(b)$. A convex combination of $\Psi(b+\epsilon)$ and $F_H(b-\epsilon)$ lies strictly below $F_H(b)$ implying the high type strictly prefers to accept. On the other hand, if $F_H'(b) < \Psi'(b)$ than she will strictly prefer to wait.[27]

Proposition 3.2. *There exists a unique pair (a^*, b^*) for which the system of differential equations and boundary conditions is satisfied.*

This pair, along with the strategies described in Theorem 3.1 and the equilibrium belief process, fully characterize the candidate equilibrium. The last step is to verify that a profitable deviation does not exist. In the no-trade region, the low type's value is strictly greater than V_L; any offer attracting only the low type in this region loses money. The high type's value function lies strictly above $E[V_\theta|p]$ for all $p < b$; any offer attracting the high type (also attracts the low type since $F_L(p) < F_H(p)$ for all $p < b$) will lose money on average. Therefore no buyer can deviate profitably.

To see that the same is true for the seller, consider a third-party intermediary who offers to "buy" the seller's problem for $F_\theta(P_t)$ at any time t. This transforms (SP) into an optimal stopping

[26]See Harrison (1985) for a discussion of necessary boundary conditions for a function of a reflected process.

[27]Although (5) and (9) imply the same condition on F_L and F_H, the former is a smooth pasting condition, while the latter follows immediately from the reflective behavior of P.

problem, similar to that in § 2.4, except F_θ is payoff upon stopping rather than Ψ. Clearly the intermediary cannot make the seller any worse off. The next result says that the reverse is also true.

Proposition 3.3. *In the optimal stopping problem, a seller of type θ can do no better than to accept F_θ immediately. Furthermore, any strategy that accepts the intermediary's offer after $T(b) = \inf\{t : P_t \geq b\}$ with positive probability does strictly worse.*

This completes the proof of Theorem 3.1.

3.3 Equilibrium when Information Quality is Extreme

When the news is completely uninformative ($s = 0$), the equilibrium is only slightly different from the description in the theorem. Recall that $\underline{p}$ is the belief that satisfies $E[V_\theta|\underline{p}] = K_H$. At $p = \underline{p}$ the high type is indifferent between trading at expected market value and holding the asset forever. If $s = 0$ then $a = b = \underline{p}$. At $p = \underline{p}$ the buyers mix such that $w = E[V_\theta|\underline{p}]$ is offered after some random waiting time. Because the news is uninformative, beliefs do not change once they jump to $\underline{p}$. The mixing is necessary to keep the low type indifferent between accepting or rejecting V_L for beliefs below $\underline{p}$.

As s increases, a and b increase, as does $b-a$ initially. It can be shown that as $s \to \infty : a, b \to 1$, the expected time to trade goes to zero, $F_H \to V_H$ for all $p \in (0, 1]$, $F_L \to V_L$ for all $p < 1$. The value functions for $s = 0$ and as $s \to \infty$ are shown in Figure 9.

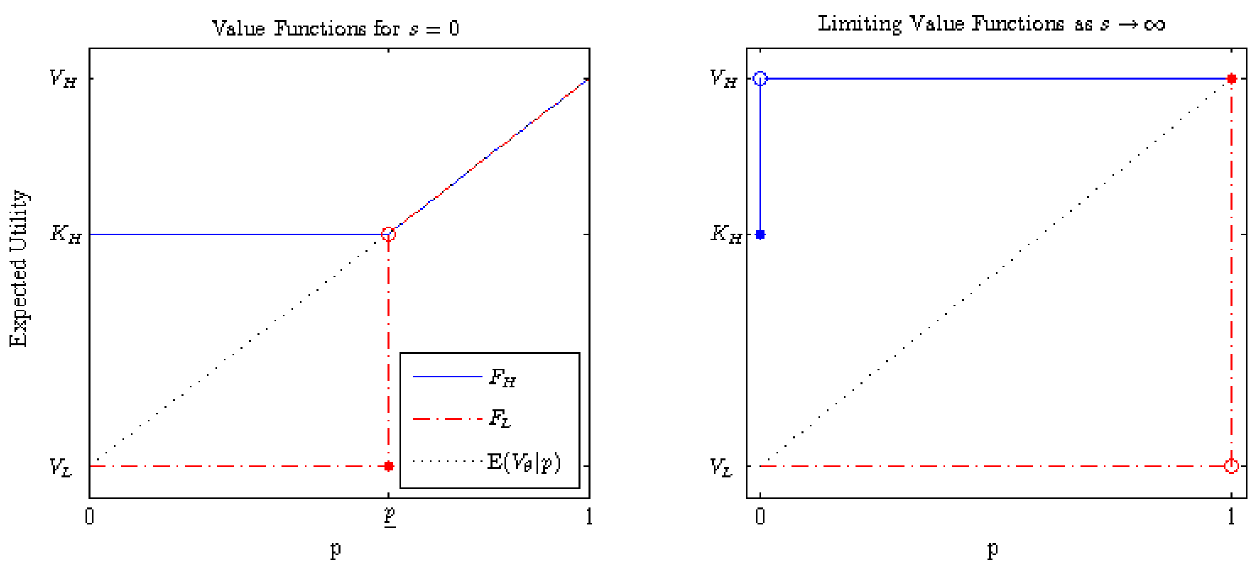

Figure 9: Value functions for extreme s

4 Welfare and Market Efficiency

This section discusses how news affects market efficiency and welfare. Efficiency is measured by total expected surplus to all players. Because buyers make zero expected profit, efficiency depends only on the seller's expected payoff. If $p \geq b$, both types trade immediately, and the fully efficient outcome is obtained. For all $p < b$, there is a strictly positive probability of delay before trade and a loss of potential surplus. Intuition suggests that a market with better information should be more efficient. After all, asymmetry of information between buyers and the seller is the sole cause of delay. This intuition proves correct only for s large enough. Starting from $s = 0$, the market initially becomes less efficient as the quality of news increases. As $s \to \infty$, the expected time to trade goes to zero, and the inefficiency disappears.

The comparative statics for the discount rate are exactly the opposite. In fact, all that matters is the ratio $\frac{r}{s^2}$. As r increases, so too does the expected cost of delaying trade over any increment of time. Increasing s speeds things up: reducing the expected amount of time it takes to reach a certain belief level. Payoffs are discounted for a shorter period of time, thereby decreasing the expected cost of delay.

The total surplus of the game given the state p is the appropriately weighted average of seller value functions

$$TS(p) = pF_H(p) + (1-p)F_L(p)$$

The efficient outcome is to trade immediately, resulting in total potential surplus equal to the expected value of the asset. The percentage loss in efficiency as a function of p is:

$$L(p) = \frac{E[V_\theta|p] - TS(p)}{E[V_\theta|p]}$$

Because trade happens eventually, L measures the expected percentage of total surplus lost due to delay. The left panel of Figure 10 shows a contour of the inefficiency as it depends on p and s. Notice that the inefficiency is most severe when $(p, s) = (\underline{p}, 0)$. As s increases the inefficiency becomes more diffuse and less severe.

Integrating L over all p yields a measure of *expected* loss.[28] Repeating this exercise over values of s allows us to numerically compare how the inefficiency varies with the quality of the news. The right panel in Figure 10 shows the expected inefficiency after integrating over all p. When $s = 0$, the seller's welfare is exactly the same as in Akerlof's static model; below $\underline{p}$ the high type's payoff is K_H and the low type gets V_L, while above $\underline{p}$ the fully efficient outcome is attained with both types trading immediately. The next proposition summarizes how welfare and efficiency are affected as the news becomes more informative.

Proposition 4.1. *Let $(\tilde{a}, \tilde{b}) > (a, b)$ be the boundaries of the no-trade regions for the signal-to-noise ratios $\tilde{s} > s$. Let $\tilde{F}_\theta$, F_θ and $\tilde{L}$, L denote the respective value functions and percentage loss in equilibrium.*

[28] This is akin to assuming a uniform distribution over all priors.

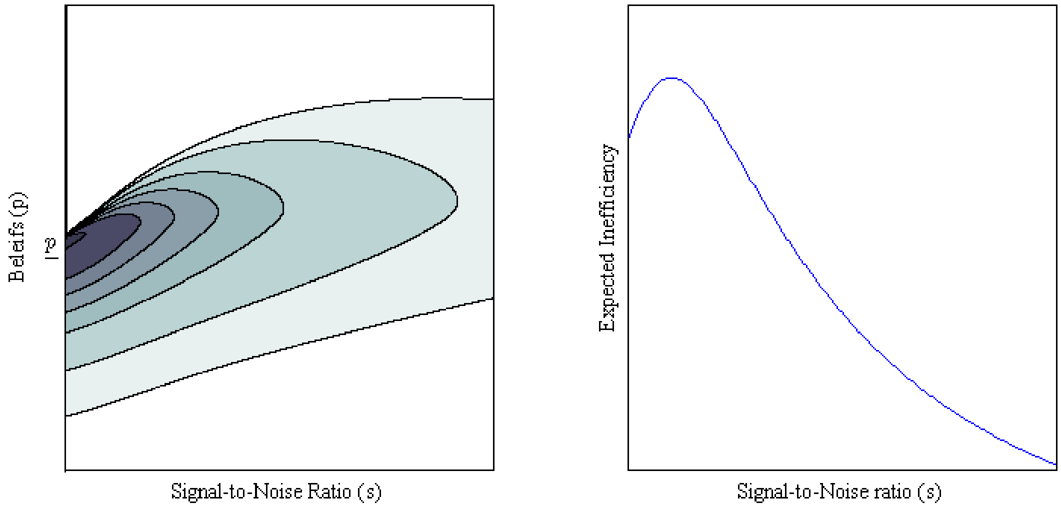

Figure 10: Inefficiency as it depends on p and s, darker indicates more expected loss from delay (left panel). Expected inefficiency as it depends on s (right panel).

1. *$\tilde{F}_H(p) \geq F_H(p)$ where the inequality is strict for $p < \tilde{b}$*

2. *$\tilde{F}_L(p) \leq F_L(p)$ where the inequality is strict for $p \in (a, \tilde{b})$*

Moreover, there exists a $\hat{p} \in (a, b)$ such that

3. *$\tilde{L}(p) < L(p)$ for $p < \hat{p}$*

4. *$\tilde{L}(p) > L(p)$ for $p \in (\hat{p}, \tilde{b})$*

The high type benefits from informative news, while the low type loses. More informative news mitigates inefficiency in a pessimistic market, but *increases* inefficiency for intermediate beliefs. Why? As the information arrives faster the high types' expected benefit from waiting shifts up. She becomes willing to wait at a larger set of beliefs and therefore b increases. a increases to keep the low type indifferent. The region of beliefs for which the inefficiency is largest shifts upward as illustrated in Figure 11.

5 Dynamic Signaling

Under the signaling condition, the high type would prefer to sell at V_L rather than never trade. As a result, she has less incentive to delay trade, especially for p close to zero or one. When $s = 0$, our model is the continuous-time analog of Swinkels (1999). Correspondingly, the unique equilibrium involves no delay: buyers offer $E[V_\theta|p_0]$ at time zero, which is accepted by both types. This result holds even for positive s below some threshold. For s large enough, this equilibrium cannot be sustained. Why? Suppose that Ψ is always offered, and belief monotonicity is satisfied. The equilibrium calls for the high type to accept everywhere, but this strategy is not optimal (recall

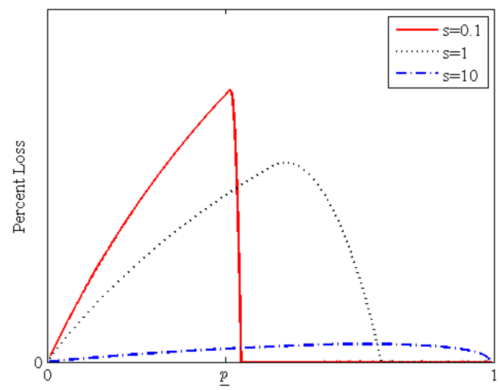

Figure 11: Percentage loss in surplus due to delay for three different levels of news quality.

the optimal policy from § 2.4). She can profitably deviate by rejecting for intermediate beliefs, and so the equilibrium fails.

The equilibrium in Theorem 3.1 (hereafter Ξ^*) exists under the signaling condition provided that news is sufficiently informative and/or players are sufficiently patient. While other forms of equilibria can exist, uniqueness is restored if we restrict attention to equilibria in which the seller's value function is non-decreasing (NDVF hereafter).

Theorem 5.1. *There exists an $\underline{\eta} > 0$ such that for all $r/s^2 < \underline{\eta}$, Ξ^* is an equilibrium under the* ***signaling condition****. Moreover, it is the unique equilibrium satisfying NDVF.*

The construction follows the same arguments as given in § 3.2 making use of two restrictions that were not needed under the lemons condition. The first restriction, $r/s^2 < \underline{\eta}$, is needed for three reasons. As mentioned above, it implies that immediate trade is not an equilibrium. It also ensures that a pair (a, b) exists that solves the six boundary conditions.[29] Finally, we use it to prove that any other form of equilibrium will violate NDVF (Proposition 5.2).

NDVF has intuitive appeal (*should more favorable beliefs ever hurt a seller?*). It can be motivated further by considering the following modification. Suppose that μ_θ is only an upper-bound on the drift of the type θ seller's news process, and the seller can choose any drift below μ_θ at no cost. Any equilibrium of this game *must* satisfy NDVF. If F_θ was decreasing over a no-trade interval, then the seller would prefer to "sabotage" herself by choosing an arbitrarily large negative drift.[30]

However, we cannot immediately rule out that the high type accepts Ψ at the bottom of the no-trade region. In fact, such a region must be an element of any other equilibrium.

[29] It can be shown that the former implies the latter.

[30] It can be shown that a no-trade region must be a component of any equilibrium for $r/s^2 < \underline{\eta}$. Even without NDVF, F_θ must be non-decreasing on any interval where trade occurs with positive probability almost everywhere.

Proposition 5.1. *Under NDVF and for r/s^2 small enough, the only possible equilibria besides Ξ^* must involve a no-trade region where both type sellers accept Ψ at both the upper and lower boundary.*

The final step of the uniqueness argument for Theorem 5.1 is to show that any equilibrium involving the second type of no-trade region will fail NDVF. The reason is that the boundaries of the no-trade region are pinned down by solving for the high type's optimal policy in the Ψ-game. There is no regard for the low type in the construction. As r/s^2 decreases, the no-trade region expands. The low type is worse off by being at $p = \underline{p}_H + \epsilon$ than she is at $p = \underline{p}_H$ because she is only waiting for the inevitable.

Proposition 5.2. *In any equilibrium with a no-trade region in which the high type accepts $\Psi(p)$ on either side, the low types' value function fails NDVF for small enough r/s^2.*

6 Robustness

Our modeling assumptions enable us to express results in a clean and simple manner. Nevertheless, it is important to investigate alternative specifications that may provide new insights and, at the very least, serve as a robustness check. In this section we argue that Ξ^* is robust to the following alternatives: public offers, divisible assets, single-buyer arrivals and a market of buyers arriving randomly. For brevity, we will discuss the results only for the lemons condition.[31]

If offers are private, inference is based only on the fact that the offer was rejected. When offers are public, buyers can also make inference from the *level* of the rejected offer. In addition, a buyer's information set includes the entire history of offers, making the argument for restricting attention to strategies that are stationary in beliefs less compelling. Nevertheless, Ξ^* is an equilibrium.[32] Because buyers do not condition their strategies on past offers in Ξ^*, the seller cannot improve her payoffs by conditioning her strategy on past offers. The only potential for a profitable deviation is if beliefs jump upward following an unexpected rejection. They do not in Ξ^*. Next, suppose all other players are playing the strategies in Ξ^*, and consider a deviation for a buyer. The same logic used in the verification step of § 3.2 applies; any offer that will be accepted with positive probability earns negative expected profit.

If the seller's asset is divisible, then Ξ^* is an equilibrium of the game. If quantities traded are unobservable, then it is the unique equilibrium. However, if prior trades are publicly observable, additional complications arise. Much like the case of public offers, buyers' information set is enriched to include the entire history of past quantities. In addition, suppose that the equilibrium called for 50% of the asset to be sold at time zero, but only 25% was sold. What should the market then think at $t > 0$? It seems natural to extend Belief Monotonicity in the following way: beliefs should not decrease (increase) if the quantity sold at date t was less (more) than expected. Ξ^* is robust to this refinement.

[31]Most of the results hold under the signaling condition for s large.

[32]If offers are public then the mixing below a *must* come only from the low-type seller.

Intra-temporal competition drives profits to zero when there are multiple buyers at each point in time. This immediately pins down the price given a p and the seller types trading. If a single buyer is present at τ, he is only competing against buyers who arrive at $t > \tau$. However, in a continuous-time setting, *inter*-temporal competition is capable of driving buyers' expected profits to zero. Ξ^* is still an equilibrium even if only a single buyer arrives at each instant. The proof is straightforward; if each buyer plays according to w then clearly the seller's strategies are still optimal, the value functions are as in Figure 6 and no buyer can profitably deviate.

In thin markets, it is unlikely that a market of buyers will be present at each point in time. Suppose instead that there is a market for the asset that arrives according to a Poisson process with arrival rate γ. Upon arrival, multiple buyers simultaneously make offers. The seller receives no offers between arrivals. Clearly the seller's payoff is adversely affected, but the equilibrium of the game involves three analogous regions, provided the market is not too thin (γ is sufficiently large). For beliefs below some lower threshold, V_L is offered upon arrival. The high type rejects. The low type mixes in such a way that beliefs jump to the lower threshold. Above some upper threshold, Ψ is offered upon arrival, which both types accept. No trade occurs for beliefs in between the two thresholds. In extremely thin markets (for small γ), there are only two regions. The low type accepts V_L with probability one in the lower region. In the upper region Ψ is always offered, and both types accept. This occurs when V_L is greater than the low types expected discounted value of getting V_H upon the next arrival.

Part II

We now extend the model from Part I by incorporating *liquidity shocks*, which make an asset holder more eager to sell, in order to capture the following observations.

- Assets in most liquid financial markets are traded repeatedly.
- The return on many financial investments comes from selling in the future at a higher price rather than holding to maturity.
- An investor's time value for money may change due to unexpected circumstances.

The model helps us to understand and explain the following insights.

1. A small amount news can have a drastic effect trading behavior.
2. A sell-off can help stabilize the market.
3. A more unstable environment (more frequent shocks) decreases asset prices.
4. Low-value assets will be *eventually* traded at their *fundamental* value.
5. A seller cannot attract a trading partner without a credible reason for wanting to sell.

In the next section, we present the model with liquidity shocks and describe equilibrium behavior. We then characterize equilibrium asset values. An algorithm to solve for the equilibrium is presented and used to solve a numerical example. The economic insights mentioned above are then discussed in more detail.

7 The Model with Liquidity Shocks

The structure of the model is similar to that in Part I. The primary difference is that an asset holder faces the risk of being hit by a liquidity shock. Upon arrival, the shock increases the rate at which the agent discounts future payoffs. An owner of the asset is *not* forced to sell upon arrival, but she is more eager to do so.

The game begins at $t = 0$ with an indivisible asset owned by an agent denoted by A_0. The asset may be one of two types. A_0 knows the type, potential buyers do not. An asset of type θ generates a cash flow v_θ ($v_H > v_L$), which accrues to the current owner as a flow payoff. At every t, two or more buyers arrive and make private offers to the current owner. If a buyer's offer is accepted, he becomes the new owner and immediately learns the asset's type. All rights to future cash flows are transferred to him. A buyer whose offer is rejected exits the game permanently.

Initially, all agents discount future payoffs at a rate of r. Let $V_\theta = v_\theta / r$ denote the *fundamental value* of a type θ asset, and denote the current owner of the asset by A_t. Liquidity shocks arrives according to a Poisson process with arrival rate λ. Upon arrival, the rate at which A_t discounts future payoffs increases to $\bar{r} > r$, where it remains until A_t exits the game.[33] We will refer to A_t

[33] For simplicity, only the current owner of the asset is affected.

as a *seller* or a liquidity-constrained owner if she has been hit by a shock and as a *holder* or an unconstrained owner if she has not yet been hit by a shock.

A shock that arrives at time τ is observable to all buyers arriving at times $t \geq \tau$. In addition, news about the asset is revealed via the same diffusion process as in § 2.1. The information set of an agent arriving at time τ contains:

- The entire history of news, $\{X_t : 0 \leq t \leq \tau\}$
- The arrival times of every liquidity shock before time τ (if any)
- Every time at which the asset has been traded before time τ (if any)

The beliefs about the type of asset are conditioned on all of the above. When $\lambda = 0$, the model presented in this section is equivalent to the model in Part I. As in Part I, we restrict attention to equilibria which are stationary in the market beliefs.[34] Strategies are also contingent on whether the current owner has been hit by a shock. The results are robust to the alternative specifications discussed in Remark 2.1.

At time $t = 0$, the market begins with a common prior $P_0 = \Pr_0(\theta = H)$. As in § 2.1, let $\hat{P} = \{\hat{P}_t, 0 \leq t \leq \infty\}$ denote the belief process for a Bayesian who updates only based on news, and let $P = \{P_t, 0 \leq t \leq \infty\}$ denote the stochastic process representing the equilibrium market beliefs. For the analysis in this section, it will be convenient to work with a state variable that represents beliefs in log-likelihood space rather than probability space. Define $\hat{Z} \equiv \ln \frac{\hat{P}}{1-\hat{P}}$, $Z \equiv \ln \frac{P}{1-P}$, and let $z \equiv \ln \frac{p}{1-p}$. Applying Ito's lemma to $\hat{Z}$ gives

$$d\hat{Z}_t = \frac{s}{\sigma}\left(dX_t - \frac{1}{2}(\mu_H + \mu_L)\,dt\right) \tag{10}$$

where $s \equiv \frac{\mu_H - \mu_L}{\sigma}$.[35] The mapping from p to z is strictly monotone (and therefore also one-to-one). Henceforth, we will work primarily in *z-space*.[36]

We characterize an equilibrium with a similar structure to Ξ^*. In equilibrium, a holder never trades. Beliefs evolve strictly according to news over any interval of time in which a holder is not hit by a shock. When A_t is a seller, the equilibrium is characterized by a pair $(\alpha^*_\lambda, \beta^*_\lambda)$ such that when buyers are *pessimistic* ($z < \alpha^*_\lambda$), the low-type sell trades with positive probability causing the equilibrium beliefs to jump immediately to α^*_λ. When buyers are *optimistic* ($z > \beta^*_\lambda$), the market is efficient. An owner with a credible reason to liquidate does so immediately. The asset is never traded for $z \in (\alpha^*_\lambda, \beta^*_\lambda)$, both sides of the market wait for more information to be revealed.

[34]See § 2.2.

[35]See Appendix A.1.1 for more details.

[36]We refer to z as the "state" and p as the "market belief". The statement "the state is z" is equivalent to "the market belief is p" where z is defined as above.

8 Asset Values in Equilibrium

Here we delineate a set of necessary conditions for the equilibrium we are interested in. Let $F_\theta(z)$ denote the value of a type θ asset to a seller given the state z. Similarly, $G_\theta(z)$ denotes the value to a holder. $B(z)$ denotes the *expected* value of the asset to an arriving buyer were he to purchase it. The value of the asset to both a seller and a holder depends on θ (because both know the asset type), while the buyer's value does not. In Part I (or if $\lambda = 0$), a buyer's expected value for the asset is simply the expected fundamental value.

$$B(z) = \Psi(z) \equiv E\left[v_\theta \int_0^\infty e^{-rs} ds \big| z\right]$$

When $\lambda > 0$, a buyer's value depends not only on his current beliefs but also on his ability to sell the asset in the future when he is hit by a shock.

$$B(z) = E[G_\theta(z)|z] \tag{11}$$

This makes solving the equilibrium a recursive problem. A buyer's value depends on an owner's value, which depends on a seller's value, which in turn depends on a buyer's value. To characterize the equilibrium, we derive a system of interdependent differential equations and specify the necessary boundary conditions using equilibrium arguments.

The Bellman equation for a liquidity-constrained seller with an asset of type θ is

$$F_\theta(z) = \max\left\{w(z), v_\theta dt + e^{-\bar{r}dt} E^\theta\left[F_\theta(z + dZ_t)\right]\right\} \tag{SPλ}$$

where $w(z)$ denotes the maximal offer made by buyers in the state z. Applying Ito's lemma to F_θ, using the law of motion of Z and taking the expectation, (SPλ) implies a differential equation that F_θ must satisfy for all $z \in (\alpha^*_\lambda, \beta^*_\lambda)$. Namely, for a high-type seller

$$\frac{s^2}{2}(F''_H + F'_H) - \bar{r}F_H + v_H = 0 \tag{12}$$

and for a low-type seller

$$\frac{s^2}{2}(F''_L - F'_L) - \bar{r}F_H + v_L = 0 \tag{13}$$

There are six necessary boundary conditions (analogous to § 3.2). A low-type's value at α^*_λ corresponds with a buyer's value, and she is indifferent between accepting or taking her continuation payoff at that point.

$$F_L(\alpha^*_\lambda) = B(\alpha^*_\lambda) \tag{14}$$

$$F'_L(\alpha^*_\lambda) = 0 \tag{15}$$

Both types accept an offer of $w = B(\beta_\lambda^*)$ at $z = \beta_\lambda^*$

$$F_\theta(\beta_\lambda^*) = B(\beta_\lambda^*) \qquad \theta \in \{L, H\} \tag{16}$$

The high type is indifferent at this point, and the belief process is purely reflecting at $z = \alpha_\lambda^*$

$$\begin{aligned} F_H'(\beta_\lambda^*) &= B'(\beta_\lambda^*) && (19)\\ F_H'(\alpha_\lambda^*) &= 0 && (20) \end{aligned}$$

For all $z > \beta_\lambda^*$, both type sellers trade immediately at $w = B(z)$, therefore $F_H(z) = F_L(z) = B(z)$. For all $z < \alpha_\lambda^*$, the equilibrium beliefs jump instantaneously to α_λ^* (as long as the current owner is liquidity constrained), therefore $F_H(z) = F_H(\alpha_\lambda^*)$, $F_L(z) = F_L(\alpha_\lambda^*)$ for all $z < \alpha_\lambda^*$.

Given B, the differential equations (12)-(13) along with the boundary conditions (14)-(20) completely determine $(\alpha_\lambda^*, \beta_\lambda^*)$ and a seller's value functions for all z. However, as indicated by equation (11), B is determined from the asset value to a holder.

In equilibrium, a holder does not trade prior to being hit by a shock. This is because she is privately informed about the asset, and (since the shocks are observable) it is common knowledge that there are no gains from trade (Milgrom and Stokey, 1982). A holder simply consumes her flow payoff until a shock transforms her into a seller. The value of the asset comes from the instantaneous flow payoff and the discounted expected value of the asset an instant later. With probability $1 - \lambda dt$, A_t is not hit by a shock and has a value of $G_\theta(z + dZ_t)$. With probability λdt, the shock comes and A_t becomes a liquidity-constrained seller with value $F_\theta(z + dZ_t)$. Thus, the holder's value for the asset must satisfy the following recursive equation for all z

$$G_\theta(z) = v_\theta dt + e^{-rdt} E^\theta \big[(1 - \lambda dt) G_\theta(z + dZ_t) + \lambda dt F_\theta(z + dZ_t)\big] \tag{21}$$

Using similar methods as above, (21) implies the following differential equation for the value of each type holder

$$\begin{aligned} \frac{s^2}{2}(G_H'' + G_H') - (r + \lambda) G_H + \lambda F_H + v_H &= 0 && (22)\\ \frac{s^2}{2}(G_L'' - G_L') - (r + \lambda) G_L + \lambda F_L + v_L &= 0 && (23) \end{aligned}$$

The final step in pinning down equilibrium asset values is to determine the boundary conditions for G_L and G_H. To do so, we make use of the fact that as $p \to 1$ ($z \to \infty$), the effect of news on equilibrium beliefs goes to zero. A holder is simply waiting for the shock to come, at which point she has a seller's value for the asset. The same is true as $p \to 0$. In the limit, a holder's value for the asset is a weighted average of the fundamental value and a seller's value. The following boundary conditions complete the characterization of asset values

$$\lim_{z\to\infty} G_\theta(z) = \frac{rV_\theta + \lambda \lim_{z\to\infty} F_\theta(z)}{r+\lambda} \qquad \theta \in \{L,H\} \tag{24}$$

$$\lim_{z\to-\infty} G_\theta(z) = \frac{rV_H + \lambda \lim_{z\to-\infty} F_\theta(z)}{r+\lambda} \qquad \theta \in \{L,H\} \tag{25}$$

Only the holder is directly affected by the arrival rate of the shocks. Buyer and seller values are affected indirectly through G_θ.[37]

9 Computing the Equilibrium: An Algorithmic Approach

In the previous section, we characterized equilibrium asset values through a system of differential equations and boundary conditions. In this section, we focus on solving that system. In Part I, we were able to derive closed-form solutions for the sellers' value functions and use them to verify the existence of an equilibrium candidate and solve for it. This approach is no longer tractable for two reasons. First, a buyer's value for the asset is endogenous to the system, whereas in Part I it was known to be the expected fundamental value. Second, the system of differential equations is interdependent, whereas in Part I they were independent. Theoretically it is still possible to derive closed form solutions for the system. However, an algorithmic approach is more practical. We present an algorithm for solving the equilibrium that allows us to numerically compute the equilibrium asset values and boundaries.

The algorithm follows an iterative process.[38] In each iteration, the asset values and boundaries are computed. Let B^k, F^k_θ, G^k_θ denote these values for iteration k and let $(\alpha^k_\lambda, \beta^k_\lambda)$ denote the no-trade boundaries.

Step 0: Initialize $B^0 = \Psi$ and let $k = 0$.

Step 1: Using B^k, solve for $\{F^k_L, F^k_H\, \alpha^k_\lambda, \beta^k_\lambda\}$ using the differential equations (12) and (13), along with the boundary conditions (14)-(20).

Step 2: Using $\{F^k_L, F^k_H\}$, solve for $\{G^k_L, G^k_H\}$ using the differential equations (22) and (23), along with the boundary conditions (24)-(25).

Step 3: Define $B^{k+1}(z) = E[G^k_\theta|z]$ for all z. Increment $k = k+1$.

Step 4: Repeat Steps 1-3 until convergence.

The algorithm has an intuitive appeal in the following sense. Iteration k of the algorithm solves for the equilibrium asset values in a world where there can be at most k shocks. This implies that the asset values in the first iteration correspond with those in Part I. We conjecture (and numerical results seem to confirm) that the following are true:

[37] Setting $\lambda = 0$ implies that $G_H = V_H$ and $G_L = V_L$.

[38] We thank Yuliy Sannikov for suggesting this approach.

- The algorithm converges to a [unique] fixed point $\{\alpha^*_\lambda, \beta^*_\lambda, B^*, F^*_L, F^*_H, G^*_L, G^*_H\}$
- The fixed point constitutes an equilibrium of the game where the strategy of buyers and sellers is identical to Ξ^* after replacing $\{\Psi, \alpha^*, \beta^*\}$ with $\{B^*, \alpha^*_\lambda, \beta^*_\lambda\}$

9.1 Numerical Example

Using the algorithm described above, we compute equilibrium asset values and boundaries for a particular set of parameters. We vary λ to illustrate how equilibrium asset values depend on the arrival rate of shocks. The following parameters remain fixed: $v_L = 1$, $v_H = 2$, $r = 10\%$, $\bar{r} = 25\%$ and $s = 2$. Figure 12 shows the asset values for each type of seller and holder as well as a buyer's value. First, notice that the structure $\{B^*, \alpha^*_\lambda, \beta^*_\lambda\}$ is just as in Part I, with the exception that a

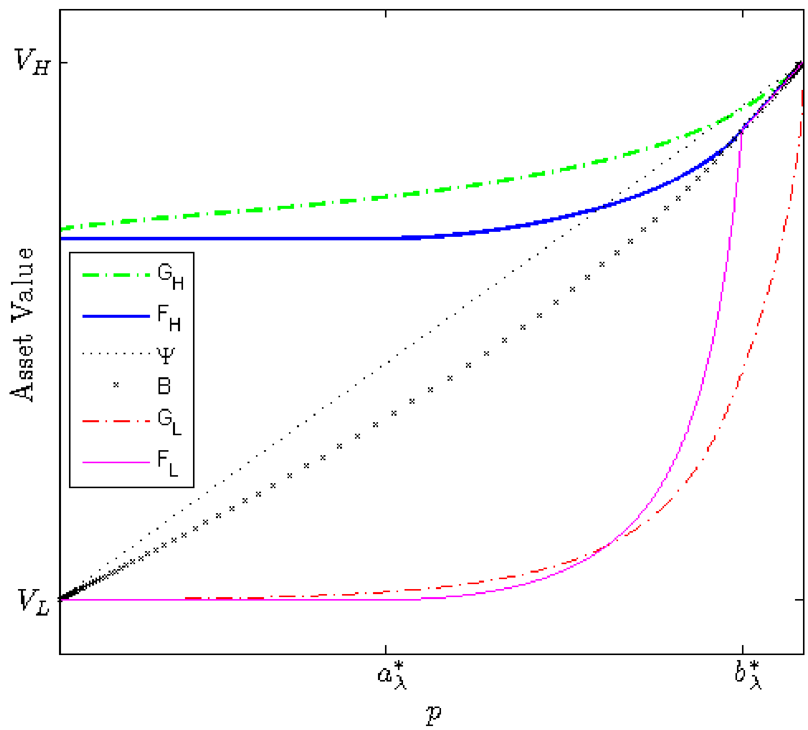

Figure 12: Asset values for $\lambda = 2$. $(a^*_\lambda, b^*_\lambda) = (0.44, 0.92)$ are the corresponding beliefs for $(\alpha^*_\lambda, \beta^*_\lambda)$

buyer's value now lies everywhere below Ψ. Not surprisingly, the value to the owner of a high-value asset is strictly higher before she is hit buy a liquidity shock. However, the same is not true for the low type asset. When beliefs are favorable, a low-type holder would prefer to become a seller because the low-type seller can trade at a price above V_L. This highlights the importance of the assumption that liquidity shocks are observable.

Figure 13 illustrates how the asset values depend on the arrival rate of the shocks. The right panel illustrates that higher λ correspond to lower asset values. As $\lambda \to 0$, the equilibrium asset values convergence to the values in Part I (left panel).

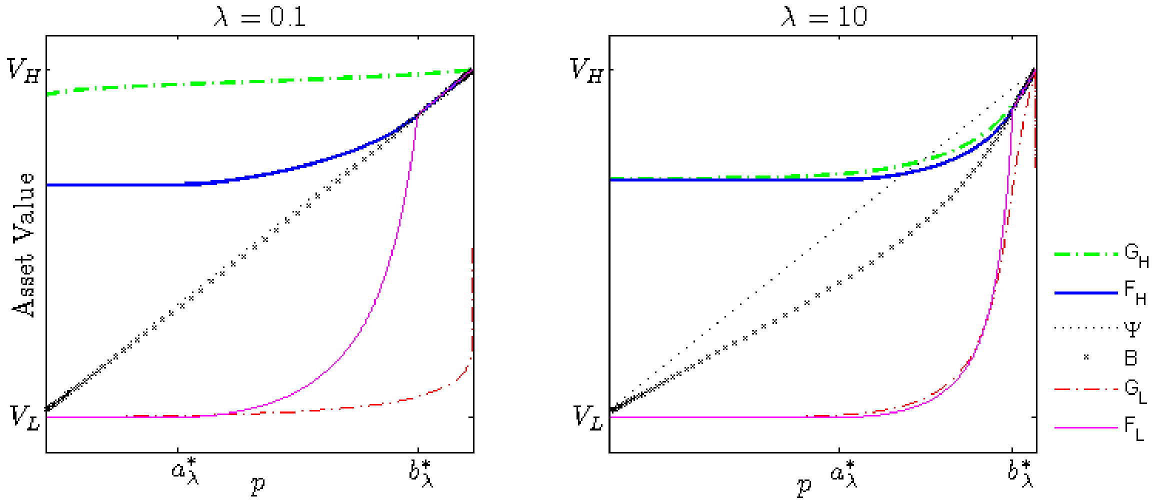

Figure 13: Equilibrium asset prices for $\lambda = 0.1$ (left panel) and $\lambda = 10$ (right panel)

10 Economic Insights

We begin with insight gained from the mere description of the equilibrium. The market value of an asset varies with news and current market conditions even though fundamentals never change. Dramatic price volatility can correspond to equilibrium behavior depending on $(\alpha^*_\lambda, \beta^*_\lambda, s)$. For example, the price path of an asset receiving a steady stream of bad news will drop from above $B(\beta^*_\lambda)$ to V_L with no trades (at intermediate prices) in between.

We now elaborate on some of the results mentioned earlier.

1. In equilibrium a very small amount of bad news can cause a market to move from a fully liquid state to one in which liquidity completely "dries up." As in Part I, when $z \in (\alpha^*_\lambda, \beta^*_\lambda)$ any price that a buyer is willing to offer is unacceptable to the seller.

2. A pessimistic market stabilizes through a partial sell-off at the lowest fundamental value. This creates a lower bound on market beliefs when news is continuous and trading opportunities are frequent.

3. In comparison to values in Part I, liquidity shocks decrease the asset value to both sellers and buyers. This result is fairly intuitive. The buyer's value depends on the price at which he can sell it in the future when he is hit with a shock. If beliefs lie below b^*_λ when this occurs, then he may not trade immediately upon arrival of a shock. This distorts the buyer's value downward (below Ψ) for all beliefs. More frequent shocks decrease the price at which buyers are willing to pay as investors face costly liquidation more frequently—a less stable world is

bad for asset trading. However, the price at which assets are traded never drops below the fundamental value of a low-value asset.

4. With an infinite horizon, a low-value asset will eventually be “found out”—with probability one, it will trade at a price of V_L at $Z_t = \alpha^*_\lambda$ for some finite t. Afterward buyers will only be willing to pay the fundamental value. This result is similar to Bar-Isaac (2003) and Hendel and Lizzeri (1999).

5. The observability of the shocks is an important assumption in the model. When beliefs are favorable, the low-type holder is anxiously waiting for the shock to come allowing her to pool with a high-type seller at a price higher than the fundamental value of the asset. If shocks were unobservable, then a holder of a low-value asset would sometimes prefer to sell before being hit by a shock, breaking the equilibrium. An observable shock provides the owner with a credible reason to liquidate. Without this, buyers face more severe exposure to the lemon’s problem. Formal analysis of a model with unobservable shocks is left for future research.

11 Final Remarks

In this paper, we presented a continuous-time framework to analyze the effect of news in a dynamic market with asymmetric information. The equilibrium consists of three distinct regions: immediate trade, no trade, and partial sell-off. The equilibrium is robust to public offers, divisible assets and thinner markets. The model encompasses both dynamic lemons markets and dynamic signaling markets and established the strong connection between both the strategic settings and their equilibria. We introduced liquidity shocks to the model characterized a similar equilibrium. The predictions explain several phenomenon observed in financial markets.

Intuitively, one might think that news reduces the asymmetry between buyers and sellers and mitigates inefficiencies. We demonstrate that this logic is only partially true. The welfare results have policy implications. For example, would a social planner ever suppress or censor informative news? Are markets more efficient when information is revealed gradually, or all at once? A comparison to Daley and Green (2008) can help answer these questions and is an investigation we plan to conduct in future research.

The observability of liquidity shocks was crucial in providing sellers with a credible reason to sell. If shocks are unobservable, holder's of low-type assets would prefer to sell, decreasing the proportion of high-value assets and breaking the equilibrium. We leave this analysis for future work.

References

Admati, Anat and Michael Perry, "Strategic Delay in Bargaining," *Review of Economic Studies*, 1987, *54*, 345–364.

Akerlof, George, "The Market for Lemons: Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics*, 1970, *84*, 488–500.

Banks, Jeffrey and Joel Sobel, "Equilibrium Selection in Signaling Games," *Econometrica*, 1987, *55*, 647–662.

Bar-Isaac, Heski, "Reputation and Survival: Learning in a Dynamic Signalling Model," *Review of Economic Studies*, 2003, *70*, 231–252.

Cox, Jeff, "Higher Volume Indicates Push Toward Turning Point," October 10 2008.

Daley, Brendan and Brett Green, "Market Signaling with Grades," *Working Paper*, 2008.

Dixit, Avinash, *The Art of Smooth Pasting*, Harwood Academic Publishers, 1993.

Dixit, Avinash K. and Robert S. Pindyck, *Investment under Uncertainty*, Princeton University Press, 1994.

Harrison, Michael J., *Brownian Motion and Stochastic Flow Systems*, Robert E. Kreiger Publishing Company, 1985.

Hendal, Igal and Alessandro Lizzeri, "The Role of Leasing under Adverse Selection," *Journal of Political Economy*, 2002, *110*, 113–143.

Hendel, Igal, Alessandro Lizzeri, and Marciano Siniscalchi, "Efficient Sorting in a Dynamic Adverse-Selection Model," *Review of Economic Studies*, 2005, *72*, 467–497.

_ **and** _ , "Adverse Selection in Durable Goods Markets," *American Economic Review*, 1999, *89*, 1097–115.

Hörner, Johannes and Nicolas Vieille, "Public vs. Private Offers in the Market for Lemons," *Econometrica (forthcoming)*, 2008.

Janssen, Maarten C. W. and Santanu Roy, "Dynamic Trading in a Durable Good Market with Asymmetric Information," *International Economic Review*, 2002, *43:1*, 257–82.

Johnson, Justin and Michael Waldman, "Leasing, Lemons, and Buy-backs," *Rand Journal of Economics*, 2003, *34*, 247–265.

Karatzas, Ioannis and Steven Shreve, *Brownian Motion and Stochastic Calculus*, Springer, 2004.

Kohlberg, Elon and Jean-Francois Mertens, "On the Strategic Stability of Equilibria," *Econometrica*, 1986, *54*, 1003–1039.

Kremer, Ilan and Andrzej Skrzypacz, "Dynamic Signaling and Market Breakdown," *Journal of Economic Theory*, 2007, *133.*

Kreps, David and Robert Wilson, "Sequential Equilibria," *Econometrica*, 1982, *50*, 863–894.

Lee, Jihong and Qingmin Liu, "Reputation in Repeated Settlement Bargaining," 2008.

Milgrom, Paul and Nancy Stokey, "Information Trade and Common Knowledge," *Journal of Economic Theory*, 1982, *26:1*, 17–27.

Nöldeke, Georg and Eric Van Damme, "Signalling in a Dynamic Labour Market," *The Review of Economic Studies*, January 1990, *57* (1), 1–23.

Oksendal, Bernt, *Stochastic Differential Equations*, Springer; 6th Edition, 2007.

Reuters, "Prices Nosedive, Liquidity Dries Up," September 17 2008.

Smith, Hannah, "EM Debt Latest Asset Class in Freefall as Liquidity Dries up," November 3 2008.

Spence, Michael A., "Job Market Signaling," *Quarterly Journal of Economics*, 1973, *90*, 225–243.

Swinkels, Jeroen M., "Education Sginalling with Preemptive Offers," *Review of Economic Studies*, 1999, *66*, 949–970.

Taylor, Curtis, "Time-on-the-Market as a Signal of Quality," *Review of Economic Studies*, 1999, *66*, 555–578.

Weiss, Andrew, "A Sorting Cum-Learning Model of Signaling," *Journal of Political Economy*, 1983, *91*, 420–442.

Zweig, Jason, "Capitulation: When the Market Throws in the Towel — Surprisingly, Bear Markets Don't Always End With a Bang – Sometimes It's Just a Whimper," The Wall Street Journal October 25 2008.

A Appendix

The proof of the results in § 2.4 and § 3.1 are postponed to § A.2 and § A.3.

A.1 Proofs for Results under the Lemons Condition

A.1.1 Existence and Uniqueness of a Candidate

From this point forward it will be convenient to make a change of variables from probability to log-likelihood ratio: define a new state variable $z \equiv \ln(p/(1-p))$ and the process $\hat{Z} \equiv \ln(\hat{P}/(1-\hat{P}))$. From equation (2) we can write $\hat{Z}$ in terms of X:

$$\begin{aligned} \hat{Z}_t &= \ln\left(\frac{p_0}{1-p_0} e^{\frac{s}{\sigma}\left(X_t - \frac{\mu_H+\mu_L}{2}t\right)}\right) \\ &= \hat{Z}_0 + \frac{s}{\sigma}X_t - \frac{1}{2}\frac{s}{\sigma}\left(\mu_H+\mu_L\right)t \end{aligned} \tag{26}$$

where $s \equiv \frac{\mu_H-\mu_L}{\sigma}$. Note that p has a range on the unit interval while the range of z is $\mathbb{R}$. Using Ito's lemma, the SDE for $\hat{Z}$ is

$$d\hat{Z}_t = \frac{s}{\sigma}\left(dX_t - \frac{1}{2}\left(\mu_H+\mu_L\right)dt\right) \tag{27}$$

Inserting dX_t^θ from (1), the process for each type can be written as

$$\begin{aligned} d\hat{Z}_t^H &= \frac{s}{\sigma}\left(\mu_H dt + \sigma dB_t - \frac{1}{2}\left(\mu_H+\mu_L\right)dt\right) \\ &= \frac{s^2}{2}dt + s dB_t \end{aligned} \tag{28}$$

$$\begin{aligned} d\hat{Z}_t^L &= \frac{s}{\sigma}\left(\mu_L dt + \sigma dB_t - \frac{1}{2}\left(\mu_H+\mu_L\right)dt\right) \\ &= -\frac{s^2}{2}dt + s dB_t \end{aligned} \tag{29}$$

Because the mapping from p to z is injective, there is no loss by making this transformation and the analysis is more tractable under the Brownian diffusion $(\hat{Z}_t)$ process rather than the more complicated non-linear $(\hat{P}_t)$ process. Abusing notation, we continue to use $F_\theta(z)$ to denote the seller's value function and Ψ to denote the expected value of V_θ given z: $\Psi(z) \equiv E[V_\theta|z] = V_L + (V_H - V_L)e^z(1+e^z)^{-1}$. We omit function arguments when there can be no confusion by doing so.

In the no-trade region, the seller's value function must satisfy

$$F_\theta(z) = k_\theta dt + e^{-rdt}E\left[F_\theta(z + dZ^\theta)\right] \tag{30}$$

by Ito's lemma

$$dF_\theta = \frac{\partial F_\theta}{\partial z}dZ^\theta + \frac{1}{2}\frac{\partial^2 F_\theta}{\partial z^2}(dZ^\theta)^2$$

and so (30) becomes

$$F_\theta(z) = k_\theta dt + e^{-rdt}E\left[F_\theta(z) + F'_\theta(z)dZ^\theta + \frac{1}{2}F''_\theta(z)\left(dZ^\theta\right)^2\right]$$

Recall that the equilibrium belief process P_t follows $\hat{P}_t$ for all $p \in (a,b)$. Hence the equilibrium log-likelihood process Z follows $\hat{Z}$ for all $z \in (\alpha, \beta)$, where $\alpha \equiv \ln(a/(1+a))$ and $\beta \equiv \ln(b/(1-b))$. Inserting the law of motion and taking the expectation for the low type gives

$$\begin{aligned} F_L &= k_L dt + e^{-rdt}\left(F_L + E\left\{F_L' dZ^L + \frac{1}{2}F_L''\left(dZ^L\right)^2\right\}\right) \\ &= k_L dt + e^{-rdt}\left(F_L + -\frac{s^2}{2}F_L' dt + \frac{1}{2}s^2 F_L'' dt\right) \end{aligned}$$

Subtracting $e^{-rdt}F^L$ from both sides and dividing by dt

$$F_L\frac{(1-e^{-rdt})}{dt} = k_L + e^{-rdt}\frac{s^2}{2}\left(F_L'' - F_L'\right)$$

taking the limit as $dt \to 0$ gives the following differential equation

$$\frac{s^2}{2}F_L'' - \frac{s^2}{2}F_L' - rF_L + k_L = 0 \tag{31}$$

The differential equation for the high type is derived similarly. It satisfies

$$\frac{s^2}{2}F_H'' + \frac{s^2}{2}F_H' - rF_H + k_H = 0 \tag{32}$$

Letting $\eta \equiv \frac{r}{s^2}$, rewrite (31) as

$$F_L'' - F_L' - 2\eta(F_L - K_L) = 0 \tag{33}$$

which has a solution of the form

$$F_L(z) = c_1 e^{u_1 z} + c_2 e^{u_2 z} + K_L \tag{34}$$

where $(u_1, u_2) = \frac{1}{2}\left(1 \pm \sqrt{1+8\eta}\right)$ and c_1, c_2 are coefficients yet to be determined. The boundary conditions (4),(5) and (6) become

$$\begin{aligned} u_1 c_1 e^{u_1\alpha} + u_2 c_2 e^{u_2\alpha} &= 0 && (35) \\ c_1 e^{u_1\beta} + c_2 e^{u_2\beta} + K_L &= \Psi(\beta) && (36) \\ c_1 e^{u_1\alpha} + c_2 e^{u_2\alpha} + K_L &= V_L && (37) \end{aligned}$$

For the high type, the solution to (32) is of the form

$$F_H(z) = d_1 e^{q_1 z} + d_2 e^{q_2 z} + K_H \tag{38}$$

where $(q_1, q_2) = -\frac{1}{2}\left(1 \pm \sqrt{1+8\eta}\right)$. The boundary conditions from (7), (8) and (9) become

$$\begin{aligned} q_1 d_1 e^{q_1\beta} + q_2 d_2 e^{q_2\beta} &= \Psi'(\beta) && (39) \\ d_1 e^{q_1\beta} + d_2 e^{q_2\beta} + K_H &= \Psi(\beta) && (40) \\ q_1 d_1 e^{q_1\alpha} + q_2 d_2 e^{q_2\alpha} &= 0 && (41) \end{aligned}$$

The astute reader has noticed that at this point we have six equations and six unknowns. Proposition 3.2 states that there exists a unique a solution to this system of equations. To prove this result, it will be useful to show a series of lemmas. First we will show that there exists two increasing, C^1 functions: B_L, which maps a lower boundary into an upper boundary that satisfies the low-type boundary conditions (35)-(37) and B_H, which maps a lower boundary into an upper boundary satisfying the high-type boundary conditions (39)-(41). Then we express the intersection of these two curves as the root of a polynomial and show that a unique root exists.

Lemma A.1. *B_L is a well-defined, increasing and continuous and differentiable function.*

Proof. Given a lower boundary (α) solve (35) and (37) to get c_1, c_2

$$c_1(\alpha) = \frac{-u_2(V_L - K_L)}{(u_1 - u_2)} e^{-u_1\alpha} \tag{42}$$

$$c_2(\alpha) = \frac{u_1(V_L - K_L)}{(u_1 - u_2)} e^{-u_2\alpha} \tag{43}$$

Notice that c_1 and c_2 are continuous in α. Recall that $u_1 > 0 > u_2$, hence both c_1 and c_2 are positive, c_1 is decreasing in α, and c_2 is increasing in α.

$$\frac{\partial c_1}{\partial\alpha} = \frac{u_1 u_2(V_L - K_L)}{(u_1 - u_2)} e^{-u_1\alpha} = -u_1 c_1 < 0 \tag{44}$$

$$\frac{\partial c_2}{\partial\alpha} = \frac{-u_2 u_1(V_L - K_L)}{(u_1 - u_2)} e^{-u_2\alpha} = -u_2 c_2 > 0 \tag{45}$$

Using boundary condition (36) gives an implicit expression for B_L

$$c_1(\alpha)e^{u_1 B_L(\alpha)} + c_2(\alpha)e^{u_2 B_L(\alpha)} + K_L - \Psi(B_L(\alpha)) = 0 \tag{46}$$

Ψ is continuous and therefore so too is B_L. Using the implicit function theorem, differentiate the expression above to get

$$\frac{\partial c_1}{\partial\alpha}e^{u_1 B_L} + \frac{\partial c_2}{\partial\alpha}e^{u_2 B_L} + c_1 e^{u_1 B_L}\left(u_1\frac{\partial B_L}{\partial\alpha}\right) + c_2 e^{u_1 B_L}\left(u_1\frac{\partial B_L}{\partial\alpha}\right) - \frac{\partial\Psi}{\partial\beta}\frac{\partial B_L}{\partial\alpha} = 0$$

then plugging (44) and (45) and rearranging terms gives

$$\frac{\partial B_L}{\partial\alpha} = \frac{u_1 c_1 e^{u_1 B_L} + u_2 c_2 e^{u_2 B_L}}{u_1 c_1 e^{u_1 B_L} + u_2 c_2 e^{u_2 B_L} - \Psi'(B_L)} = \frac{F_L'(B_L)}{F_L'(B_L) - \Psi'(B_L)} > 0$$

Where the last inequality follows from the fact that F_L is convex, $F_L'(\alpha) = 0$ and that the low type strictly prefers to accept at β (i.e. $F_L'(\beta) - \Psi'(\beta) > 0$). See Lemma A.6) for a proof of these statements. ■

Lemma A.2. *There exists a lower bound ($\underline{\beta}_H$) such that for all $\beta < \underline{\beta}_H$, there does not exists a corresponding lower boundary that solves (39), (40) and (41). However, a solution exists for all $\beta \geq \underline{\beta}_H$ and the function $B_H : (-\infty, \infty) \to \left(\underline{\beta}_H, \infty\right)$ is well-defined, increasing, continuously and differentiable.*

Proof. In any equilibrium, $\Psi(\beta)$ must be greater than K_H since the high type cannot be accepting an offer below his outside option. Thus, we can restrict attention to search for a $\beta \in [\underline{\beta}, \infty)$ where $\underline{\beta} \equiv \Psi^{-1}(K_H)$.[39] For a given a β, we can solve (39) and (40) to get

$$d_1(\beta) = \frac{\Psi'(\beta) + q_2(K_H - \Psi(\beta))}{(q_1 - q_2)} e^{-q_1\beta} \tag{47}$$

$$d_2(\beta) = \frac{q_1(\Psi(\beta) - K_H) - \Psi'(\beta)}{(q_1 - q_2)} e^{-q_2\beta} \tag{48}$$

Note that $d_1 > 0$ for all $\beta \in [\underline{\beta}, \infty)$. Using (41) we get an implicit expression for B_H^{-1}

$$d_1(\beta)q_1 e^{q_1 B_H^{-1}(\beta)} + d_2(\beta)q_2 e^{q_2 B_H^{-1}(\beta)} = 0 \tag{49}$$

In order for (49) to hold, a new lower bound on β arises: the first term on the LHS is always positive and thus the second must be negative. Since $q_2 < 0$, (49) requires that $d_2 > 0$, which from (48) then requires that $\Gamma(\beta) \equiv q_1(\Psi(\beta) - K_H) - \Psi'(\beta) > 0$. This expression is negative for β small (note $\Gamma(\underline{\beta}) < 0$), eventually

[39] Under the signaling condition $\underline{\beta}$ is not well defined.

increasing and tends to $q_1(V_H - K_H) > 0$ as β goes to infinity. It has a unique real root which we denote by $\underline{\beta}_H$. Solving (49) we get

$$B_H^{-1}(\beta) = \frac{1}{q_1 - q_2} \ln\left(-\frac{q_2}{q_1}\frac{d_2(\beta)}{d_1(\beta)}\right)$$

Finally we show that $\frac{d_2(\beta)}{d_1(\beta)}$ is increasing for $\beta > \underline{\beta}_H$ which completes the proof of monotonicity. Let $G(\beta) \equiv -q_2(\Psi(\beta) - K_H) + \Psi(\beta)$ and rewrite $d_2(\beta)/d_1(\beta)$ as $\frac{\Gamma(\beta)}{G(\beta)}e^{(q_1-q_2)\beta}$. Differentiate with respect to β to get

$$\frac{d}{d\beta}\left(\frac{d_2(\beta)}{d_1(\beta)}\right) = \left((q_1 - q_2)\frac{\Gamma}{G} + \frac{\Gamma' G - G'\Gamma}{G^2}\right)e^{(q_1-q_2)\beta}$$

It suffices to show that $(q_1 - q_2)\Gamma G + \Gamma' G - G'\Gamma > 0$, and $\Gamma' > 0$ for all $\beta > \underline{\beta}_H$. $G > 0$ everywhere so the second term is strictly positive. Combining the first and the third gives

$$\begin{aligned}
\Gamma((q_1 - q_2)G - G') &= \Gamma\left((q_1 - q_2)(-q_2(\Psi - K_H) + \Psi') + q_2\Psi' - \Psi''\right) \\
&= \Gamma\left(\underbrace{-q_2(q_1 - q_2)(\Psi - K_H) + q_2\Psi'}_{+} + \overbrace{(q_1 - q_2)\Psi' - \Psi''}^{+}\right) \\
&> 0
\end{aligned}$$

Where the first (+) follows from the fact that $q_1(\Psi - K_H) > \Psi'$ above $\underline{\beta}_H$ and $(q_1 - q_2) > q_1 > 0$ and the second (+) follows from $\Psi' > \Psi''$. The above shows that d_2/d_1 is increasing and $-q_2/q_1$ is positive therefore B_H^{-1} is increasing and so too is B_H. ■

Lemma A.3. *The two curves B_H, B_L intersect exactly once: specifically B_L crosses B_H from below.*

Proof. Let $\underline{\alpha}_\epsilon \equiv B_H^{-1}(\underline{\beta}_H + \epsilon)$, for $\epsilon > 0$, note that $\underline{\alpha}_\epsilon \to -\infty$ and $B_H(\underline{\alpha}_\epsilon) \to \underline{\beta}_H$ as $\epsilon \to 0$ where $\underline{\beta}_H$ is finite for all $s > 0$. As α goes to $-\infty$, $c_1 \to \infty$ and $c_2 \to 0$ therefore $\lim_{\alpha\to-\infty} B_L(\alpha) = -\infty$ and so clearly $B_L(\underline{\alpha}_\epsilon) < B_H(\underline{\alpha}_\epsilon)$ for ϵ sufficiently small. Hence the curves must intersect and B_L crosses from below.

For uniqueness of the intersection we make another change of variables from log-likelihood space to likelihood space in order to express the intersection as the root of a polynomial. We use $\tilde{z}$ to denote e^z (similarly for $\tilde{\alpha}, \tilde{\beta}$) and $\tilde{\Psi}(y) \equiv \Psi(\ln(y))$ so that $\tilde{\Psi}(\tilde{z}) = \Psi(z)$ (similarly for $\tilde{F}_\theta$). Making the change of variables and solving (47), (48) and (49) gives an expression for $\tilde{\alpha}$ in terms of $\tilde{\beta}$ that solves the high types equations.

$$x_H(\tilde{\beta}) \equiv \tilde{\beta}\left(\frac{q_2}{q_1}\left(\frac{\tilde{\Psi}' + q_1(K_H - \tilde{\Psi})}{\tilde{\Psi}' + q_2(K_H - \tilde{\Psi})}\right)\right)^{\frac{1}{q_1-q_2}} \tag{50}$$

Let $f(y) \equiv \frac{y}{x_H(y)}$ and note that

$$\begin{aligned}
sign(f'(y)) &= sign\left(\frac{d}{dy}\left(\frac{q_2(\tilde{\Psi} - K_H) - \tilde{\Psi}'}{\tilde{\Psi}' + q_1(K_H - \tilde{\Psi})}\right)\right) \\
&= sign\left((q_2 - q_1)(\tilde{\Psi}''(K_H - \tilde{\Psi}) + (\tilde{\Psi}')^2)\right) \\
&< 0
\end{aligned}$$

Solving the first two low type equations (42), (43) gives

$$\begin{aligned}
c_1(\alpha) &= \frac{-u_2(V_L - K_L)}{u_1 - u_2}e^{-\alpha u_1} = A\tilde{\alpha}^{-u_1} \\
c_2(\alpha) &= \frac{u_1(V_L - K_L)}{u_1 - u_2}e^{-\alpha u_2} = B\tilde{\alpha}^{-u_2}
\end{aligned}$$

where $A \equiv \frac{-u_2(V_L - K_L)}{u_1 - u_2}$ and $B \equiv \frac{u_1(V_L - K_L)}{u_1 - u_2}$. Plugging these into (46) gives

$$A\left(\frac{\tilde{\beta}}{\tilde{\alpha}}\right)^{u_1} + B\left(\frac{\tilde{\beta}}{\tilde{\alpha}}\right)^{u_2} + K_L - \tilde{\Psi}(\tilde{\beta}) = 0 \tag{51}$$

We have reduced the problem to solving two equations: (50) and (51). Substituting f for $\frac{\tilde{\beta}}{\tilde{\alpha}}$ into (51) gives a polynomial function, which we denote by Λ. Any real root of Λ is a solutions to the six boundary conditions.

$$\Lambda(y) \equiv A(f(y))^{u_1} + B(f(y))^{u_2} + K_L - \tilde{\Psi}(y) \tag{52}$$

Let $\underline{\tilde{\beta}}_H \equiv \exp(\underline{\beta}_H)$ denote the lowest possible upper boundary in likelihood space and note that $f \to \infty$ as $y \to \underline{\tilde{\beta}}_H$. Since $A > 0$, Λ also goes to infinity as $y \to \underline{\tilde{\beta}}_H$. On the other hand, $f(y) \to 1$ as $y \to \infty$ and so the limit of Λ is

$$\begin{aligned} \lim_{y \to \infty} \Lambda(y) &= A + B + K_L - V_H \\ &= V_L - V_H < 0 \end{aligned}$$

Λ has at least one root greater than $\underline{\tilde{\beta}}_H$ implying existence of a solution. To prove its uniqueness, we show that Λ is decreasing for all $y > \underline{\tilde{\beta}}_H$. Taking the derivative and simplifying gives

$$f' f^{-1} \left(Au_1 f^{u_1} + Bu_2 f^{u_2}\right) - \tilde{\Psi}'$$

Since $f' < 0$, $f > 0$ and $\tilde{\Psi}' > 0$, it is sufficient to show that the term inside the parentheses is positive. Since $Au_1 = -Bu_2 > 0$, this requires only that $f^{u_1} - f^{u_2} > 0$ which follows from $f > 0$, $u_1 > 0 > u_2$. ■

Proof of Proposition 3.2. Follows immediately from Lemmas A.1, A.2 and A.3. ■

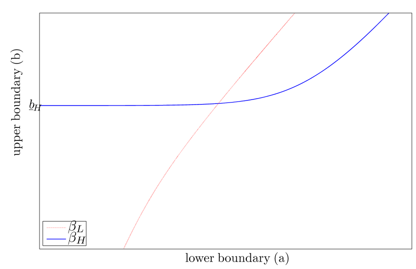

Figure 14: Beta Curves

A.2 Verification of Equilibrium Strategies

Before verifying that the strategies in Ξ^* constitute an equilibrium, we return to the Ψ-game. The next lemma proves that the optimal policy for the high type in the Ψ-game is to accept for all $z > \underline{\beta}_H$.

Lemma A.4. *The optimal policy in the single decision maker's problem is to accept Ψ for all $z > \underline{\beta}_H$ and reject at all $z < \underline{\beta}_H$.*

Proof. The form of the optimal policy is a standard result in optimal stopping (Dixit and Pindyck, 1994) under the following two conditions: (1) the marginal benefit of waiting an additional dt before accepting is (in the limit) monotonic and (2) the CDF of $\hat{Z}_{t+s}$ for shifts uniformly to the right when $\hat{Z}_t$ increases for all $s, t \geq 0$. In the problem stated above, the second condition is immediate. However, the marginal benefit of waiting an instant of time before accepting is

$$\frac{s^2}{2}(\Psi' + \Psi'') - r(\Psi - K_H) \tag{53}$$

where the first term represents the expected increase in the offer tomorrow and the second term is the loss in flow from delaying trade. This expression is single peaked and *eventually* negative. Let β^{**} denote the argument that maximizes (53) and $\bar{\beta}$ denote maximal point for which it is non-negative. Applying the result to the restricted state space $\Omega_1 = (-\infty, \beta^{**})$ implies that if accepting is optimal at some $\hat{\beta} < \beta^{**}$ then the optimal policy is to accept everywhere below $\hat{\beta}$. Of course this cannot be optimal because it can be improved upon by never accepting below $\underline{\beta}$. Thus the optimal policy is to always reject in Ω_1. In $\Omega_2 = (\beta^{**}, \infty)$ the result tells us that if it is ever optimal to accept than it must be optimal to accept everywhere above. Acceptance *must* be optimal somewhere hence the policy stated in the lemma is optimal.[40] To see that the optimal boundary must be $\underline{\beta}_H$, notice that the value matching and smooth pasting conditions at the upper boundary are exactly the same as in the high type's problem (39) and (40). The differential equation is also the same. The only difference is that there is not a lower boundary. Taking $\alpha \to -\infty$ in the high type's problem implies that the upper boundary must be $\underline{\beta}_H$. ■

Proof of Lemma 2.1. Follows immediately from Lemma A.4. p_H^* is equivalent to $\underline{\beta}_H$ after making the change of variables back to probability space. ■ A step in verifying that the seller's strategy is optimal is to show that the high type cannot do better by rejecting Ψ over some interval of $z > \beta^*$.

Lemma A.5. $\beta^* > \bar{\beta}$

Proof. Lemmas A.2 and A.3 showed that $\beta^* > \underline{\beta}_H$, the result follows immediately if we show that $\underline{\beta}_H \geq \bar{\beta}$. Suppose the opposite is true. Then the single decision maker would prefer to reject for all $\beta \in (\underline{\beta}_H, \bar{\beta})$ contradicting lemma A.4. ■

Lemma A.6. *$F_H(z) > \Psi(z)$ and increasing for all $z \leq \beta$ and $F_L(z) > V_L$ and increasing for all $z > a$.*

Proof. Twice differentiating F_H gives $F_H''(z) = d_1 q_1^2 e^{q_1 z} + d_2 q_2^2 e^{q_2 z} > 0$ because both terms are strictly positive. F_H convex and $F_H'(\alpha) = 0$ implies $F_H'(z) > 0$ for all $z \in (\alpha, \beta)$. The change of variable to likelihood space preserves the monotonicity (recall we use the ˜ to denote functions in likelihood space). $\tilde{F}$ is also convex because $\tilde{F}_H''(x) = d_1 q_1^2 x^{q_1} + d_2 q_2^2 x^{q_2} > 0$. Since $\tilde{F}_H'(\alpha) = 0 < \tilde{\Psi}'(\alpha)$ and $\tilde{F}_H'(\beta) = \tilde{\Psi}'(\beta)$, the convexity of $\tilde{F}_H$ and concavity of $\tilde{\Psi}$ implies that $\tilde{F}_H' < \tilde{\Psi}'$ in the no trade region. Therefore, $\tilde{F}_H$ is *flatter* than $\tilde{\Psi}$ and since they are equal at β, F_H lies everywhere above for all $z \in (\alpha, \beta)$. Twice differentiating F_L gives $F_L''(z) = c_1 u_1^2 e^{u_1 z} + c_2 u_2^2 e^{c_2 z} > 0$. Therefore F_L is also increasing (recall $F_L(\alpha) = 0$) and convex over (α, β). The result then follows since $F_L(\alpha) = V_L$. ■

Lemma A.7. *Buyers cannot deviate profitably.*

Proof. Recall that $F_H > \Psi > F_L$ for all $z < \beta$. Any offer accepted by the high type in this region will also be accepted by the low type and hence earn negative expected profits. Since $F_L > V_L$ for $z \in (\alpha, \beta)$, an offer which attracts only low types also loses money. ■

The final step in the verification is to prove that the sellers strategy is optimal.

Proof of Proposition 3.3. Start with the high type at $Z_0 = z$. Consider an intermediary who offers to pay the high type $F_H(z)$ immediately or $F_H(Z_t)$ at any point in the future. At each point in time, the seller can choose whether to accept an offer from *either* the intermediary, a buyer or neither. Clearly the high type must do at least as well in the game with both the intermediary and buyers as she does in the game with only buyers. A strategy in this game is simply a stopping time. Let $G_H(z)$ denote the maximal value to the high type of this stopping problem starting from the initial state z. Let τ be an arbitrary stopping

[40]The expression in (53) is strictly decreasing for η large. In this case $\beta^{**} = \bar{\beta} = -\infty$, $\Omega_1 = \emptyset$, $\Omega_2 = \mathbb{R}$.

time. It is immediate that $G_H(z) \geq F_H(z)$.

$$
\begin{aligned}
G_H(z) &= \sup_{\tau} E_z^H \left\{ \int_0^{\tau} k_H e^{-rt} dt + e^{-r\tau} [w(Z_\tau) \vee F_H(Z_\tau)] \right\} \\
&= \sup_{\tau} E_z^H \left\{ \int_0^{\tau} k_H e^{-rt} dt + e^{-r\tau} F_H(Z_\tau) \right\}
\end{aligned}
$$

where the second equality follows from $F_H(z) \geq w(z)$ for all z. Let $f_H(z,t) \equiv \int_0^t k_H e^{-rs} ds + e^{-rt} F_H(z)$. The payoff from following an arbitrary stopping rule τ is $E_z^H\{f(Z_\tau, \tau)\}$. $F_H \in C^1$ everywhere and it is twice continuously differentiable at all points other than $\{a, b\}$, therefore $f_H \in C^2$ almost everywhere. Hence, we can apply Ito's lemma to f_H

$$
f_H(Z_\tau, \tau) = f_H(Z_0, 0) + e^{-r\tau} \left[\int_0^{\tau} \mathcal{A}^H F_H(Z_t) dt + \int_0^{\tau} s F_H'(Z_t) dB_t \right] \tag{54}
$$

where $\mathcal{A}^H$ is the differential operator satisfying $\mathcal{A}^H F_H \equiv \frac{s^2}{2}(F_H' + F_H'') - r(F_H - K_H)$. By construction $\mathcal{A}^H F_H(z) = 0$ for all $z \in [a, b]$. $F_H(z) = \Psi(z)$ for all $z > \beta$, and Lemma A.5 implies that $\mathcal{A}^H F_H(z) = \mathcal{A}^H \Psi(z) < 0$ for all $z > b$. Using equation (54), the expected payoff to the high type under a stopping rule τ as

$$
\begin{aligned}
E_z^H[f_H(Z_\tau, \tau)] &= f_H(z, 0) + E_z^H \left[\int_0^{\tau} \mathcal{A}^H F_H(Z_t) dt \right] \\
&= F_H(z) + e^{-r\tau} E_z^H \left[\int_0^{\tau} \mathcal{A}^H F_H(Z_t) dt \right] \\
&\leq F_H(z)
\end{aligned} \tag{55}
$$

where the final inequality is strict if $\tau > T(b) \equiv \inf\{t : P_t \geq b\}$ with positive probability. The above is true for any arbitrary policy, therefore the high type seller can do no better than accept the intermediary's offer immediately and does strictly worse by waiting longer than $T(b)$. Clearly then in the game with no intermediary the high type's maximal payoff is no greater than $F_H(z)$, which proves that the high type is playing optimally in Ξ^*.

The proof for the low type follows a similar argument. The expected payoff to the low type under a stopping rule τ is analogous to (55) where $\mathcal{A}^H F_H$ is replaced with $\mathcal{A}^L F_L \equiv \frac{s^2}{2}(-F_L' + F_L'') - r(F_L - K_L)$. Of course when Ψ is offered, any optimal policy for the low type must accept immediately. Moreover, since $\mathcal{A}^L F_L(z) = 0$ for all $z \in (\alpha, \beta)$, a low type can do no better than F_L in the game with the intermediary. The final detail is to note that any policy that stops immediately at all $z \geq \beta$ achieves an expected payoff of $F_L(z_0)$.[41] Therefore the low type mixing at $z \leq a$ is consistent with playing a best response. ∎

A.2.1 Reflection of Equilibrium Beliefs

Proof of Lemma 3.1. Z_t must have a reflecting barrier at α (see Proposition 3.1), which requires the pushing process $\hat{Y}_t = \max\{\alpha - \hat{M}_t, 0\}$ (Harrison, 1985). The claim is that if the low type rejects with probability $e^{-\hat{Y}_t}$ then P_t will follow $\hat{P}_t$ with a reflecting barrier at a implying the result. Suppose first that $\hat{Y}_t = 0$. The low type has always rejected and P_t follows $\hat{P}_t$ which lies everywhere above a since $\alpha - \hat{M}_t \leq 0 \implies \inf_{0 \leq s \leq t} \hat{Z}_t \geq \alpha \implies \inf_{0 \leq s \leq t} \frac{e^{\hat{Z}_t}}{1+e^{\hat{Z}_t}} \geq \frac{e^{\alpha}}{1+e^{\alpha}} \implies \inf_{0 \leq s \leq t} \hat{P}_t \geq a$. Next suppose that $\hat{Y}_t > 0$. By definition $\hat{Y}_t$ only increases when $Z_t = \alpha$ implying that P_t follows $\hat{P}$ everywhere except when

[41]This is statement is true regardless of the policy for $z \leq \beta$. Since the process is reflecting, the low type is indifferent between accepting the intermediaries offer at any $z < \beta$ or continuing.

$P_t = a$. Moreover, if the low type accepts rejects with probability $e^{-\hat{Y}_t}$ then by Bayes' rule

$$\begin{aligned}\inf_{0\le s\le t} P_t &= \frac{e^{\hat{M}_t}/(1+e^{\hat{M}_t})}{e^{\hat{M}_t}/(1+e^{\hat{M}_t}) + e^{-\hat{Y}_t}/(1+e^{\hat{M}_t})} \\ &= \frac{e^{\hat{M}_t}}{e^{\hat{M}_t}+e^{\hat{M}_t-\alpha}} \\ &= \frac{e^{\alpha}}{1+e^{\alpha}} \equiv a\end{aligned}$$

■

Proof of Proposition 3.1. Let $\hat{Y}_t$ denote the pushing process coming from the low type acceptance. Four conditions must be satisfied in equilibrium:

1. $Z_t = \hat{Z}_t + \hat{Y}_t$: the equilibrium beliefs are affected only by the news and the strategy of the low type seller.
2. $\hat{Y}_t$ is increasing and continuous with $\hat{Y}_0 = y_0$: $y_0 > 0$ if $Z_0 < \alpha$ (in which case Z jumps instantaneously) but there are no atoms following an initial jump. Moreover, the low type cannot "add more" of herself and so the pushing process can only increase Z.
3. $Z_t \ge \alpha\ \forall t > 0$ prior to acceptance: in equilibrium beliefs cannot drop below a (see § 3.1)
4. $\hat{Y}_t$ is only increasing at times t such that $Z_t = \alpha$: the equilibrium belief process is pushed by low type acceptance only at $p = a$ for all $t > 0$.

The above is known as the Skorohod problem (Karatzas and Shreve (2004)). The solution is a process, $\hat{Y}$, known as the "local time of Z on the boundary," where $Z = \hat{Z} + \hat{Y}$ follows reflected B.M. with a lower boundary at α. The process $\hat{Y}$ that is constructed in Lemma 3.1 solves this problem. ■
Proof of Corollary 3.1. Let $t = 0$, $\hat{Z}_0 = \alpha \implies \hat{Y}_0 = 0$. Apply the formula from Proposition (3.1) to get $E_{\hat{Y}_0=0}[1-e^{-\hat{Y}_\Delta}] = E_{L_0=0}[1-e^{-\hat{M}_\Delta}] \approx 1 - E_{\hat{M}_0=0}[\hat{M}_\Delta] = 1 + E_0[\inf_{0\le s\le\Delta}\hat{Z}_s] = 1 - ks\sqrt{\Delta}$ where $k = \sqrt{2/\pi}$. Next, let $f(x) = \Pr_x\{\text{Low type reaches } b\}$. Starting from x, at $p = b$ beliefs must be consistent, $b = \frac{x}{x+(1-x)f(x)}$. Solving for $f(x)$ yields the result. ■

A.2.2 Welfare Results

Proof of Proposition 4.1. For (1), $\tilde{F}_H(p) > \Psi(p)\ \forall p < \tilde{b}$ and so immediately we have $\tilde{F}_H(p) > F_H(p)$ $\forall p \in [b, \tilde{b})$. Of course $\tilde{F}_H(p) = F_H(p) = \Psi(p)\ \forall p \ge \tilde{b}$. All that is left is to show that $\tilde{F}_H(p) > F_H(p)\ \forall p < b$. Starting from any $p \in (\tilde{a}, b)$ the high type is waiting to get to the payoff at b. $\tilde{F}_H(b) > F_H(b) = \Psi(b)$ and the high type gets to b faster under $\tilde{s}$ implying that $\tilde{F}_H(p) > F_H(p)\ \forall p \in (\tilde{a}, b)$. F_H is increasing from a to $\tilde{a}$ while $\tilde{F}_H$ remains constant implying the strict inequality for all $p \le \tilde{a}$. (2) is immediate since $\tilde{F}_L(p) = V_L$ $\forall p \le \tilde{a}$, $\tilde{a} > a$, and $F_L(p)$ begins increasing from V_L at a.

From (1) and (2), $p\tilde{F}_H(p) + (1-p)\tilde{F}_L(p) > pF_H(p) + (1-p)F_L(p)$ for all $p \le a$. By continuity there exists some $\hat{p}$ such that $\tilde{L}(\hat{p}) > L(\hat{p})$ for all $p < \hat{p}$ yielding (3). $L(p) = 0$ for all $p \ge b$ whereas $\tilde{L}(p) > 0$ for $p \in [b, \tilde{b})$. By continuity there must exist a $p' \in (a, b)$ such that $\tilde{L}(p') > L(p)$ for all $p \in (p', \tilde{b})$. $p' = \hat{p}$ follows from $F'_H > \tilde{F}'_H$ and $F'_L > \tilde{F}'_L$ for all $p \in (a, b)$, which implies single crossing of L and $\tilde{L}$ and completes the proof. ■

A.3 Proofs for Results under the Signaling Condition

Definition A.1. *Let $\eta_1 > 0$ be s.t. $\frac{(q_1-1)^2}{4q_1} = \frac{V_L-K_H}{V_H-V_L}$, where $q_1 \equiv 1/2(\sqrt{1+8\eta}-1)$ as defined earlier.*

The following can be shown using analysis similar to that in § A.1.1 (also see the proof of Theorem 5.1).

Fact A.1. *Under the **signaling condition**, for all $\eta < \eta_1$:*

- *The "marginal benefit" to delay for a high type seller is positive for at least some beliefs.*
- *A pair (α^*, β^*), $\alpha^* < \beta^*$ exists which solves the six boundary conditions (35)-(41)*

Proof of Lemma 2.2. There are two possible cases. The first case is when $\frac{s^2}{2}(\Psi' + \Psi'') - r(\Psi - K_H) < 0$ for all z. In this case, it is optimal for the high type to accept right away, $F_H(z) = \Psi(z)$ for all z. The standard argument for optimality applies since $\mathcal{A}^H F_H < 0$ for all z and $F_H \in C^2$ (see Proof of Proposition 3.3).

The second case is when $\frac{s^2}{2}(\Psi' + \Psi'') - r(\Psi - K_H) > 0$ for at least some z. It is straightforward to verify that for all $\eta < \eta_1$ the condition for the second case is satisfied. Stopping right away cannot be optimal because then $F_H = \Psi$ and $\mathcal{A}^H \Psi(z) > 0$ for at least some z, violating a necessary condition for optimality (Oksendal, 2007).

The high type must accept for at least some $z \in \Omega_1$ (see Proof of Lemma A.4), otherwise $\lim_{z \to -\infty} F_H(z) = K_L$, and the policy that accepts Ψ everywhere that F_H is less than V_L does strictly better. Let $\underline{\beta}_s$ denote the largest $z \in \Omega_1$ such that the high type accepts. By the same argument given in Lemma A.4, it is optimal to accept for all $z < \underline{\beta}_s$. For $z \in \Omega_2$, the argument made in Lemma A.4 also applies. The unique pair $(\underline{\beta}_s, \overline{\beta}_s)$ that solves the high types' optimal stopping problem is characterized by the differential equation in (32) along with the boundary conditions: $F_H'(\underline{\beta}_s) = \Psi'(\underline{\beta}_s)$, $F_H'(\overline{\beta}_s) = \Psi'(\overline{\beta}_s)$ (smooth-pasting) and $F_H(\underline{\beta}_s) = \Psi(\underline{\beta}_s)$, $F_H(\overline{\beta}_s) = \Psi(\overline{\beta}_s)$ (value-matching). ■

Proof of Proposition 5.1. *Existence of a no-trade region*: We argue by contradiction. First, suppose the high type traded with positive probability for almost all p. Then $F_H = \Psi$. For $\eta < \eta_1$, there exists a p such that a deviation will be profitable. To see this, consider the off-path belief process $P = \hat{P}$. Then the high type's (SP) is equivalent to the Ψ-game. Appealing to Lemma 2.2 yields the result. Belief monotonicity implies that $dP \geq d\hat{P}$, making rejection always at least as profitable as in the Ψ-game. Therefore, there exists at least one interval in $[0, 1]$ where the high type trades with probability zero.

Second, any offer which the high type accepts will also be accepted by the low type. Therefore, if there does not exist a no-trade region, then the low type is trading with positive probability for almost all $p \in [0, 1]$. Consider a belief p where the high type trades with probability zero. The low type is trading at a price of V_L with positive probability implying $F_L(p) = V_L$. As argued in the proof of Theorem 3.1, the low type is not trading with probability one or at a flow rate over an interval. Rejection at p causes a discrete increase in P to $j(p)$. From the assumption that trade occurs with positive probability almost everywhere, the low type must have a positive probability of trading at a price arbitrarily close to $\Psi(j(p))$ in the next dt. It follows that $F_L(j(p)) > V_L$, breaking the indifference condition necessary for the low type's mixing behavior at p. This contradiction implies that there must exist an interval (a, b) such that trade occurs with probability zero for all $p \in (a, b)$.

Equilibrium properties: Consider any no-trade interval (a, b). Lemma **??** and NDVF imply that $F_L(b) = F_H(b) = \Psi(b)$. Further, either $F_L(a) = V_L$ or $F_L(a) = \Psi(a)$. If $F_L(a) = V_L$, then (a, b) must satisfy all of the conditions set forth in the proof of Theorem 3.1. If $F_L(a) = \Psi(a)$, then $F_H(b) = \Psi(b)$ as well. The argument in Step 4 in the proof of Theorem 3.1 implies that at both a and b the high type must be indifferent between accepting and rejecting given (1) $dP = d\hat{P}$ conditional on rejection, and (2) $F_H(p) = \Psi(p)$ for all $p \in (a - \varepsilon, a) \cup (b, b + \varepsilon)$.

Because the systems of equations that govern no-trade regions have finitely many solutions (see proofs of Lemma 2.2 and Theorem 5.1), there are finitely many no-trade regions. Consider the no-trade region $(\overline{a}, \overline{b})$ such that $\overline{a} \geq b$ for all other trade regions (a, b). Trade occurs with positive probability at all $p > \overline{b}$. By F_L non-decreasing, trade cannot occur at a price of V_L for $p > b$, therefore $F_L(p) = F_H(p) = \Psi(p)$ for all $p > b$.

Finally, there are two possibilities for $F_L(\overline{a})$.

1. If $F_L(\overline{a}) = V_L$, then F_L non-decreasing and Lemma 3.2 imply that $F_L(p) = V_L$ for all $p < a$. The equilibrium is therefore Ξ^* (see proof of Theorem 5.1 for details).

2. If $F_L(\overline{a}) = \Psi(\overline{a})$ then there exists a *type 2* no-trade region $(\overline{a}, \overline{b})$ where both types of sellers trade at both $\overline{a}$ and $\overline{b}$.

■

Proof of Proposition 5.2. It is sufficient to show that $\lim_{s \to \infty} F_L'(a) \to 0^-$. Let (x, y) denote the pair of boundaries in a *type 2* no-trade region (in likelihood space). The pair is determined solely by the high types optimal stopping problem therefore x must be below the point at which the marginal benefit of the high type is positive. Recall the sign of the marginal benefit of the high type is the same as the sign of

$\tilde{\Psi}' + \tilde{\Psi}'' - 2\eta\left(\tilde{\Psi} - K_H\right)$. This expression is single peaked with two real roots for $\eta < \eta_1$. Let $\xi(\eta)$ denote the lower root. The high type can only possibly be accepting above the upper root and below ξ. WLOG normalize $V_L = 0$, $V_H = 1$ (and hence the signaling condition implies $K_H < 0$). Straightforward calculation shows that $\lim_{\eta \to 0} \frac{\xi(\eta)}{\eta} = \frac{-K_H}{2}$. Therefore the lower boundary must converge to zero at a rate at least proportional to η ($O(\eta)$ hereafter). Recall that $F_L'(x) = u_1 c_1 x^{u_1} + u_2 c_2 x^{u_2}$. As $\eta \to 0$: $u_1 \to 1^+$, $c_1 \to 0^+$ (at rate $\sqrt{\eta}$), $x^{u_1} \to 0^+$ at a rate weakly bigger than the rate at which $\xi(\eta)^{u_1}$ goes to zero, which from above is $\eta_1^u > \eta$. This implies that $u_1 c_1 x^{u_1}$ is $O(\eta^{3/2})$. In the second term: $u_2 \to 0^-$ (at rate η), $c_2 \to (V_L - K_L)^-$, $x^{u_2} \to 1^-$ implying that the second term is $O(\eta)$. The first term goes to zero faster and hence the second term dominates the sign of the derivative for small η, and since the second term is negative the derivative converges from below. To see this, note first that $sign(F_L'(a)) = sign(F_L'(a)/u_1 c_1 x^{u_1})$ then take the limit as $\eta \to 0$, $\lim_{\eta \to 0} 1 + \frac{u_2 c_2 x^{u_2}}{u_1 c_1 x^{u_1}} \le 1 + \frac{-O(\eta)}{O(\eta^{3/2})} = -\infty$. Hence for all η small enough, $F_L'(a) < 0$ violating NDVF. ■

Proof of Theorem 5.1. The proof follows in three steps. Steps (2) and (3) are shown above, the proof of (1) (below) is similar to the proof under the lemons condition and references a number of results from § A.2 and § A.1.1.

1. We prove the existence of a unique candidate pair (α, β) and verify it is an equilibrium.
2. Proposition 5.1 implies that either the equilibrium is Ξ^* or it has a *type 2* no-trade region.
3. Proposition 5.2 shows why a *type 2* no-trade region violates NDVF.

Recall from the proof of Lemma A.2, if (α, β) is a solution then $\Gamma(\beta) \equiv q_1(\Psi(\beta) - K_H) - \Psi'(\beta)$ must be non-negative. Let $\tilde{\Gamma}(y) \equiv \Gamma(\ln(y))$, which is quadratic in y and has two real roots provided $\eta < \eta_1$. From now on consider only $\eta < \eta_1$. Let $\tilde{z}_1 < \tilde{z}_2$ denote the two real roots of $\tilde{\Gamma}$ and note that $\tilde{\Gamma}$ is negative for $\tilde{z} \in (\tilde{z}_1, \tilde{z}_2)$. Lemmas A.2 and A.3 can be applied to show that there exists exactly one solution for which $\tilde{\beta} > \tilde{z}_2$.[42] Moreover, any other solution (of which there can only finitely many since it can be expressed as the root of a polynomial) involves $\tilde{\beta} < \tilde{z}_1$. Proposition 5.1 guarantees that a no-trade region below $\tilde{z}_1$ must occur in tandem with a type 2 no-trade region, which is ruled out by Proposition 5.2 as $\eta \to 0$. The verification steps follows exactly from § A.2 provided that $\eta < \eta_1$. ■

[42] Note that $y_2 = \exp(\underline{\beta}_H)$

B Appendix: Discrete-Time Analog

The purpose of this appendix is to establish the strong connection between the continuous-time game in the body of the paper and a discrete-time analog. First, when time periods are short, there exists a unique equilibrium with nearly identical structure to the unique equilibrium in continuous time (Ξ^*). Second, as we take the period length to zero, the sequence of discrete-time equilibria converges to Ξ^*.

There are three reasons to carry out this exercise. First, some concepts and intuitions may be easier to capture in a discrete-time world. The reader can appeal to whichever framework he finds most useful for understanding any aspect of the game. Second, we can regard this as a robustness check. It is re-assuring that the limit of the equilibrium of the discrete-time model is the equilibrium of the model posed directly in continuous time. Finally, after seeing the "integer problems" encountered in discrete time, the reader should be convinced that working in continuous time has considerable advantages.

For the sake of brevity, some arguments given here are less formal than the proofs in Appendix A. There is, however, enough content that the interested reader should understand the mechanics of the discrete-time equilibrium and be able to fill in desired details.

B.1 Setup

Everything is as in § 2 with the following exceptions.

1. Buyers arrive to make offers only at times $t = \Delta, 2\Delta, ...$, where Δ is "small".
2. For each type, the news process follows a random walk.

Beginning at any $X^\theta_{n\Delta}$,

$$X^H_{(n+1)\Delta} = \begin{cases} X^H_{n\Delta} + \sigma\sqrt{\Delta} & \text{with probability } q^\Delta \\ X^H_{n\Delta} - \sigma\sqrt{\Delta} & \text{with probability } 1 - q^\Delta \end{cases}$$

$$X^L_{(n+1)\Delta} = \begin{cases} X^L_{n\Delta} + \sigma\sqrt{\Delta} & \text{with probability } 1 - q^\Delta \\ X^L_{n\Delta} - \sigma\sqrt{\Delta} & \text{with probability } q^\Delta \end{cases}$$

Where $q^\Delta = \frac{1}{2}\left(1 + \frac{\mu}{\sigma}\sqrt{\Delta}\right)$. It is well known that these processes converge to the continuous-time processes in the body (Karatzas and Shreve, 2004). The assumption that the high type will get good news ($X_{(n+1)\Delta} > X_{n\Delta}$) with the same probability that the low type gets bad news is convenient. Starting from any prior, the belief of a Bayesian who observes one instance of good news followed by one instance of bad news will arrive back at his prior. Therefore, given parameters μ, σ and Δ, for any z, there exists a unique grid G_z on which beliefs based only on observation of X will reside.[43] G_z is a countable sequence of points that span the real line. z is an element of G_z, and a point $z' \in G_z$ that is n points above z is the posterior a Bayesian would arrive at starting from a prior of z after observing n greater instances of good news than bad news. We denote such a z' as z_{+n}. z_{-n} is defined analogously.

Unlike the continuous-time model, it is no longer without loss to study equilibria where the buyers always play pure strategies. We will see that in some cases the unique equilibrium will require precise mixing of w in certain states. The final alteration we make for the discrete-time model is a minor tightening of our restriction on off-path beliefs. In continuous time we restricted beliefs to be at least as favorable as the process that ignores off-path events (see § 2.2). For analysis of the discrete-time model, we restrict off-path beliefs to be exactly those that ignore off-path events. That is, beliefs are unchanged by the act of rejecting unexpectedly. This simplifies matters by not having to consider unexpected rejection leading to a higher belief level that resides on a different grid.

B.2 Main Results

Assume the *lemons condition*: $k_H > v_L$. We will show two main results.

[43] As in Appendix A, it will be easier to work in Z-space. The statement, "Beliefs are z," literally means beliefs are $p = \frac{e^z}{1+e^z}$.

Theorem B.1. *When Δ is small, there exists a unique equilibrium. It is characterized by a unique pair (α^*, β^*). Equilibrium play is the same as in Ξ^* with one exception: $w(\beta^*)$ is offered at $z = \beta^*$ with probability γ, which may or may not be uniquely determined, depending on parameters.*

Theorem B.2. *As $\Delta \to 0$, the unique discrete-time equilibrium converges to Ξ^* in the following ways:*

1. *Equilibrium payoffs converge: $F_\theta^\Delta(z) \to F_\theta(z)$ for all z.*
2. *Equilibrium strategies converge for almost all z.*

B.3 B^Δ Functions

The arguments for equilibrium existence, uniqueness and convergence will rely on the discrete-time analogs of the B_θ functions analyzed in Appendix A.

Definition B.1. *Fix a pair (α, β), $\beta \geq \alpha$, $\beta \in G_\alpha$. Define $F_L^\Delta(z|\alpha, \beta)$ as the value to the low type of being at $z \in G_\alpha$, $\alpha \leq z \leq \beta$ given that*

1. *Z evolves based on observation of X.*
2. *$F_L^\Delta(\alpha_{-1}) = V_L$*
3. *$F_L^\Delta(\beta) = \Psi(\beta)$*

Definition B.2. *Fix a pair (α, β), $\beta \geq \alpha$, $\beta \in G_\alpha$. Define $F_H^\Delta(z|\alpha, \beta)$ as the value to the high type of being at $z \in G_\alpha$, $\alpha \leq z \leq \beta$ given that*

1. *Z evolves based on observation of X, except as stated in (3).*
2. *$F_H^\Delta(\beta_{+1}) = \Psi(\beta_{+1})$*
3. *Z reflects off α in the manner discussed in § 3.1.*

Definition B.3. *$B_L^\Delta(\alpha)$ is the maximum $\beta \geq \alpha$, $\beta \in G_\alpha$, such that $F_L^\Delta(\alpha|\alpha, \beta) \geq V_L$.*

Definition B.4. *$B_H^\Delta(\alpha)$ is the minimum $\beta \geq \alpha$, $\beta \in G_\alpha$, such that $F_L^\Delta(\beta|\alpha, \beta) \leq \Psi(\beta)$.*

Recall the definitions from Appendix A:

- $B_L(\alpha)$: The $\beta \geq \alpha$ such that, fixing $F_L(\beta) = \Psi(\beta)$ and $F_L(z) = V_L$ for all $z < \alpha$, the low type is indifferent between accepting or rejecting V_L at $z = \alpha$. It can be found solving the following three equations, which have four unknowns, leaving β in terms of α.

    1. *Value Matching (lower)*: $F_L(\alpha) = V_L$
    2. *Value Matching (upper)*: $F_L(\beta) = \Psi(\beta)$
    3. *Smooth Pasting*: $F_L'(\alpha) = 0$

- $B_H(\alpha)$: The $\beta \geq \alpha$ such that, fixing $F_H(\beta) = \Psi(\beta)$ for all $z > \beta$ and with Z having a reflecting boundary at α, the high type is indifferent between accepting or rejecting $\Psi(\beta)$ at $z = \beta$. It can be found solving the following three equations, which have four unknowns, leaving β in terms of α.

    1. *Value Matching*: $F_H(\beta) = \Psi(\beta)$
    2. *Reflection*: $F_H'(\alpha) = 0$
    3. *Smooth Pasting*: $F_H'(\beta) = \Psi'(\beta)$

From the analysis of the continuous-time game in Appendix A, B_H and B_L are continuous and increasing. They intersect at a single point α^*, below which $B_L < B_H$ and above which $B_L > B_H$. The following convergence result is very useful.

Lemma B.1. *For $\theta \in \{L, H\}$, B_θ^Δ converges to B_θ pointwise, as $\Delta \to 0$.*

The proof of Lemma B.1 is found in § B.6.

B.4 Equilibrium Existence

To establish equilibrium existence and uniqueness, we need to delve a little deeper into the structures of B_L^Δ and B_H^Δ.

Fact B.1. *B_L^Δ is continuous and increasing almost everywhere. There exists an infinite sequence spanning $\mathbb{R}$ such that B_L^Δ is discontinuous at every point in the sequence. B_L^Δ increases at every point of discontinuity.*

Understanding the property is straightforward. Consider an α and its corresponding $\beta = B_L^\Delta(\alpha)$, such that $F_L^\Delta(\alpha|\alpha,\beta) = V_L$. Let $\beta = \alpha_{+n}$. Now, for an $\alpha' = \alpha - \varepsilon$, its corresponding $\beta' = B_L^\Delta(\alpha')$ must be $\alpha'_{+(n-1)}$. Why? Suppose it maintained the same distance in grid points as the original α and β. Then every path of X_t^L results in the same behavior up and down a grid of equal size, endowing the same flow payoff and payoff at the lower terminal node. However, every path that ends at the upper terminal node receives a slightly lower terminal payoff. Because the original terminal payoff of $\Psi(\beta)$ was the exact amount to make $F_L^\Delta(\alpha|\alpha,\beta) = V_L$, the new lower terminal payoff of $\Psi(\beta')$ will cause $F_L^\Delta(\alpha'|\alpha',\beta')$ to fall below V_L. Because waiting is always bad for the low type in expectation, the solution is to decrease β' by one grid point to $\alpha'^{+(n-1)}$. This problem only occurs when $F_L^\Delta(\alpha) = V_L$. If $F_L^\Delta(\alpha) > V_L$, then α' and β' will also be separated by the same number of grid points as the original α, β were. This yields the continuous intervals in B_L^Δ. To summarize, let us explain it from the reverse perspective, one of increasing α. Start with some (α,β) pair separated by n grid points. As we continuously increase α, if we keep β the same number of grid points away, we increase $F_L^\Delta(\alpha)$ as well. Eventually, it raised so high that now $F_L^\Delta(\alpha|\alpha,\alpha_{n+1}) = V_L$, and $B_L^\Delta(\alpha)$ therefore jumps up by one grid point.

Fact B.2. *B_H^Δ is continuous and increasing almost everywhere. There exists an infinite sequence spanning $\mathbb{R}$ such that B_H^Δ is discontinuous at every point in the sequence. B_H^Δ <u>decreases</u> at every point of discontinuity.*

The argument for this property is analogous to the one above. The difference is that as we continuously increase (α,β), keeping their distance in grid points fixed, the payoff to waiting at β_{-1} falls, until $F_\Delta^H(\beta_{-1}) = \Psi(\beta_{-1})$. Therefore, $B_H^\Delta(\alpha)$ must jump down by one grid point when this occurs.

Not surprisingly, the equilibrium results will rely on intersection between B_L^Δ and B_H^Δ. Generically, the two functions intersect. The pointwise convergence of these functions to their continuous-time analogs gives that $B_L^\Delta < B_H^\Delta$ for small α and $B_L^\Delta > B_H^\Delta$ for large α. The fact that the discontinuities in both functions are always on the order of one grid point implies that the only way the two functions will fail to intersect is if they are both discontinuous at some α such that $B_L^\Delta < B_H^\Delta$ for all $\alpha' < \alpha$ and $B_L^\Delta > B_H^\Delta$ for all $\alpha' > \alpha$. Designate this occurrence a "near intersection." In addition to being non-generic, a "near intersection" does not pose a problem for equilibrium existence or uniqueness (see case 3 below).

Identifying (α^*,β^*) and γ

There are two cases to consider: (1) B_L^Δ and B_H^Δ intersect, or (2) they do not. Start with (1), and let α be a point of intersection. α is an element of two half-open intervals $[\underline{\alpha}_L,\overline{\alpha}_L)$ and $[\underline{\alpha}_H,\overline{\alpha}_H)$, where B_θ^Δ is continuous on $[\underline{\alpha}_\theta,\overline{\alpha}_\theta)$. Given that the B_θ^Δ curves intersect at α, $\max\{\underline{\alpha}_L,\underline{\alpha}_H\}$ is also a point of intersection. Let $\alpha^* = \max\{\underline{\alpha}_L,\underline{\alpha}_H\}$.

1. If $\underline{\alpha}_L \le \underline{\alpha}_H$, then $\beta^* = B_\theta^\Delta(\alpha^*)$ for $\theta = L,H$. Define γ to be the probability that $w(\beta^*) = \Psi(\beta^*)$. Recall that $F_L^\Delta(\alpha^*|\alpha^*,B_L^\Delta(\alpha^*)) \ge V_L$ and $F_L^\Delta(\alpha^*|\alpha^*,(B_L^\Delta(\alpha^*))_{+1}) < V_L$. There is then a unique probability $\gamma \in (0,1]$ such that if $w(B_L^\Delta(\alpha^*)) = \Psi(B_L^\Delta(\alpha^*))$ with probability γ and $w((B_L^\Delta(\alpha^*))_{+1}) = \Psi((B_L^\Delta(\alpha^*))_{+1})$ with probability one, the low type's value at α^* will be exactly V_L. Let this be the equilibrium γ.
2. If $\underline{\alpha}_L > \underline{\alpha}_H$, then $(B_\theta^\Delta(\alpha^*))_{-1} < \beta^* < B_\theta^\Delta(\alpha^*)$ for $\theta = L,H$, and any $\gamma \in [0,1]$ is consistent with equilibrium. We will describe how to precisely pin down β^* in the verification stage of this argument.
3. Finally, if B_L^Δ and B_H^Δ do not intersect. As discussed above, there is then a z where they "nearly intersect." This z is α^*. $\beta^* = B_H^\Delta(\alpha^*) = (B_L^\Delta(\alpha^*))_{-1}$ and $\gamma = 0$.

Equilibrium Verification

Taking the three cases in turn:

1. Given the definitions of B_L^Δ, B_H^Δ and our restriction on off-path beliefs, verification of equilibrium strategies is immediate for all $z < \alpha^*$ and for all $z \in G_{\alpha^*}$, including the precise characterization of γ which is necessary to make the low type indifferent for all $z < \alpha^*$. $\alpha^* = \underline{\alpha}_H$ implies that the high type is just indifferent between accepting or rejecting $\Psi(\beta^*)$ at $z = \beta^*$, which is necessary and sufficient for $w(\beta^*)$ to be randomized. The only states where the proposed strategies must be verified are those z off the grid and above α^*. It is sufficient to verify that

   - $F_L^\Delta(z) \geq V_L$ for all $z \in (\alpha^*, \beta^*)$
   - $F_H^\Delta(z) \geq \Psi(z)$ for all $z \in (\alpha^*, \beta^*)$
   - The high type does not wish to deviate by rejecting when $z > \beta^*$

   All three of these follow arguments based on the structure of the Bayesian updating and the "marginal benefit" of delay in the seller's problem (§ 2.4), which are completely analogous to the ones given for the verification of the equilibrium of the continuous-time game.

2. The only difference from the first case is that at $B_H^\Delta(\alpha^*)$ the high type *strictly* prefers to accept, which will imply $\beta^* < B_H^\Delta(\alpha^*)$. In this case, $w(B_H^\Delta(\alpha^*)) = \Psi(B_H^\Delta(\alpha^*))$ with probability one. $\alpha^* = \underline{\alpha}_L$ implies that this makes $F_L^\Delta(\alpha^*) = V_L$ creating the necessary indifference between acceptance and continuation for $z < \alpha^*$. The high type strictly prefers to accept at $z = B_H^\Delta(\alpha^*)$, but he strictly prefers to reject $\Psi((B_H^\Delta(\alpha^*))_{-1})$ at $z = (B_H^\Delta(\alpha^*))_{-1}$. It follows from the structure of the high type's seller's problem that there exists a unique $\beta^* \in ((B_\theta^\Delta(\alpha^*))_{-1}, B_\theta^\Delta(\alpha^*))$ where the high type will reject $\Psi(z)$ for $z < \beta^*$ and accept $\Psi(z)$ for $z > \beta^*$. At β^*, the high type is indifferent between accepting and rejecting $\Psi(\beta^*)$. Because $\beta^* \notin G_{\alpha^*}$ the behavior at β^* does not affect the low type's value at α^*, and therefore, γ can be any element of $[0, 1]$.

3. The case of "near intersection" is very much like the second case. The high type strictly prefers to accept $\Psi(z)$ at $z > B_H^\Delta(\alpha^*) = \beta^*$, but he is indifferent at $z = B_H^\Delta(\alpha^*)$. $B_L^\Delta(\alpha^*) = (B_H^\Delta(\alpha^*))_{+1}$ implies that $B_H^\Delta(\alpha^*)$ is an element of G_{α^*}. Therefore, the probability with which $\Psi(\beta^*)$ is offered at β^* does affect $F_L^\Delta(\alpha^*)$. To maintain $F_L^\Delta(\alpha^*) = V_L$, γ must be zero.

B.5 Equilibrium Uniqueness

Here we provide an argument that the equilibrium in Theorem B.1 is unique. For the sake of brevity, we will assume that any equilibrium must satisfy F_L^Δ and F_H^Δ non-decreasing.[44]

Consider a grid G defined by an arbitrary z. The high type will never trade when beliefs are $z < \underline{z} = \frac{e^{\underline{p}}}{1+e^{\underline{p}}}$. Therefore, just as in the continuous-time game, there exists a $z \in G$ such that the low type trades at a price of V_L with positive probability when beliefs are z. Further, in discrete time, when the market belief is high, delay carries a discrete cost for both types. Hence, there exists a z high enough such that both types trade with probability one for all beliefs above z.

Now, let $\widetilde{z}$ be the greatest element of G such that the low type trades at a price of V_L with positive probability when beliefs are $\widetilde{z}$. First, in equilibrium, if the high type rejects at $\widetilde{z}$, he will receive some continuation payoff greater than K_H. Second, because $\widetilde{z}$ is the maximum element of G where the low type accepts offers that the high type does not, rejection at any $z \in G$, $z > \widetilde{z}$, will not affect Z_t. Beliefs will be governed only by news when $Z_t > \widetilde{z}$ on G.

Therefore the minimum $z \in G$, $z \geq \widetilde{z}$, where the high type will be weakly willing to accept $\Psi(z)$ will coincide to the minimum z where the high type is weakly willing to accept $\Psi(z)$ in the Ψ-game posed in § 2.4 modified to the discrete-time setup and with the value of continuation at $\widetilde{z}$ set at its equilibrium continuation value. Just as in Lemma 2.1, there will exist a unique cutoff z_H such that the high type strictly prefers to accept (and trade occurs with probability one) when beliefs are greater than z_H, and strictly prefers to reject (and there will be no trade) when beliefs are below z_H. In the discrete setup, he may or may not be

[44] A full proof, following much the same argument given in the proof of Theorem 3.1, illustrates that this assumption is not necessary.

indifferent at the cutoff. For small Δ, z_H must be bounded away from $\widetilde{z}$ otherwise the low type would prefer to reject V_L at $\widetilde{z}$.

Index each $\widetilde{z}$ by the grid G it resides as $\widetilde{z}_G$ and let $\mathcal{G}$ be the set of all grids. Let $\alpha_L \equiv \sup_{\mathcal{G}} \widetilde{z}_G$. Given that only the low trades with positive probability for beliefs below α_L, if the market belief is initially $z_0 < \alpha_L$, it will jump to some level $j(z_0) \geq \alpha_L$. For small Δ, $j(z_0)$ must be bounded away (by multiple grid points) from the beliefs where the both types accept. To see this, let the minimum $z \in G_{j(z_0)}$ such that both types trade with probability one be $(j(z_0))_{+n}$. Then

$$F_L^{\Delta}(j(z_0)) \geq (q^{\Delta})^n \left[k_L \int_0^{n\Delta} e^{-rs} ds + e^{-rn\Delta}\Psi(z_H)\right] + (1-(q^{\Delta})^n)\left[k_L \int_0^{\Delta} e^{-rs} ds + e^{-r\Delta} V_L\right]$$

The second term on the right hand side is the minimal continuation value for the low type since $F_L^{\Delta}(z) \geq V_L$ for all z. For Δ small, $n \gg 1$ is necessary for the left hand side term to be no greater than V_L, which is needed to maintain the low type's indifference.

Let $\alpha_H \equiv sup_{z_0<\alpha_L}\{j(z_0)\}$. Because $j(z_0) \geq \alpha_L$, $F_L(j(z_0))$ is driven by continuation value, which must be V_L in order to satisfy the low type's indifference condition at z_0. Because $F_L(j(z_0)) = V_L$ for all z_0 and F_L non-deceasing, $F_L(z) = V_L$ for all $z < \alpha_H$.

Consider any two states z and z' such that the low type trades at price V_L with positive probability at both z and z'. Our next step is to show that $j(z) = j(z')$. Let $\mathcal{J}$ be the set of states where the low type does not accept V_L, but $F_L(j) = V_L$ for all $j \in \mathcal{J}$. To establish that $\mathcal{J}$ contains a single element, suppose it does not. For clarity, assume that B_L^{Δ} and B_H^{Δ} intersect.[45] Recall that B_L^{Δ} and B_H^{Δ}, therefore, coincide on a half-open interval $[\underline{\alpha}, \overline{\alpha})$.

Consider $\underline{j} \equiv \min j \in \mathcal{J}$ (if the minimum does not exist, there is a straightforward extension of the argument considering an element of $\mathcal{J}$ arbitrarily close to the $\inf_{\mathcal{J}}$). We claim that $\underline{j} \geq \underline{\alpha}$. Suppose that it were not so. The low type's value at $\underline{j}$ is V_L, where $\underline{j}$ is the lowest point of a no-trade region on grid $G_{\underline{j}}$. Hence, there is a unique profile of offers consistent with $F_L(\underline{j}) = V_L$:

- No trade at any $z < B_L^{\Delta}(\underline{j})$.
- At $B_L^{\Delta}(\underline{j})$, $\Psi(B_L^{\Delta}(\underline{j}))$ is offered with a uniquely determined $\gamma_{\underline{j}} \in (0, 1]$.
- At $z > B_L^{\Delta}(\underline{j})$, $\Psi(z)$ is offered with probability one.

However, from our analysis of the B_θ^{Δ} curves, we know this cannot hold in equilibrium. Because $\underline{j} < \underline{\alpha}$, $B_L^{\Delta}(\underline{j}) < B_H^{\Delta}(\underline{j})$. Therefore, if "reflection" occurred as in the definition of B_H^{Δ}, the high type prefers to reject $\Psi(B_L^{\Delta}(\underline{j}))$ in favor of continuation when $Z_t = B_L^{\Delta}(\underline{j})$. The reflection, however, may not take place as in the definition of B_H^{Δ}. But, because $\underline{j}$ is the minimum $j \in \mathcal{J}$, the reflection process must be at least as favorable to the high type as the one in the definition of B_H^{Δ} (if Z_t falls below $\underline{j}$, it will jump to some $j' \geq \underline{j}$). He is then at least as willing to reject when $Z_t = B_L^{\Delta}(\underline{j})$. If the high type will not accept $\Psi(B_L^{\Delta}(\underline{j}))$ if it is offered, it will not be offered. This contradiction implies that $\underline{j} \geq \underline{\alpha}$. The analogous argument demonstrates that every element of $\mathcal{J}$ is less than $\overline{\alpha}$.

We now know that every element of $\mathcal{J}$ is an element of $[\underline{\alpha}, \overline{\alpha})$. The final step is to establish that $\mathcal{J} = \{\underline{\alpha}\}$. Consider the maximum $\overline{j} \equiv \max j \in \mathcal{J}$.[46] If $\overline{j} > \underline{\alpha}$, it implies two things: (1) that $\gamma_{\overline{j}} < 1$ to ensure low-type indifference, and (2) that, if reflection took place as in the definition of B_H^{Δ}, the high type would strictly prefer to accept at $B_L^{\Delta}(\overline{j})$. Again, $\overline{j}$ being the maximum element of $\mathcal{J}$ implies that the reflection process is no more favorable to the high type than the reflection process in the definition of B_H^{Δ}. Hence, the high type strictly prefers to accept at $B_L^{\Delta}(\overline{j})$. It is immediate that these are inconsistent—if the high type prefers to accept Ψ than to continue, Ψ must be offered with probability one. Hence, $\mathcal{J} = \{\underline{\alpha}\}$. The rest of the details establishing uniqueness of equilibrium follow from the subsection above on existence and verification.

[45] The argument for the case when the curves only "nearly intersect" is analogous, though slightly more nuanced.

[46] Again, if the maximum does not exist, the argument is easily extended by considering an element arbitrarily close to the sup of $\mathcal{J}$.

B.6 Equilibrium Convergence

Both statements in Theorem B.2 are simply corollaries to the equilibrium analysis we have conducted and Lemma B.1. Recall that in the continuous-time game, the equilibrium is characterized by the unique pair (α^*, β^*) such that $B_H(\alpha^*) = B_L(\alpha^*) = \beta^*$. The discrete-time game is similarly characterized by a pair $(\alpha^{\Delta*}, \beta^{\Delta*})$ where B_H^{Δ} and B_L^{Δ} interest (or "nearly intersect"). It is immediate that $(\alpha^{\Delta*}, \beta^{\Delta*}) \to (\alpha^*, \beta^*)$ as $\Delta \to 0$. The theorem follows.

Proof of Lemma B.1. We begin with a fact that follows easily from the convergence of the discrete-time news processes to their continuous-time counterparts. Let $F_\theta(\cdot|\alpha, \beta)$ be defined in continuous time analogously to how $F_\theta^{\Delta}(\cdot|\alpha, \beta)$ was defined in discrete time.

Fact B.3. *Let $\{\beta^{\Delta}\}$ be a sequence that converges to a limit $\widehat{\beta}$ as $\Delta \to 0$. Then, as $\Delta \to 0$*

- $F_L^{\Delta}(z|\alpha, \beta^{\Delta}) \to F_L(z|\alpha, \widehat{\beta})$ *pointwise.*
- $F_H^{\Delta}(z|\alpha, \beta^{\Delta}) \to F_H(z|\alpha, \widehat{\beta})$ *pointwise.*

In both discrete and continuous time, the value functions are determined by the values at absorbing states, the likelihood of reaching each absorbing state, and the distribution on first hitting times of the absorbing states. The discrete-time versions of each of these components are converging to their continuous-time analogs, giving the result.

Now we show that $B_L^{\Delta}(\alpha) \to B_L(\alpha)$ for any fixed α. Recall that $B_L(\alpha)$ is the *unique* β such that, given α, three conditions hold:

1. $F_L(\alpha|\alpha, \beta) = V_L$
2. $F_L(\beta|\alpha, \beta) = \Psi(\beta)$
3. $F_L'(\alpha|\alpha, \beta) = 0$

Fix an α. Let $\widehat{\beta}_L^{\Delta}$ be the limit of $B_L^{\Delta}(\alpha)$ as $\Delta \to 0$. It is sufficient to show that

1. $F_L(\alpha|\alpha, \widehat{\beta}_L^{\Delta}) = V_L$
2. $F_L(\widehat{\beta}_L^{\Delta}|\alpha, \widehat{\beta}_L^{\Delta}) = \Psi(\widehat{\beta}_L^{\Delta})$
3. $F_L'(\alpha|\alpha, \widehat{\beta}_L^{\Delta}) = 0$

The first two follow easily from the definition of $F_L^{\Delta}(\cdot|\alpha, \beta)$. $F_L^{\Delta}(B_L^{\Delta}(\alpha)|\alpha, B_L^{\Delta}(\alpha)) = \Psi(B_L^{\Delta}(\alpha))$ for all Δ, giving the second point. For the first, the continuity of $F_L(\cdot|\alpha, \beta)$ implies that

$$\left|F_L^{\Delta}(\alpha|\alpha, B_L^{\Delta}(\alpha)) - F_L^{\Delta}(\alpha_{-1}|\alpha, B_L^{\Delta}(\alpha))\right| \to 0$$

$F_L^{\Delta}(\alpha_{-1}|\alpha, B_L^{\Delta}(\alpha)) = V_L$ for all Δ gives the result.

We now verify the third point. We need that

$$\lim_{\Delta \to 0} \frac{F_L^{\Delta}(\alpha|\alpha, \beta_L^{\Delta}(\alpha)) - F_L^{\Delta}(\alpha_{-1}|\alpha, \beta_L^{\Delta}(\alpha))}{\alpha - \alpha_{-1}} = 0$$

It is routine to show that $(\alpha - \alpha_{-1})$ converges to zero at the same rate as $\sqrt{\Delta}$ does. By definition of B_L^{Δ}, $F_L^{\Delta}(\alpha|\alpha, B_L^{\Delta}(\alpha)) > F_L^{\Delta}(\alpha_{-1}|\alpha, B_L^{\Delta}(\alpha)) = V_L$ for all Δ. Therefore, $F_L'(\alpha|\alpha, \widehat{\beta}_L^{\Delta}) \geq 0$. To complete the verification, suppose that $F_L'(\alpha|\alpha, \widehat{\beta}_L^{\Delta}) > 0$. This implies that as $\Delta \to 0$, $F_L^{\Delta}(\alpha|\alpha, \beta_L^{\Delta}(\alpha)) - V_L$ is of an order at least as great as $\sqrt{\Delta}$. However, this contradicts the definition of $B_L^{\Delta}(\alpha)$ being the *maximum* $\beta \in G_\alpha$ that maintains $F_L^{\Delta}(\alpha|\alpha, \beta) \geq V_L$. To see this, consider an increase in β by one grid point from $B_L^{\Delta}(\alpha)$ to $B_L^{\Delta}(\alpha)_{+1}$ when Δ is small. A lower bound on the new lower value of $F_L^{\Delta}(\alpha|\alpha, B_L^{\Delta}(\alpha)_{+1})$ is

$$F_L^{\Delta}(\alpha|\alpha, B_L^{\Delta}(\alpha))(1 - r\Delta) - k_L\Delta = F_L^{\Delta}(\alpha|\alpha, B_L^{\Delta}(\alpha)) - \Delta(rF_L^{\Delta}(\alpha|\alpha, B_L^{\Delta}(\alpha)) - k_L)$$

This lower bound corresponds to simply increasing the length of each path by one period of time, with no increase in payoffs from reaching the new, higher, upper boundary and only writing terms that are first order in Δ. Because $\Delta(rF_L^\Delta(\alpha|\alpha, B_L^\Delta(\alpha)) - k_L)$ is second order compared to $\sqrt{\Delta}$, the decrease in value from increasing β is not large enough for $F_L^\Delta(\alpha|\alpha, B_L^\Delta(\alpha)_{+1}) < V_L$. Therefore, the original $B_L^\Delta(\alpha)$ did not stratify the definition of B_L^Δ. This establishes that $F_L'(\alpha|\alpha, \widehat{\beta}_L^\Delta) = 0$, and $B_L^\Delta(\alpha) \to B_L(\alpha)$ for any fixed α.

Now we show that $B_H^\Delta(\alpha) \to B_H(\alpha)$ for any fixed α. Recall that $B_H(\alpha)$ is the *unique* β such that, given α, three conditions hold:

1. $F_H(\beta|\alpha, \beta) = \Psi(\beta)$
2. $F_H'(\alpha|\alpha, \beta) = 0$
3. $F_H'(\beta|\alpha, \beta) = \Psi'(\beta)$

Fix an α. Let $\widehat{\beta}_H^\Delta$ be the limit of $B_H^\Delta(\alpha)$ as $\Delta \to 0$. It is sufficient to show that

1. $F_H(\widehat{\beta}_H^\Delta|\alpha, \widehat{\beta}_H^\Delta) = \Psi(\widehat{\beta}_H^\Delta)$
2. $F_H'(\alpha|\alpha, \widehat{\beta}_H^\Delta) = 0$
3. $F_H'(\widehat{\beta}_H^\Delta|\alpha, \widehat{\beta}_H^\Delta) = \Psi'(\widehat{\beta}_H^\Delta)$

Again, the first point is immediate. To see the second point, from the definition of value function, the nature of the reflection of Z at α and only writing terms that are first order in Δ,

$$F_H^\Delta(\alpha|\alpha, B_H^\Delta(\alpha)) = k_H\Delta + (1 - r\Delta)[q^\Delta F_H^\Delta(\alpha_{+1}|\alpha, B_H^\Delta(\alpha)) + (1 - q^\Delta)F_H^\Delta(\alpha|\alpha, B_H^\Delta(\alpha))]$$

Subtract $F_H^\Delta(\alpha)$ from both sides, and suppress the dependence on $(\alpha, B_H^\Delta(\alpha))$,

$$0 = k_H\Delta + q^\Delta[F_H^\Delta(\alpha_{+1}) - F_H^\Delta(\alpha)] - r\Delta[q^\Delta F_H^\Delta(\alpha_{+1}) + (1 - q^\Delta)F_H^\Delta(\alpha)]$$

Divide by $(\alpha_{+1} - \alpha)$ and take limits as $(\alpha_{+1} - \alpha) \to 0$. Again, $(\alpha_{+1} - \alpha)$ is of the order $\sqrt{\Delta}$. We are left with $F_H'(\alpha|\alpha, \widehat{\beta}_H^\Delta) = 0$.

To verify the third point, we start by bounding $F_H^\Delta(B_H^\Delta(\alpha)|\alpha, B_H^\Delta(\alpha))$, which we shorten to $F_H^\Delta(B_H^\Delta)$ hereafter. From the definition, $F_H^\Delta(B_H^\Delta) \leq \Psi(B_H^\Delta)$. Also, $F_H^\Delta((B_H^\Delta)_{-1}) > \Psi((B_H^\Delta)_{-1})$, and therefore,

$$F_H^\Delta(B_H^\Delta) > k_H\Delta + (1 - r\Delta)[q^\Delta\Psi((B_H^\Delta)_{+1}) + (1 - q^\Delta)\Psi((B_H^\Delta)_{-1})]$$

Setting $F_H^\Delta(B_H^\Delta)$ to its upper bound of $\Psi(B_H^\Delta)$ for all Δ and taking the limit gives that $F_H'(\widehat{\beta}_H^\Delta|\alpha, \widehat{\beta}_H^\Delta) = \Psi'(\widehat{\beta}_H^\Delta)$. Now set $F_H^\Delta(B_H^\Delta)$ to its lower bound for all Δ and evaluate

$$\frac{F_H^\Delta((B_H^\Delta)_{+1}) - F_H^\Delta(B_H^\Delta)}{(B_H^\Delta)_{+1} - B_H^\Delta} = \frac{\Psi((B_H^\Delta)_{+1}) - k_H\Delta + (1 - r\Delta)[q^\Delta\Psi((B_H^\Delta)_{+1}) + (1 - q^\Delta)\Psi((B_H^\Delta)_{-1})]}{(B_H^\Delta)_{+1} - B_H^\Delta}$$

Re-arranging and eliminating additive terms that tend to zero as $\Delta \to 0$,

$$\begin{aligned} \frac{F_H^\Delta((B_H^\Delta)_{+1}) - F_H^\Delta(B_H^\Delta)}{(B_H^\Delta)_{+1} - B_H^\Delta} &= (1 - q^\Delta)\frac{\Psi((B_H^\Delta)_{+1}) - \Psi(B_H^\Delta)_{-1}}{(B_H^\Delta)_{+1} - B_H^\Delta} \\ &= (1 - q^\Delta)2\frac{\Psi((B_H^\Delta)_{+1}) - \Psi((B_H^\Delta)_{-1})}{2((B_H^\Delta)_{+1} - B_H^\Delta)} \\ &= (1 - q^\Delta)2\frac{\Psi((B_H^\Delta)_{+1}) - \Psi((B_H^\Delta)_{-1})}{((B_H^\Delta)_{+1} - (B_H^\Delta)_{-1})} \end{aligned}$$

Taking the limit as $\Delta \to 0$, we get that $F_H'(\widehat{\beta}_H^\Delta|\alpha, \widehat{\beta}_H^\Delta) = \Psi'(\widehat{\beta}_H^\Delta)$. Therefore, and $B_H^\Delta(\alpha) \to B_H(\alpha)$ for any fixed α. ∎