Market-making with Search and Information Frictions*

Benjamin Lester
Federal Reserve Bank of Philadelphia

Ali Shourideh
Carnegie Mellon University

Venky Venkateswaran
NYU Stern School of Business

Ariel Zetlin-Jones
Carnegie Mellon University

March 23, 2018

Abstract

We develop a dynamic model of trading through market-makers that incorporates two canonical sources of illiquidity: trading (or search) frictions, which imply that market-makers have some amount of *market power*; and information frictions, which imply that market-makers face some degree of *adverse selection.* We use this model to study the effects of various technological innovations and regulatory initiatives that aim to reduce trading and/or information frictions. We show that conventional predictions, derived from models that study either of these frictions in isolation, do not necessarily hold when both frictions are present. A key result is that reducing trading frictions slows down the rate at which market-makers learn about asset quality, amplifying the effects of adverse selection and ultimately leading to more illiquidity, as measured by, e.g., bid-ask spreads.

Keywords: Adverse selection, trading frictions, bid-ask spreads, liquidity, learning.

*We thank Pierre-Olivier Weill, Burton Hollified, Dimitri Vayanos, Marzena Rostek, Giuseppe Moscarini, Edouard Schaal, Willie Fuchs, participants at the 2017 Money, Banking and Asset Markets Conference (Madison, WI), the Search and Matching Conference (Philadelphia), Cornell-PSU Macro Conference, Annual Meeting of the Society for Economic Dynamics as well as seminar audiences at CREI, European University Institute and Carnegie-Mellon for useful comments and discussions.

1 Introduction

In recent years, financial markets have undergone significant changes as a result of both technological innovations and regulatory initiatives. Many of these changes have targeted one of two fundamental sources of illiquidity in financial markets, which we succinctly summarize as either *trading* frictions or *information* frictions. By trading frictions, we mean any friction that prevents traders from having immediate access to a wide variety of trading partners. By information frictions, we mean a situation in which one side of a transaction has more information about the quality of the asset than the other, which creates a standard adverse selection problem.

Consider, for example, technological developments such as the personal computer, electronic trading platforms, and faster bandwidth. As a result of these innovations, investors now have the opportunity to trade much more quickly with a wider set of dealers.[1] In addition, consider the effects of recent regulatory initiatives that advocate a shift from opaque over-the-counter markets towards more transparent, standardized exchanges. As these initiatives are implemented, investors gain much greater access to information regarding prevailing prices and quantities traded. [2] As these changes continue—and, in some cases, accelerate—a natural question arises: how does market liquidity respond to changes in the level of trading and/or information frictions?

The goal of this paper is to provide some answers to this question. To do so, we develop a model that incorporates both trading frictions and information frictions into a single, unified framework. We characterize equilibrium prices, trading decisions, and the corresponding evolution of beliefs, and then perform comparative statics to understand how the underlying frictions in the model determine market liquidity. We focus much of our attention on one particular measure of market liquidity, the bid-ask spread, though we also discuss implications for other measures, such as trading volume and price impact.

The effect of trading frictions and information frictions on bid-ask spreads have been studied extensively in separate literatures, which provide fairly stark predictions about the consequences of reducing either friction in isolation. In particular, the literature that formalizes OTC markets using search and matching models of trade, such as Duffie et al. (2005), offers a simple answer to the question posed above:

[1]The transition of financial markets from dealer-based platforms to electronic platforms, and the ever-increasing execution speeds, have been well documented. Appendix A of Pagnotta and Philippon (2015) offers an excellent summary.

[2]In the U.S., for example, the Dodd-Frank Wall Street Reform and Consumer Protection Act has called for the introduction of Swap Execution Facilities in the market for interest rate swaps; according to this legislation, an investor's request to trade must be circulated to at least three dealers for price quotes before the trade can be executed. Similar regulatory requirements for pre-trade price transparency have been implemented in European markets as a consequence of MiFID II. Requirements for post-trade transparency are equally as prevalent. A notable example is legislation introduced by the Financial Industry Regulatory Authority (FINRA), which collects post-trade price information in markets for asset-backed securities and corporate bonds through the Trade Reporting and Compliance Engine (TRACE).

if investors can contact dealers more easily, then competition among dealers will increase and bid-ask spreads will fall. Likewise, the literature that rationalizes bid-ask spreads as a consequence of asymmetric information, such as Glosten and Milgrom (1985), offers an equally straightforward prediction: if dealers have access to better information about the payoffs of the asset being traded, adverse selection will be less severe and bid-ask spreads again should fall.

Our main result is that these conventional predictions are not necessarily true when both frictions are present. Instead, reducing one of these frictions can make the other more severe. We illustrate this result, primarily, by showing that reducing trading frictions can slow down the process by which dealers learn about the quality of the assets being traded, which exacerbates the effects of asymmetric information and can ultimately lead to *wider* bid-ask spreads.

To understand the intuition behind this result, it's helpful to describe a few key features of the model. There are two types of agents—traders and dealers—who trade a homogeneous asset that is either high or low quality. We introduce information frictions by assuming that traders know the quality of the asset but dealers do not. We introduce trading frictions by assuming that, in each period, traders are matched with a stochastic number of dealers. In particular, a given trader might not match with any dealers, in which case they can't trade; they might be matched with a single dealer, in which case the dealer has some degree of market power; or they might be matched with two or more dealers, in which case the dealers compete a la Bertrand. After a dealer is matched with a trader, and observes whether the trader has matched with any other dealers, he offers bid and ask prices at which he's willing to buy or sell, respectively. The trader then decides whether or not to trade based on her *reservation value* for the asset, along with the realization of contemporaneous (aggregate and idiosyncratic) preference shocks. The trader's reservation value can be decomposed into two pieces: one that depends on the fundamental value of the asset and another that depends on the expected gains from trading the asset in the future. Dealers observe aggregate trading volume, which depends on the true quality of the asset and the (uncorrelated, unobserved) aggregate preference shock. Hence, volume is a noisy signal of asset quality, and dealers update their beliefs accordingly.

The first key piece of intuition is that dealers learn quickly when investors' behavior—which is summarized by their reservation value—is very different in the two states of the world, whereas they learn more slowly when investors behave similarly regardless of asset quality. The second key result is that reducing trading frictions implies that investors' reservation values depend more heavily on the payoffs associated with trading the asset in the future, and less on the fundamental value of the asset. Having frequent opportunities to trade, with potentially many dealers, implies that traders can expect to take advantage of favorable prices when their preference shocks dictate a desire to buy or sell. Alternatively, when traders meet dealers infrequently, expected gains from trade shrink and reservation values depend more heavily on the fundamental value.[3] As a result, reducing trading frictions makes investors' reservation values more similar across asset qualities, so that the (endogenously generated) signals that dealers observe are less informative. The last piece of the puzzle relates the speed of learning to bid-ask spreads: when learning slows down, dealers set wider bid-ask spreads to compensate for being more uncertain about the quality of the asset. Hence, we conclude that reducing trading frictions can actually reduce market liquidity, as measured by the bid-ask spread.[4]

We think our analysis constitutes a contribution to the literature for several reasons. First, it provides a single, unified framework that incorporates three ingredients that have been identified as crucial factors in financial markets: trading frictions, which have been studied extensively in search-based models of OTC markets; adverse selection, which lies at the heart of information-based models of market microstructure and the bid-ask spread; and learning, which is the focal point of dynamic models of information revelation. Second, by identifying novel interactions between these three ingredients—and the resulting, counter-intuitive implications for observable outcomes—our model offers a framework to interpret the ambiguous effects of reducing either trading frictions or information frictions that have been documented in the existing empirical evidence. For example, given the tradeoff discussed above, our model can easily rationalize an increase in the bid-ask spread after a reduction in trading or information frictions, as reported in, e.g., Huang and Stoll (1997), Madhavan et al. (2005), Acharya and Johnson (2007), and Hendershott and Moulton (2011). Lastly, if our model can help us understand the affects of trading frictions and information frictions in the past, it may also prove for anticipating the effects of future changes to OTC market structures, as a number of current regulatory proposals are aimed at either eliminating trading frictions or reducing information asymmetries in financial markets.

The rest of the paper is organized as follows. After reviewing the related literature below, we introduce the model in Section 2, characterize optimal behavior, and define an equilibrium. In Section 3, we consider a special case of the model that admits an analytical solution, and use this special case to illustrate the key results. Then, in Section 4, we consider a more flexibile specification and assign parameter values that are roughly consistent with those in the existing literature and/or the data. This section is not intended as a serious calibration exercise, per se, but rather as a vehicle that serves two purposes. First, it should

[3] In the limit, when traders live in autarky, the fundamental value is the *only* component of the reservation value.

[4] In Section **??**, we also show that reducing the degree of asymmetric information can reduce market liquidity as well.

assure the reader that there is nothing particularly unique about the more tractable special case in Section 3. Second, it allows us to derive additional results that cannot be derived analytically. Section 5 concludes.

1.1 Related Literature

This paper is related to several strands of the literature. First, it is closely related to the large body of work that uses search frictions to model decentralized trading. Duffie et al. (2005), Lagos and Rocheteau (2009), and Hugonnier et al. (2014) focus on implications for bid-ask spreads under full information.[5] Gehrig (1993), Spulber (1996), and, more recently, Lester et al. (2015) analyze pricing under asymmetric information about preferences, i.e., about the traders' private values of holding the asset.[6] In our paper, the traders possess private information about their preferences *and* about a common value component of the asset, which leads to adverse selection. Moreover, since all assets have the same common component, there is a role for learning over time by the uninformed market-makers.[7]

This combination of adverse selection, learning, and decentralized trade in our model is also present in papers such as Wolinsky (1990), Blouin and Serrano (2001), Duffie and Manso (2007), Duffie et al. (2009), Golosov et al. (2014) and Lauermann and Wolinsky (2016). The key difference between these papers and our own is the source of learning: in these papers, agents learn only from their own trading experiences, while in our paper, learning occurs from observing market-wide outcomes, which we feel is a realistic feature of many financial markets.[8]

In analyzing the effects of reducing trading frictions, we also make contact with the literature that studies the effects of high frequency trading, such as Biais et al. (2015), Pagnotta and Philippon (2015), Menkveld and Zoican (2017), and Du and Zhu (2017). A key distinction between our work and these papers—in addition to the many different modeling assumptions—is the crucial role that is assigned to the dealers' learning process in our framework.

Finally, our analysis also contributes to several strands of a large literature that focuses on the effects of asymmetric information in settings without trading and/or search frictions. One strand of this

[5]In fact, Lagos and Rocheteau (2009) also find that bid-ask spreads can widen when trading frictions ease, though the mechanism is different. In particular, they document this property in an environment with no information frictions and investors who can hold arbitrary portfolios, whereas we establish our results in an environment with information frictions and investors who can only hold zero or one unit of the asset. See also Afonso (2011), who shows that reducing trading frictions can have counter-intuitive effects because of congestion externalities.

[6]Another, more recent, example is Bethune et al. (2016)

[7]This latter feature distinguishes our work from papers that study adverse selection stemming from private information about the idiosyncratic quality of an asset; a non-exhaustive list of papers in this tradition includes Camargo and Lester (2014), Guerrieri and Shimer (2014), Kaya and Kim (2015), Fuchs and Skrzypacz (2015), Chiu and Koeppl (2016), Choi (2016), and Kim (2017). In these papers, information revealed from a particular trade is asset-specific and therefore, is typically not useful in future trades.

[8]A related literature also studies learning and information diffusion in network settings; see, e.g., Babus and Kondor (2016).

literature focuses on the effects of asymmetric information on the bid-ask spread, such as the seminal contributions of Copeland and Galai (1983), Glosten and Milgrom (1985), and Kyle (1985). Our focus on the informational content of endogenous market signals is shared by the stand of this literature that studies information aggregation in rational expectations equilibrium (REE) models, pioneered by Grossman and Stiglitz (1980) and Hellwig (1980). In contrast to these earlier papers, however, our analysis highlights novel interactions between asymmetric information and search frictions, and shows how these interactions can lead to surprising and counter-intuitive implications for liquidity and prices.

2 Model

2.1 Environment

Agents, Assets, and Preferences. Time is discrete and indexed by t. There are two types of risk neutral, infinitely-lived agents that we call "traders" and "dealers." Neither traders nor dealers discount future payoffs. There is a single asset of quality $j \in \{l, h\}$. Traders can hold either zero or one unit of the asset, while dealers' positions are unrestricted, i.e., dealers can take on arbitrarily long or short positions.

At the beginning of each period, the asset matures with probability $1 - \delta$, in which case the game ends. A trader who owns a unit of the asset receives a payoff c_j if the asset matures, with $c_l < c_h$. If the asset does not mature in period t, a trader who owns a unit of the asset receives a flow payoff $\omega_t + \varepsilon_{i,t}$, which we interpret as a liquidity shock. The aggregate portion of the shock, ω_t, is an i.i.d. draw each period from a distribution $F(\cdot)$. The idiosyncratic portion of the liquidity shock, $\varepsilon_{i,t}$, is an i.i.d. draw for each trader in each period from a distribution $G(\cdot)$. We assume that $F(\cdot)$ and $G(\cdot)$ have full support and mean zero, and we normalize the payoffs to a trader without an asset to zero.[9]

Finally, we assume that there is a large mass of dealers. Dealers receive a payoff ν_j when the asset matures, with $\nu_h > \nu_l$, but they do not receive any flow payoff from the asset before it matures. Given our assumption that dealers can take unrestricted positions, it follows that the payoff to a dealer from buying or selling a unit of the asset of quality $j \in \{l, h\}$ is ν_j and $-\nu_j$, respectively.

Trading and Frictions. There are two key frictions in the model. The first is an information friction: we assume that traders know more about the quality of the asset than dealers. For the sake of simplicity, we make the extreme assumption that all traders are endowed with perfect information about the quality of

[9] There are several alternative interpretations of what it means for an asset to "mature" in period t. For example, one interpretation is that a trader stops actively trading the asset in period t, i.e., he stops checking current bid and ask prices, and simply retains his current position (owner or non-owner) until the asset actually matures (or uncertainty about the asset's payoffs are resolved) at some future date $t' > t$. We discuss this at greater length in Section **??**.

the asset, $j \in \{l, h\}$, while dealers only know the *ex-ante* probability that the asset is of quality h at $t = 0$, which we denote by μ_0.

The second key friction in the model is a trading friction: in every period, each trader meets with a stochastic number of dealers. In particular, let p_i denote the probability that a trader meets $i \in \{0, 1, ...\}$ dealers. As we describe below, any meeting with $i \geqslant 2$ dealers will have the same outcome. Hence, the trading frictions can be succinctly summarized by two statistics, say p_0 and p_1. However, for the purpose of our analysis, it will be convenient to summarize the trading frictions instead by the probability that a trader meets with at least one dealer,

$$\pi \equiv 1 - p_0,$$

and the probability of meeting with one dealer (a "monopolist" meeting), conditional on meeting with $i \geqslant 1$ dealers,

$$\alpha_m = \frac{p_1}{\pi}.$$

We define the probability of meeting with $i \geqslant 2$ dealers (a "competitive" meeting), conditional on meeting with at least one dealer, by $\alpha_c = 1 - \alpha_m$.[10]

After meetings occur, each dealer quotes a bid and an ask price, i.e. prices at which she's willing to buy and sell a unit of the asset, respectively. Importantly, we assume that the number of dealers that a trader meets is common knowledge when dealers choose prices.[11] Hence, the dealer in a monopolist meeting can extract rents, whereas the dealers in a competitive meeting drive the bid price up and the ask price down until expected profits from the trade are zero. We denote the prices quoted by the dealer when she is a monopolist by (B_t^m, A_t^m), and the prices quoted by competing dealers by (B_t^c, A_t^c).

Information and Learning. We assume that dealers observe the aggregate volume of trade at the end of each trading round. As we describe in detail below, this will turn out to be a noisy signal about asset quality, which the dealers will use to update their beliefs over time. This assumption will play a crucial role in making our analysis tractable. In particular, as we will show, it implies that (i) all dealers have identical beliefs at the beginning of each period and (ii) the actions of an individual trader and/or dealer will not alter the evolution of future beliefs. In what follows, we let μ_t denote the beliefs of (all) dealers at the beginning of trading at time t that the asset is of quality h.

[10]This transformation is convenient in that it allows us separately study the effects of more frequent meetings (say, from an increase in trading speed) and the effects of more competition in each meeting (say, from an increase in pre-trade price transparency).

[11]This is in contrast to the literature on price dispersion that follows, e.g., Burdett and Judd (1983).

2.2 Traders' Optimal Behavior

Let $W^q_{j,t}$ denote the expected discounted value of a trader who owns $q \in \{0,1\}$ unit of the asset at the beginning of period t when the asset is of quality $j \in \{l,h\}$. Then, for an investor who does not own the asset, we have

$$W^0_{j,t} = \delta\, \mathbb{E}_{\omega,\varepsilon} \left[\pi \sum_{k=c,m} \alpha_k \max\left\{-A^k_t + \omega_t + \varepsilon_{i,t} + W^1_{j,t+1}, W^0_{j,t+1}\right\} + (1-\pi) W^0_{j,t+1} \right]. \tag{1}$$

Note that the expectation is taken over ω_t and $\varepsilon_{i,t}$, which are drawn from $F(\omega)$ and $G(\varepsilon)$, respectively. All objects inside the brackets—including the current ask prices A^k_t and future payoffs $W^q_{j,t+1}$—can be calculated using the information available to a trader at time t, which would include the true quality of the asset, along with the current beliefs of dealers. We describe in detail below how the trader uses this information to formulate beliefs.

In words, the first expression in equation (1) represents the expected payoff if the asset does not mature and the trader meets at least one dealer, whereupon he may either purchase a unit of the asset at price A^k_t or reject the offer and continue searching in period $t+1$. The second expression represents the expected payoff if the asset does not mature but the trader fails to meet a dealer. Recall that a trader with $q = 0$ assets receives zero payoff if the asset matures, which occurs with probability $1-\delta$.

Similar logic can be used to derive the expected payoff of a trader who owns one unit of the asset,

$$\begin{aligned} W^1_{j,t} &= (1-\delta)c_j + \\ &\delta \mathbb{E}_{\omega,\varepsilon}\left[\pi \sum_{k=c,m} \alpha_k \max\left\{\omega_t + \varepsilon_{i,t} + W^1_{j,t+1}, B^k_t + W^0_{j,t+1}\right\} + (1-\pi)\left(\omega_t + \varepsilon_{i,t} + W^1_{j,t+1}\right)\right]. \end{aligned} \tag{2}$$

Note that, when the asset matures, a trader who owns one unit receives a payoff c_j.

We conjecture, and later confirm, that an individual trader's decision to accept or reject an offer has no effect on dealers' beliefs, and hence no effect on the path of future prices. An immediate consequence is that traders' decisions to buy or sell follow simple cutoff rules: given asset quality $j \in \{l,h\}$, a trader who does not own the asset will buy in a meeting of type $k \in \{m,c\}$ if $\varepsilon_{i,t} \geqslant \bar{\varepsilon}^k_{j,t}$, while a trader who owns the asset will sell if $\varepsilon_{i,t} \leqslant \underline{\varepsilon}^k_{j,t}$, where these cutoffs satisfy

$$\begin{aligned} -A^k_t + \omega_t + \bar{\varepsilon}^k_{j,t} + W^1_{j,t+1} &= W^0_{j,t+1} \\ \omega_t + \underline{\varepsilon}^k_{j,t} + W^1_{j,t+1} &= B^k_t + W^0_{j,t+1}. \end{aligned}$$

Let us denote the *reservation value* of an investor at time t given asset quality $j \in \{l,h\}$ by

$$R_{j,t} \equiv W^1_{j,t} - W^0_{j,t}.$$

Then, the optimal behavior of traders is succinctly summarized by the cutoffs

$$\underline{\varepsilon}^k_{j,t} = B^k_t - \omega_t - R_{j,t+1} \tag{3}$$

$$\overline{\varepsilon}^k_{j,t} = A^k_t - \omega_t - R_{j,t+1} \tag{4}$$

defined for $k \in \{m, c\}$, along with the reservation value

$$R_{j,t} = (1-\delta)c_j + \tag{5}$$
$$\delta \mathbb{E}_\omega \left\{ \pi \sum_{k=c,m} \alpha_k \left[B_t G(\underline{\varepsilon}^k_{j,t}) + \int_{\underline{\varepsilon}^k_{j,t}}^{\overline{\varepsilon}^k_{j,t}} \left[\omega_t + \varepsilon_{i,t} + R_{j,t+1}\right] dG\left(\varepsilon_{i,t}\right) + A_t \left[1 - G(\overline{\varepsilon}^k_{j,t})\right] \right] + (1-\pi) R_{j,t+1} \right\},$$

which is obtained by subtracting (1) from (2) and using the cutoff rules described above. Again, the expectation operator in (5) is taken over the aggregate shock, ω_t, as well as the prices and future payoffs, which we describe below.

Demographics. Given the trading rules described above, we can now describe the evolution of the distribution of asset holdings across traders over time. To do so, let N^q_t denote the measure of traders who have asset holdings $q \in \{0, 1\}$ at time t. When the asset is of quality $j \in \{l, h\}$, then, we have

$$\begin{aligned} N^1_{j,t+1} &= N^1_t \left[1 - \pi + \pi \left(1 - \sum_{k=c,m} \alpha_k G(\underline{\varepsilon}^k_{j,t}) \right) \right] + N^0_t \pi \left[1 - \sum_{k=c,m} \alpha_k G(\overline{\varepsilon}^k_{j,t}) \right] \\ N^0_{j,t+1} &= N^1_t \pi \sum_{k=c,m} \alpha_k G(\underline{\varepsilon}^k_{j,t}) + N^0_t \left([1 - \pi + \pi \sum_{k=c,m} \alpha_k G(\overline{\varepsilon}^k_{j,t}) \right]. \end{aligned}$$

Naturally, the measure of investors that own an asset in period $t+1$ is equal to the measures of investors that owned an asset in period t and did not sell, plus the measure of investors that did not own an asset but chose to buy. The intuition behind the law of motion for the measure of investors that don't own an asset follows the same logic. We assume that the initial distribution of owners and non-owners, (N^0_0, N^1_0), is common knowledge. Hence, as we describe below, dealers will know (N^0_t, N^1_t) at the beginning of each period, but they will not be able to perfectly infer $j \in \{l, h\}$.

2.3 Dealers' Optimal Behavior

Monopolist Pricing. We first consider the optimal price offered in a meeting between a trader and a single dealer. When formulating this offer, the dealer takes as given the trader's optimal behavior derived above. We will show that, under our assumptions, the dealer's pricing problem is static: neither the price that she sets nor the trader's response affects payoffs in future periods (e.g., through beliefs). We treat this as a conjecture, for now, and verify it later.

Under this conjecture, a monopolist dealer's optimal prices (A_t^m, B_t^m) solve

$$\max_{A,B} \quad \mathbb{E}_{j,\omega}\left[N_t^0\left[1-G\left(\bar{\varepsilon}_{j,t}\right)\right](A-\nu_j)+N_t^1 G\left(\underline{\varepsilon}_{j,t}\right)(\nu_j-B)\right],$$

where the expectations operator is taken over the quality $j \in \{l, h\}$ of the asset—using the dealer's current beliefs, μ_t—as well as the aggregate liquidity shock, ω_t, and future reservation values, $R_{j,t+1}$, that determine the thresholds $\underline{\varepsilon}_{j,t}$ and $\bar{\varepsilon}_{j,t}$. Again, we postpone the derivation of how these latter expectations are formed until Section 2.4, below.

The optimal prices can be summarized by the two first-order conditions:

$$0 = \mathbb{E}_{j,\omega}\left[1-G\left(\bar{\varepsilon}_{j,t}^m\right)-g\left(\bar{\varepsilon}_{j,t}^m\right)(A_t^m-\nu_j)\right] \quad (6)$$

$$0 = \mathbb{E}_{j,\omega}\left[-G\left(\underline{\varepsilon}_{j,t}^m\right)+g\left(\underline{\varepsilon}_{j,t}^m\right)(\nu_j-B_t^m)\right]. \quad (7)$$

Re-arranging equation (6), we can write the optimal ask price as

$$A_t^m = \mu_t\nu_h+(1-\mu_t)\nu_l + \underbrace{\frac{1-\mathbb{E}_{j,\omega}\left[G\left(\bar{\varepsilon}_{j,t}^m\right)\right]}{\mathbb{E}_{j,\omega}\left[g\left(\bar{\varepsilon}_{j,t}^m\right)\right]}}_{\text{market power}} + \underbrace{\mu_t(1-\mu_t)(\nu_h-\nu_l)\frac{\mathbb{E}_{\omega}\left[g\left(\bar{\varepsilon}_{h,t}^m\right)-g\left(\bar{\varepsilon}_{l,t}^m\right)\right]}{\mathbb{E}_{j,\omega}\left[g\left(\bar{\varepsilon}_{j,t}^m\right)\right]}}_{\text{asymmetric information}}. \quad (8)$$

This equation expresses the ask price as the sum of the expected fundamental value of the asset and two additional components. The first component derives from the dealer's *market power* and is the inverse of the semi-elasticity of expected demand, akin to the standard markup in a monopolist's optimal price. The second component is a premium that dealers charge to compensate for the presence of *asymmetric information.* Note that this second component can be re-written as

$$\mu_t(1-\mu_t)(\nu_h-\nu_l)\frac{\mathbb{E}_{\omega}\left[g\left(\bar{\varepsilon}_{h,t}^m\right)-g\left(\bar{\varepsilon}_{l,t}^m\right)\right]}{\mathbb{E}_{j,\omega}\left[g\left(\bar{\varepsilon}_{j,t}^m\right)\right]} = \text{Cov}\left(\frac{g\left(\bar{\varepsilon}_{j,t}^m\right)}{\mathbb{E}_{j,\omega}\left[g\left(\bar{\varepsilon}_{j,t}^m\right)\right]}\ ,\ \nu_j\right).$$

Hence, the asymmetric information component is essentially an adjustment that accounts for the relationship between the density of marginal buyers and the dealer's valuation of the asset; it implies that dealers will adjust their asking price upward if the density of marginal buyers is relatively large when the asset is of high quality.[12] Also note that this component disappears if there is no uncertainty over the quality of the asset, i.e., if $\mu_t = 0$, $\mu_t = 1$, or $\nu_h = \nu_l$.

Similar logic reveals that the bid price is equal to the expected fundamental value of the asset, adjusted downwards by the two components discussed above:

$$B_t^m = \mu_t\nu_h+(1-\mu_t)\nu_l - \frac{\mathbb{E}_{j,\omega}\left[G\left(\underline{\varepsilon}_{j,t}^m\right)\right]}{\mathbb{E}_{j,\omega}\left[g\left(\underline{\varepsilon}_{j,t}^m\right)\right]} - \mu_t(1-\mu_t)(\nu_h-\nu_l)\frac{\mathbb{E}_{\omega}\left[g\left(\underline{\varepsilon}_{l,t}^m\right)-g\left(\underline{\varepsilon}_{h,t}^m\right)\right]}{\mathbb{E}_{j,\omega}\left[g\left(\underline{\varepsilon}_{j,t}^m\right)\right])}. \quad (9)$$

[12]Note that this asymmetric information component can be positive or negative.

Competitive Pricing. Next, we solve for equilibrium prices when a trader meets two or more dealers. This situation corresponds almost exactly to the pricing problem in the canonical setting of Glosten and Milgrom (1985), where equilibrium bid and ask prices are set so that expected (static) profits are zero. In other words, when two or more dealers compete, the bid price B_t^c (ask price A_t^c) is equal to the expected value of the asset conditional on a trader selling (buying) at that price. Formally, this zero profit condition can be written

$$0 = A_t^c - \frac{\mathbb{E}_{j,\omega}\left[\nu_j\left(1 - G(\bar{\varepsilon}_{j,t}^c)\right)\right]}{\mathbb{E}_{j,\omega}\left[\left(1 - G(\bar{\varepsilon}_{j,t}^c)\right)\right]} \tag{10}$$

$$0 = B_t^c - \frac{\mathbb{E}_{j,\omega}\left[\nu_j G(\underline{\varepsilon}_{j,t}^c)\right]}{\mathbb{E}_{j,\omega}\left[G(\underline{\varepsilon}_{j,t}^c)\right]}. \tag{11}$$

Re-arranging yields

$$A_t^c = \mu_t\nu_h + (1-\mu_t)\nu_l + \mu_t(1-\mu_t)(\nu_h - \nu_l)\frac{\mathbb{E}_\omega\left[G\left(\bar{\varepsilon}_{l,t}^c\right) - G\left(\bar{\varepsilon}_{h,t}^c\right)\right]}{\mathbb{E}_{j,\omega}\left[1 - G\left(\bar{\varepsilon}_{j,t}^c\right)\right]} \tag{12}$$

$$B_t^c = \mu_t\nu_h + (1-\mu_t)\nu_l - \mu_t(1-\mu_t)(\nu_h - \nu_l)\frac{\mathbb{E}_\omega\left[G\left(\underline{\varepsilon}_{l,t}^c\right) - G\left(\underline{\varepsilon}_{h,t}^c\right)\right]}{\mathbb{E}_{j,\omega}\left[G\left(\underline{\varepsilon}_{j,t}^c\right)\right]}. \tag{13}$$

Again, it is worth noting that, e.g.,

$$\mu_t(1-\mu_t)(\nu_h - \nu_l)\frac{\mathbb{E}_\omega\left[G\left(\bar{\varepsilon}_{l,t}^c\right) - G\left(\bar{\varepsilon}_{h,t}^c\right)\right]}{\mathbb{E}_{j,\omega}\left[1 - G\left(\bar{\varepsilon}_{j,t}^c\right)\right]} = \mathrm{Cov}\left(\frac{1 - G\left(\bar{\varepsilon}_{j,t}^c\right)}{\mathbb{E}_{j,\omega}\left[1 - G\left(\bar{\varepsilon}_{j,t}^c\right)\right]}\ ,\ \nu_j\right).$$

These expressions show that, under competition, bid and ask prices are equal to the expected value of the asset to the dealer, adjusted for adverse selection. This adjustment depends on the covariance between the probability of trade and the value of the asset. For example, the ask (bid) price is higher (lower) than the expected value since traders are more (less) likely to buy (sell) when the state is high. This creates a positive bid-ask spread, exactly as in Glosten and Milgrom (1985).

Comparing equations (8)-(9) and (12)-(13) reveals that prices under competition are similar in structure to monopoly prices, with two key differences. First, as one might expect, competitive prices do not contain the markup component found in monopoly prices. Second, the adjustment for asymmetric information in (8)-(9) depends on the mass of traders at the appropriate thresholds in the two states (i.e., the pdf), while the corresponding adjustment in (12)-(13) depends on the difference between the probability of trade in the two states (i.e., the cdf). Intuitively, this occurs because the monopolist's optimal price is a function of the expected profit from the *marginal* trader, whereas competitive pricing is pinned down by the requirement that dealers earn zero profits on *average*.

2.4 Learning

We now explain how dealers update their beliefs about the quality of the asset, and how investors form expectations about dealers' beliefs—and hence the prices they offer—in future periods.

As noted above, we assume that dealers learn by observing aggregate trading activity in each period. Notice immediately that this is equivalent to observing the thresholds $\left(\bar{\varepsilon}^m_{j,t}, \bar{\varepsilon}^c_{j,t}, \underline{\varepsilon}^m_{j,t}, \underline{\varepsilon}^c_{j,t}\right)$. Moreover, since dealers know which prices have been offered in equilibrium, each of these thresholds ultimately contains the same information. Consider, for example, $\bar{\varepsilon}^m_{j,t}$, which depends on the ask price A^m_t, which dealers know, along with the reservation value of the trader $R_{j,t+1}$ and the aggregate shock ω_t, both of which the dealers do not observe. The reservation value clearly depends on the quality of the asset, while the aggregate liquidity shock is orthogonal to quality (by assumption). Hence, the volume of asset purchases in monopoly meetings—which dealers can perfectly infer from the total volume of asset purchases given α_c and α_m—is a noisy signal about asset quality, and the informational content can be summarized by

$$S_t = R_{t+1} + \omega_t \ , \tag{14}$$

where $R_{t+1} = R_{j,t+1}$ when the true state of the world is $j \in \{l, h\}$.

Let us conjecture, for now, that investors' reservation values depend only on dealers' beliefs, along with the true state j. Then, given current beliefs μ_t and the observed signal S_t, a dealer's updated belief μ_{t+1} depends on the likelihood of observing that signal when the asset's quality is h relative to that when it is l. To arrive at this likelihood, we first calculate the value of the aggregate shock, ω_t, that is consistent with the observed signal S_t. Formally, define

$$\omega^\star_{\iota,t} = S_t - R_{\iota,t+1}(\mu_{t+1}). \tag{15}$$

Using (14), $\omega^\star_{\iota,t}$ is the value of ω_t consistent with the signal S_t if the dealer conjectures that the asset is of quality $\iota \in \{l, h\}$ and future beliefs are μ_{t+1}. Naturally, if ι is equated to the true asset quality, then $\omega^\star_{\iota,t} = \omega_t$, i.e., the dealer's conjecture corresponds to the true value of the aggregate liquidity shock. Now, one might be concerned that the reservation values $R_{l,t+1}$ and $R_{h,t+1}$ in (15) are calculated under different information sets, but this is not the case since, *by construction*, both $\omega^\star_{l,t}$ and $\omega^\star_{h,t}$ are consistent with the signal S_t.

The law of motion for dealers' beliefs is thus a function $\mu_{t+1}(\mu_t, S_t)$ that solves the fixed point problem

$$\mu_{t+1} = \frac{\mu_t}{\mu_t + (1-\mu_t)\frac{f\left(\omega^\star_{l,t}\right)}{f\left(\omega^\star_{h,t}\right)}} = \frac{\mu_t}{\mu_t + (1-\mu_t)\frac{f(S_t - R_{l,t+1}(\mu_{t+1}))}{f(S_t - R_{h,t+1}(\mu_{t+1}))}}. \tag{16}$$

Now, even though dealers' future beliefs cannot depend directly on the true quality of the asset (since they do not observe it), traders (who know the true quality) can certainly use this information to formulate *expectations about dealers' beliefs*. In particular, it will be helpful to define the function $\tilde{\mu}_{j,t+1}(\mu_t, \omega_t)$ as the solution to the fixed point problem

$$\mu_{t+1} = \frac{\mu_t}{\mu_t + (1-\mu_t)\frac{f\left(\omega_t + R_{j,t+1}(\mu_{t+1}) - R_{l,t+1}(\mu_{t+1}))\right)}{f\left(\omega_t + R_{j,t+1}(\mu_{t+1}) - R_{h,t+1}(\mu_{t+1}))\right)}}. \tag{17}$$

In words, given current beliefs μ_t, the true quality of the asset is $j \in \{l, h\}$, and the aggregate liquidity shock ω_t , traders (correctly) anticipate that dealers' beliefs in period $t+1$ will be $\tilde{\mu}_{j,t+1}(\mu_t, \omega_t)$.

This recursive law of motion validates our earlier conjectures about the formation of beliefs in equilibrium. First, since future beliefs (and therefore, future prices) only depend on current beliefs and the realization of the aggregate liquidity shock ω_t, it follows that traders' reservation values R_{t+1} depend only on beliefs μ_{t+1} and the true quality of the asset.

Second, since future beliefs are independent of the actions of any one dealer or trader, both can formulate optimal behavior—prices for dealers and buy, sell, or don't trade for traders— without affecting future beliefs. This verifies the conjecture that the dealers' pricing problem is a static one. In other words, dealers do not have an incentive to deviate from the static optimal prices in order to experiment, i.e. to acquire information about the quality of the asset. To see why, note that an individual trader's action is measurable with respect to the sum of her reservation value R_{t+1} and the combined liquidity shock $\omega_t + \varepsilon_{i,t}$. At any quoted price, her action therefore can at best reveal $R_{t+1} + \omega_t + \varepsilon_{i,t}$. For example, if the dealer quotes a bid of B' and the trader chooses (not to) sell at that price, the dealer learns that $R_{t+1} + \omega_t + \varepsilon_{i,t}$ is (larger) smaller than B'. At the end of the period, the dealer perfectly learns $R_{t+1} + \omega_t$ by observing the market-wide volume[13]. The information contained in this signal about R_{t+1} (and therefore, about asset quality) dominates that contained in an individual trader's actions. Thus, deviating from the static optimal price involves giving up current profits but generates no additional benefit.

2.5 Definition of Equilibrium

We now define a Markov equilibrium, where the strategies of all agents are functions of (at most) current dealer beliefs, μ_t, and realizations of the aggregate liquidity shock, ω_t. Such an equilibrium can be represented recursively as a collection of functions $\left\{\underline{\varepsilon}_j^k, \overline{\varepsilon}_j^k, R_j, A^k, B^k, \mu^+, \tilde{\mu}_j^+, N_j^{0,+}, N_j^{1,+}\right\}$ for $j \in \{l, h\}$ and $k \in \{m, c\}$ such that:

[13]This follows from the assumption that liquidity shocks have full support. However, this is not essential. In Section 3, we consider a version with shocks drawn from a finite support and the no experimentation result still goes through.

1. Taking as given the way dealers set prices and update beliefs, investors' decisions to buy or sell are determined by:

$$\underline{\varepsilon}_j^k(\mu,\omega) = B^k(\mu) - \omega - R_j\left(\tilde{\mu}_j^+(\mu,\omega)\right) \quad (18)$$

$$\overline{\varepsilon}_j^k(\mu,\omega) = A^k(\mu) - \omega - R_j\left(\tilde{\mu}_j^+(\mu,\omega)\right) \quad (19)$$

$$\begin{aligned} R_j(\mu) &= (1-\delta)c_j + \delta(1-\pi)\int_\omega R_j\left(\tilde{\mu}_j^+(\mu,\omega)\right) dF(\omega) + \delta\pi\int_\omega \Bigg\{ \sum_{k\in\{m,c\}} \alpha^k \Big[B^k(\mu) G\left(\underline{\varepsilon}_j^k(\mu,\omega)\right) \\ &+ \int_{\underline{\varepsilon}_j^k(\mu,\omega)}^{\overline{\varepsilon}_j^k(\mu,\omega)} \left[\omega + \varepsilon + R_j\left(\tilde{\mu}_j^+(\mu,\omega)\right)\right] dG(\varepsilon) + A^k(\mu)\left[1 - G\left(\overline{\varepsilon}_j^k(\mu,\omega)\right)\right]\Big]\Bigg\} dF(\omega). \end{aligned} \quad (20)$$

2. Given investors' behavior and expectations about future beliefs, prices are consistent with optimal behavior and, in the competitive case, zero profits. That is, $A^m \equiv A^m(\mu)$ and $B^m \equiv B^m(\mu)$ satisfy:

$$0 = \sum_{j\in\{l,h\}} \mu_j \int_\omega \left[1 - G\left(\overline{\varepsilon}_j^m(\mu,\omega)\right) - g\left(\overline{\varepsilon}_j^m(\mu,\omega)\right)\left(A^m - \nu_j\right)\right] dF(\omega) \quad (21)$$

$$0 = \sum_{j\in\{l,h\}} \mu_j \int_\omega \left[-G\left(\underline{\varepsilon}_j^m(\mu,\omega)\right) + g\left(\underline{\varepsilon}_j^m(\mu,\omega)\right)\left(\nu_j - B^m\right)\right] dF(\omega) \quad (22)$$

where $\mu_h \equiv \mu$ and $\mu_l \equiv 1-\mu$, while $A^c \equiv A^c(\mu)$ and $B^c \equiv B^c(\mu)$ satisfy:

$$0 = A^c - \frac{\sum_{j\in\{l,h\}} \mu_j \nu_j \int \left[1 - G\left(\overline{\varepsilon}_j^c(\mu,\omega)\right)\right] dF(\omega)}{\sum_{j\in\{l,h\}} \mu_j \int \left[1 - G\left(\overline{\varepsilon}_j^c(\mu,\omega)\right)\right] dF(\omega)} \quad (23)$$

$$0 = B^c - \frac{\sum_{j\in\{l,h\}} \mu_j \nu_j \int \left[G\left(\underline{\varepsilon}_j^c(\mu,\omega)\right)\right] dF(\omega)}{\sum_{j\in\{l,h\}} \mu_j \int \left[G\left(\underline{\varepsilon}_j^c(\mu,\omega)\right)\right] dF(\omega)}. \quad (24)$$

3. Given a signal S, dealers' beliefs evolve according to $\mu^+(\mu, S)$, which is a solution to:

$$\mu^+ = \frac{\mu}{\mu + (1-\mu)\frac{f(S - R_l(\mu^+))}{f(S - R_h(\mu^+))}}. \quad (25)$$

Given the true asset quality $j \in \{l,h\}$ and aggregate shock ω, investors' expectations of dealers' beliefs evolve according to $\tilde{\mu}_j^+(\mu,\omega)$, which is a solution to

$$\mu^+ = \frac{\mu}{\mu + (1-\mu)\frac{f\left(\omega + R_j(\mu^+) - R_l(\mu^+)\right)}{f\left(\omega + R_j(\mu^+) - R_h(\mu^+)\right)}}. \quad (26)$$

Moreover, investors' expectations are consistent with the evolution of dealers' beliefs, so that

$$\tilde{\mu}_j^+(\mu,\omega) = \mu^+\left(\mu, R_j\left(\tilde{\mu}_j^+(\mu,\omega)\right) + \omega\right) \quad \text{for } j \in \{l,h\}. \quad (27)$$

4. Given true asset quality $j \in \{l,h\}$, beliefs μ, and an aggregate shock ω, the population evolves according to:

$$N_j^{1,+}(\mu,\omega) = N_j^1\left[1-\pi+\pi\left(1-\sum_{k\in\{m,c\}} G\left(\underline{\varepsilon}_j^k(\mu,\omega)\right)\right)\right] + N_j^0\pi\left(1-\sum_{k\in\{m,c\}} G\left(\overline{\varepsilon}_j^k(\mu,\omega)\right)\right) \quad (28)$$

$$N_j^{0,+}(\mu,\omega) = N_j^1\pi\sum_{k\in\{m,c\}} G\left(\underline{\varepsilon}_j^k(\mu,\omega)\right) + N_j^0\left[1-\pi+\pi\sum_{k\in\{m,c\}} G\left(\overline{\varepsilon}_j^k(\mu,\omega)\right)\right]. \quad (29)$$

Note that the laws of motion for N_j^1 and N_j^0 depend only on the thresholds $\left\{\underline{\varepsilon}_j^k, \overline{\varepsilon}_j^k\right\}$, for $j \in \{l, h\}$ and $k \in \{m, c\}$. Hence, dealers can always infer the distribution of assets across traders, even though they can't directly observe asset quality.[14]

3 Frictions, Learning, and Prices: A Tractable Case

In this section, we explore how trading friction and information frictions affect traders' reservation values, the evolution of dealers' beliefs, and, ultimately, equilibrium bid and ask prices. We show that, in isolation, each of these frictions has the expected effect: holding beliefs fixed, a reduction in trading frictions causes bid-ask spreads to narrow; and holding trading frictions constant, increasing uncertainty over the quality of the asset causes bid-ask spreads to widen.

However, the interaction between these two frictions generates novel predictions. First, we establish that reducing trading frictions slows down learning. Intuitively, when investors have the opportunity to trade more frequently, their behavior in the two states of the world is more similar, which implies that the endogenous signal in the model— aggregate volume—is less informative. Second, since slower learning implies more uncertainty, and more uncertainty implies wider spreads, we show that a reduction in trading frictions can ultimately lead to an increase in the bid-ask spread.

In order to establish these results analytically, we make a few parametric assumptions, which are described in detail below. These assumptions are not terribly special, per se, above and beyond the fact that they offer a certain amount of tractability (VV: THIS SENTENCE SEEMS UNNECESSARY - WE MIGHT GET INTO TROUBLE OVER "TERRIBLY SPECIAL". ALSO, MAYBE THE FOOTNOTE CAN BE MOVED INTO THE MAIN TEXT AT THE END OF THE PREV SECTION, AS A WAY OF MOTIVATING THE TRACTABLE CASE?).[15] In the next section, we establish that the key mechanisms derived here are preserved under more flexible specifications.

[14] Intuitively, by construction, $\omega_l^\star$ and $\omega_h^\star$ rationalize the aggregate trading volume that dealers observe, and hence the implied thresholds. As a result, the evolution of N^0 and N^1 when the asset quality is l and the aggregate shock is $\omega_l^\star$ are identical to the evolution of these variables when the asset quality is h and the aggregate shock is $\omega_h^\star$.

[15] As (18)–(27) reveal, one can see that there is a fairly complicated fixed point problem at the heart of the equilibrium: the law of motion for dealers' beliefs is a convolution of both exogenous parameters (ω) and endogenous variables (R_j), which themselves depend on future prices and beliefs. This makes it difficult to derive analytical results for arbitrary distributions of liquidity shocks.

3.1 Parametric Assumptions

We make three key assumptions, described below.

Assumption 1 (Uniform Shocks)**.** *The aggregate liquidity shock, ω, is uniformly distributed over the interval $[-m, m]$ for some $0 < m < \infty$, and the idiosyncratic liquidity shock, ε, is uniformly distributed over the interval $[-e, e]$ for some $0 < e < \infty$.*

As we will show below, the assumption that ω is uniformly distributed simplifies the dealers' learning process, while the assumption that ε is uniformly distributed simplifes the dealers' pricing problem. Note, however, that these distributions violate our maintained assumption that $F(\cdot)$ and $G(\cdot)$ have full support. One might be concerned that having finite bounds would open up the possibility that dealers would like to experiment when setting prices, e.g., that they would choose to set a (statically sub-optimal) price that would reveal to them the state of the world. We show in Appendix B.1 that this is not the case.

Assumption 2 (Interior Thresholds)**.** *The bounds on the distributions of liquidity shocks are sufficiently large:*

$$m \geqslant \frac{1}{2}(\nu_h - \nu_l)\max\left\{1, \frac{\delta(1-\pi/2)}{1-\delta(1-\pi/2)}\right\} \quad \textit{and} \quad e \geqslant \sqrt{\frac{3}{2}}(\nu_h - \nu_l).$$

This second assumption ensures that, for all prices offered in equilibrium and all realizations of ω, the thresholds $\bar{\varepsilon}^k_{t,j}, \underline{\varepsilon}^k_{t,j}$ lie in the interior of $[-e, e]$ for $j \in \{l, h\}$ and $k \in \{m, c\}$, i.e., that some traders always buy/sell in equilibrium.

Assumption 3 (Equal Valuations)**.** *On average, dealers and traders have the same valuation for an asset, i.e., $\nu_j = c_j$ for $j \in \{l, h\}$.*

This last assumption allows for a more direct comparison with models often used in finance (such as Glosten and Milgrom, 1985), and also simplifies the analysis.

3.2 Learning

The assumption that ω is uniformly distributed greatly simplifies the dealers' learning process. To see why, note from (16) that the updating process depends on current beliefs, μ, and the likelihood ratio

$$\frac{f(S - R_l(\mu^+))}{f(S - R_h(\mu^+))}.$$

When ω is uniformly distributed, $f(\omega) = \frac{1}{2m}$ for all $\omega \in [-m, m]$ and $f(\omega) = 0$ for all $\omega \notin [-m, m]$. Hence, either the signal that dealers observe is uninformative or it is fully revealing about the state $j \in \{l, h\}$.

Formally, let $\Sigma_j(\mu)$ denote the set of signals (i.e., the values of aggregate trading volume) that are only feasible when the asset is of quality j, given current beliefs μ, and let Σ_b denote the set of signals that are feasible in both states, l and h, so that

$$\mu^+(\mu, S) = \begin{cases} 0 & \text{if } S \in \Sigma_l(\mu) \\ \mu & \text{if } S \in \Sigma_b(\mu) \\ 1 & \text{if } S \in \Sigma_h(\mu). \end{cases}$$

We conjecture, and later confirm, that

$$\Sigma_l(\mu) = [-m + R_l(0), -m + R_h(\mu)) \quad (30)$$

$$\Sigma_b(\mu) = [-m + R_h(\mu), m + R_l(\mu)] \quad (31)$$

$$\Sigma_h(\mu) = (m + R_l(\mu), m + R_h(1)]\,. \quad (32)$$

In words, suppose the true asset quality is $j = h$. If the signal does not reveal the true asset quality, then $\mu^+ = \mu$. Moreover, we will show below that reservation values are increasing in μ, so that $R_h(\mu) \leqslant R_h(1)$. Therefore, under the candidate equilibrium, the minimum realization for $S = \omega + R_j$ when $j = h$ is $-m + R_h(\mu)$; any $S < -m + R_h(\mu)$ is only feasible if $j = l$. Similar reasoning can be used to explain (31)–(32). Note that

$$\Sigma_b(\mu) \neq \emptyset \;\Leftrightarrow\; R_h(\mu) - R_l(\mu) < 2m.$$

Assumption 2 ensures that valuations always satisfy this condition.

Let $p(\mu)$ denote the probability that the signal $S = \omega + R(\mu) \in \Sigma_l \cup \Sigma_h$, i.e., the probability that the quality of the asset is fully revealed to the dealers. When ω is uniformly distributed over the support $[-m, m]$, we have

$$p(\mu) = \frac{R_h(\mu) - R_l(\mu)}{2m}.$$

Since the expected number of periods before the quality is revealed is the inverse of $p(\mu)$, the following insight follows immediately.

Remark 1. *The expected speed of learning depends positively on* $R_h(\mu) - R_l(\mu)$.

Intuitively, learning occurs quickly when investors behave very differently when the asset is of high or low quality, i.e., when $R_h(\mu) - R_l(\mu)$ is relatively large. When investors' behavior is less dependent on asset quality, and $R_h(\mu) - R_l(\mu)$ is relatively small, it is more difficult for dealers to extract information from trading volume, and learning occurs more slowly.

3.3 Prices

We now derive equilibrium bid and ask prices in matches when a trader meets a single, monopolist dealer, and in matches when a trader meets competing dealers. Two aspects of our parametric specification make it possible to derive relatively simple pricing equations. First, the extreme learning process described above, which followed from the uniform distribution of ω, implies a straightforward relationship between current prices and future beliefs: beliefs are stationary until the state of the world is known with certainty. Second, given the uniform distribution over ε, the demand and supply functions that the dealers face are linear.

To start, it is helpful to define the expected continuation value of a trader when the asset quality is $j \in \{l, h\}$ and current beliefs are μ:

$$r_j(\mu) = \mathbb{E}_\omega\left[R_j\left(\mu^+\right)\right] = (1 - p(\mu)) R_j(\mu) + p(\mu) R_j(\mathbf{1}[j = h]) \tag{33}$$

Given this notation, it is straightforward to establish that the optimal price that a dealer offers when she is a monopolist is given by

$$\begin{aligned} B^m(\mu) &= \frac{\mathbb{E}_j r_j(\mu) + \mathbb{E}_j v_j - e}{2} \\ A^m(\mu) &= \frac{\mathbb{E}_j r_j(\mu) + \mathbb{E}_j v_j + e}{2} \end{aligned}$$

In words, the bid and ask prices are simply the average of the expected value of the dealer and the trader, adjusted by a markup term $\frac{e}{2}$. Importantly, the density of marginal buyers and marginal sellers is the same in both states of the world, so that the covariance between, e.g., $g\left(\overline{\varepsilon}_j\right)$ or $g\left(\underline{\varepsilon}_j\right)$ and $j \in \{l, h\}$ is zero. From the optimal pricing equations, (8) and (9), this implies that the adverse selection term disappears when the trader meets with a single dealer. Moreover, in this case, the bid-ask spread is equal to e for all values of beliefs.

When a trader meets with two dealers, the bid and ask price consistent with zero profits are given by

$$\begin{aligned} B^c(\mu) &= \frac{\mathbb{E}_j r_j + \mathbb{E}_j v_j - e}{2} + \frac{1}{2}\sqrt{\left(e + \mathbb{E}_j\left(v_j - r_j\right)\right)^2 - 4\mathrm{Cov}\left(r_j, v_j\right)} \\ A^c(\mu) &= \frac{\mathbb{E}_j r_j + \mathbb{E}_j v_j + e}{2} - \frac{1}{2}\sqrt{\left(e - \mathbb{E}_j\left(v_j - r_j\right)\right)^2 - 4\mathrm{Cov}\left(r_j, v_j\right)}. \end{aligned}$$

In the Appendix we show that, under Assumptions 1–3,

$$\mathbb{E}_j r_j(\mu) = \mathbb{E}_j v_j. \tag{34}$$

Using this property, we can simplify the bid and ask prices as follows:

$$B^m = \mathbb{E}_j v_j - \frac{e}{2} \tag{35}$$

$$A^m = \mathbb{E}_j v_j + \frac{e}{2} \tag{36}$$

$$B^c = \mathbb{E}_j v_j - \frac{e}{2} + \sqrt{\left(\frac{e}{2}\right)^2 - \text{Cov}\left(r_j, v_j\right)} \tag{37}$$

$$A^c = \mathbb{E}_j v_j + \frac{e}{2} - \sqrt{\left(\frac{e}{2}\right)^2 - \text{Cov}\left(r_j, v_j\right)}. \tag{38}$$

Equations (35)–(38) illustrate that the model with uniformly distributed shocks and two types of meetings (monopolist and competitive) is tractable enough to admit analytcial solutions, and yet rich enough to capture the key economic mechanisms at work. On the one hand, the markup term in B^m and A^m implies that the dealers capture some rents that are unrelated to asymmetric information, as in, e.g., Duffie et al. (2005). On the other hand, the adverse selection term in B^c and A^c captures the portion of the bid-ask spread that is attributed to adverse selection, as in, e.g., Glosten and Milgrom (1985). Since $\mathbb{E}_j r_j = \mathbb{E}_j v_j$, the term $\text{Cov}\left(r_j, v_j\right)$ is maximized at $\mu = \frac{1}{2}$, which implies that this portion of the bid-ask spread is also maximized at $\mu = \frac{1}{2}$, i.e., when the information asymmetry between traders and dealers is maximal.

3.4 Reservation Values

Using the optimal bid and ask prices derived above, we show in Appendix B.2 that the reservation value of an investor, given current beliefs $\mu \in (0, 1)$ and asset quality j, can be written as

$$R_j(\mu) = (1-\delta)v_j + \delta r_j(\mu) + \delta\pi \sum_{k=c,m} \alpha_k \Omega_j^k(\mu), \tag{39}$$

where $r_j(\mu)$ is defined in (33) and $\Omega_j^k(\mu)$ is what we call the *net option value* of holding a quality j asset in a type k meeting, i.e., the option value of selling the asset net of the option value of buying the asset. This net option value derives from the fact that acquiring a unit of the asset offers the investor the option value of selling it a later date but—given our assumption that investors can only hold one unit of the asset at a time—acquiring a unit of the asset also implies forfeiting the option value of buying a unit of the asset at a later date, too.

Under Assumptions 1–3, one can show that

$$\Omega_j^k(\mu) = \frac{B^k - A^k + 2e}{2e}\left(\frac{A^k + B^k}{2} - r_j(\mu)\right). \tag{40}$$

Intuitively, Ω_j^k is the expected surplus an investor earns from future trading opportunities in type $k \in \{m, c\}$ meetings, given that the asset is of quality $j \in \{l, h\}$. The first term on the left-hand side of (40) is the ex ante probability (before ε and ω are realized) that the investor will optimally choose to trade in a type k meeting, given prices A^k and B^k. The second term,

$$\frac{A^k + B^k}{2} - r_j(\mu) = \frac{B^k - r_j(\mu) - \left(r_j(\mu) - A^k\right)}{2},$$

is the expected difference between the surplus the investor will earn from selling the asset at a later date and the surplus he could have earned from buying an asset at a later date.

Since the bid and ask prices are independent of asset quality, the net option value is decreasing in the expected continuation value r_j. One can show that $r_h(\mu) > r_l(\mu)$, which implies that the net option value is larger when the asset is of quality l. We highlight this property in the remark below, as it will play a key role in the ensuing results.

Remark 2. *The net option value is decreasing in ν_j, so that $\Omega_l^k(\mu) \geqslant \Omega_h^k(\mu)$ for any $\mu \in (0,1)$, $k \in \{m, c\}$.*

3.5 Equilibrium Characterization

To characterize the equilibrium, we can use (39)–(40) to write

$$\begin{aligned} R_h(\mu) - R_l(\mu) &= (1-\delta)(\nu_h - \nu_l) + \delta(r_h(\mu) - r_l(\mu)) + \delta\pi \sum_{k=c,m} \alpha_k \left[\Omega_h^k(\mu) - \Omega_l^k(\mu)\right] \\ &= (1-\delta)(\nu_h - \nu_l) + \delta(r_h(\mu) - r_l(\mu)) - \delta\pi \sum_{k=c,m} \alpha_k \frac{B^k - A^k + 2e}{2e}(r_h(\mu) - r_l(\mu)) \end{aligned} \tag{41}$$

Then, using (35)–(38), along with

$$p(\mu) = \frac{R_h(\mu) - R_l(\mu)}{2m},$$

we establish in the Appendix that (41) can be written as a single equation in one unknown, p, given beliefs μ and parameters $\Xi \equiv (\delta, \pi, \alpha_c, \nu_h, \nu_l)$. In particular, let $p^\star(\mu)$ denote the solution to $Z(p, \mu; \Xi) = 0$, where

$$\begin{aligned} Z(p,\mu;\Xi) &= -2mp + (1-\delta)(\nu_h - \nu_l) + \delta(1-\pi)(p(\nu_h - \nu_l) + 2mp(1-p)) \\ &\quad - \frac{\delta\pi\alpha_c}{2}\sqrt{1 - \frac{4}{e^2}(\nu_h - \nu_l)\mu(1-\mu)[2m(1-p)p + p(\nu_h - \nu_l)]}\,(p(\nu_h - \nu_l) + 2mp(1-p)). \end{aligned} \tag{42}$$

Proposition 1. *Under assumptions 1–3, there exists a unique $p^\star(\mu)$ such that $Z(p^\star(\mu), \mu; \Xi) = 0$.*

Importantly, solving for the speed of learning is sufficient for a full characterization of the model: one can use $p^\star(\mu)$ to construct the reservation values, $\{R_j(\mu)\}_{j\in\{l,h\}}$, along with equilibrium prices, $\{A^k, B^k\}_{k\in\{m,c\}}$.[16] In the next section, we exploit several of the results derived above, along with the characterization afforded by equation (42), to understand how changes to the parameters affect equilibrium outcomes.

[16]A key result is that $\mathbb{E}_j r_j = \mathbb{E}_j \nu_j$. This implies that $R_j(\mathbf{1}[j = h]) = \nu_j$, i.e., that reservation values under full information are equal to the value of owning the asset and not trading it.

3.6 Comparative Statics

In this section, we explore how bid-ask spreads change in response to changes in the underlying economic environment, with a focus on understanding the interaction between search frictions, asymmetric information, and learning. We proceed in two steps. We start by examining how reservation values, the speed of learning, and the bid-ask spread depend on the dealers' beliefs, μ. The results we derive are informative for our next step, where we explore the effects of changing the degree of search frictions in our model, either by changing the frequency of trading opportunities (π) or the fraction of trading opportunities that are competitive (α_c).

Comparative Statics with Respect to Beliefs

Since $A^m - B^m = e$ for all μ, the effect of beliefs on the bid-ask spread operate exclusively through the prices in competitive meetings. In the Appendix, we establish the following results.

Lemma 1. *The bid-ask spread in competitive meetings,* $A^c - B^c$*, the difference in reservation values,* $R_h - R_l$*, and the probability that the true asset quality is revealed,* p*, are all hump-shaped in* μ *with a maximum at* $\mu = 1/2$.

Intuitively, the bid-ask spread is largest when uncertainty is maximal, i.e., when $\mu = 1/2$. When the bid-ask spread is widest, the probability of trade is smallest, as investors are more likely to draw idiosyncratic liquidity shocks that lie in the "inaction region." This causes the difference in the net option values of trading, $|\Omega_h - \Omega_l|$, to decline. Intuitively, a decrease in the probability of trade causes a disproportionate decline in the option value to sell when the asset quality is l, and in the option value of buying when the asset quality is h, making $\Omega_h - \Omega_l$ less negative. As a result, the difference in the reservation values, $R_h - R_l$, widens. As investors' behavior in the two states of the world becomes more distinguishable, the probability that the true asset quality is revealed, p, increases and learning occurs, on average, more quickly.

The Effect of Search frictions

Now consider a decrease in the severity of search frictions. In what follows, we will focus on the effect of an increase in π, though we show in the Appendix that similar results obtain for an increase in α_c. Our first result, summarized below, implies that an increase in π unambiguously slows down the learning process.

Proposition 2. *For any* $\mu \in (0, 1)$, $\frac{\partial p^*(\mu)}{\partial \pi} < 0$.

In the discussion above, more extreme values of μ caused bid-ask spreads to narrow and increased the probability of trade, causing $R_h - R_l$—and thus p—to fall. The intuition behind the result in Proposition 2 is similar: an increase in π causes a direct increase in the probability of trade, as opposed to operating through the bid-ask spread. As a result, the difference in net option values increases (i.e., becomes more negative), and $R_h - R_l \propto p$ falls. In words, in the presence of asymmetric information, an increase in the probability of trade again makes the option value of selling (buying) relatively high when the asset is quality l (h). As a result, the behavior of investors in the two states of the world becomes more similar and it takes longer, on average, to learn the state of the world.

The effect on bid-ask spreads, however, is more nuanced as there are two, opposing effects. The first, which we call the *static effect*, has the usual sign: holding beliefs constant, an increase in π causes spreads to shrink. As discussed above, an increase in π causes the difference in reservation values to narrow. This implies a decrease in the covariance of a trader's expected valuations, r_j, and a dealer valuations, ν_j. Since average bid-ask spreads are given by

$$\sum_k \alpha_k \left(A^k - B^k\right) = e - \alpha_c \sqrt{e^2 - 4\mathrm{Cov}\left(r_j, \nu_j\right)},$$

a decline in this covariance leads to a lower bid-ask spread (holding beliefs fixed). Intuitively, since the increase in π makes investors behave more similarly in the two states of the world—i.e., the likelihood of an investor buying or selling at a given price becomes more similar for $j \in \{l, h\}$—the problem of adverse selection is diminished and spreads fall.

However, even though an increase in π causes bid-ask spreads to fall for a given level of beliefs, in equilibrium beliefs are changing over time. This leads us to the second effect of increasing the frequency of trade, which we call the *dynamic effect*: since an increase in π implies that dealers will remain uncertain about the true asset quality for longer (Lemma 2), and bid-ask spreads are larger when dealers are more uncertain (Lemma 1), more frequent meetings ultimately leader to larger bid-ask spreads in the future.

To state this formally, we let $\sigma_{j,t}$ denote the average quoted bid-ask spread in period t when the asset is of quality $j \in \{l, h\}$. Formally,

$$\sigma_{t,j} \equiv \mathbb{E}_{\omega^t}\left[\sum_k \alpha_k \left(A_t^k - B_t^k\right) \mid j\right]$$

The following proposition shows that the dynamic effect of a higher π eventually dominates, leading to wider spreads.

Proposition 3. *There exists a* $\tau < \infty$ *such that* $\frac{\partial \sigma_{t,j}}{\partial \pi} > 0$ *for all* $t \geqslant \tau$.

To see this, we plug in the equilibrium bid and ask prices into the expression for $\sigma_{j,t}$:

$$\begin{aligned}
\sigma_{t,j} &= \mathbb{E}_{\omega^t}\left[\sum_k \alpha_k \left(A_t^k - B_t^k\right) \ \mid j\right] \\
&= \left(1-(1-p(\mu))^t\right)[e-\alpha_c e] + (1-p(\mu))^t \left(e - e\alpha_c \sqrt{1-\frac{4}{e^2}(v_h - v_l)\,\mu(1-\mu)(r_h(\mu)-r_l(\mu))}\right).
\end{aligned}$$

Differentiating with respect to π, we have

$$\begin{aligned}
\frac{\partial}{\partial \pi}\sigma_{t,j} =& -t(1-p(\mu))^{t-1}\frac{\partial p}{\partial \pi}\alpha_c e\left[1-\sqrt{1-\frac{4}{e^2}(v_h - v_l)\,\mu(1-\mu)(r_h(\mu)-r_l(\mu))}\right] \\
&+ (1-p(\mu))^t \alpha_c e \frac{\frac{4}{e^2}\mu(1-\mu)(v_h - v_l)\frac{\partial}{\partial \pi}(r_h(\mu)-r_l(\mu))}{2\sqrt{1-\frac{4}{e^2}(v_h - v_l)\,\mu(1-\mu)(r_h(\mu)-r_l(\mu))}}
\end{aligned}$$

The first term in the expression above is positive while the second term is negative. However, as t converges to infinity, the first term becomes large relative to the second term. Hence, $\frac{\partial}{\partial \pi}\sigma_{t,j} > 0$ for values of t that are sufficiently large.

4 General Model

In this section, we relax Assumptions 1–3 and confirm the results above numerically. Most importantly, we consider more general distributions for the aggregate and idiosyncratic liquidity shocks. This makes the analysis considerably more complicated, as the interaction between the optimal bid and ask prices, reservation values, and future prices and beliefs introduce a number of additional effects that were absent from our analysis with uniform distributions. While these complications make analytical results harder to derive, it remains fairly straightforward to solve the model on the computer. We describe the simple, iterative algorithm here:

1. Given a grid for μ and ω, we start with an initial guess for the reservation values, $R_j(\mu)$, $j = h, l$.
2. Given reservation values, we can determine traders' expectations of dealer beliefs μ^+ for each ω by solving (17).
3. Given the updating equations, we compute optimal prices for any beliefs μ.
4. The law of motion for dealer beliefs and the pricing formulae can then be combined to yield an updated guess for the reservation value functions $R_j(\mu)$, using the expression in (5).
5. We repeat this iterative process until convergence.

Using this algorithm, we solve the model and show that the main insights from Section 3 continue to hold, despite the presence of the additional feedback effects mentioned above.

4.1 Full Information Benchmark

VV: WE SHOULD LOSE THIS SUBSECTION, NO? As a benchmark, we first analyze the full information case, i.e., where dealers know the quality of the asset (but not the traders' liquidity shocks, which are still privately observed). In this case, reservation values and prices are constant across time, and it is fairly straightforward to conduct comparative statics on the matching frequency, π. We show that, under mild restrictions on the distributions of shocks, the standard relationships between trading frequency and bid-ask spreads emerge.

Condition 1. *Let $\tilde{G}$ be the distribution of $\omega + \epsilon$. $\tilde{G}$ is symmetric around 0 and $\frac{\tilde{G}}{\tilde{g}}$ is increasing and convex.*

This condition is satisfied by a number of commonly used distributions, including the normal, logistic and the Pareto distributions. It implies that the full information markup charged by monopolist dealers increases with volume at an increasing rate. This assumption is sufficient to ensure that full information spreads shrink as trading opportunities become more frequent.

Proposition 4. *When asset quality is common knowledge, and Condition 1 holds, the spread charged by monoplist dealers, $A_j^m - B_j^m$, is decreasing in π for $j \in \{l, h\}$.*

To see the intuition, we focus on the case with $\nu_j > c_j$, i.e. when dealers value the asset more than the average obtained by the traders upon exit. This implies that, on average, traders are more likely to be sellers than buyers. As a result, the option to sell is worth more than the option to buy, i.e. the net option value is positive, $\Omega_j > 0$, . Indeed, in the proof of Proposition 4, we show that

$$\frac{dR_j}{d\pi} = \Omega_j + \pi \frac{\partial \Omega_j}{\partial \pi} > 0,$$

so that more frequent trading opportunities raise reservation values. Moreover, under Condition 1, we also show that the inverse semi-elasticity of the investors that are buying the asset is greater than the inverse semi-elasticity of the investors that are selling the asset. As a result, the increase in π causes the bid price to increase more than the ask price, and the spread falls. Since this analogous to the relationship between the frequency of trading opportunities and spreads in the canonical Duffie-Garleanu-Pedersen framework, we refer to it as the "DGP effect."

4.2 Asymmetric Information and Non-Uniform Distributions

Let us now return to the model with asymmetric information, where asset quality is privately observed by the traders. As noted above, this model has to be solved numerically. We impose assumptions on the distributions of liquidity shocks and the values of the model's key parameters. While all of our results depend, in some sense, on the parametric assumptions that are imposed, it is important to note that there is nothing special about *these* particular assumptions; that is, the results reported below obtain in a large range of the parameter space.

We assume that both aggregate and idiosyncratic liquidity shocks are drawn from a mean-zero normal distribution, i.e. $\omega \sim N(0, \sigma_\omega^2)$ and $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. The normality of ε implies that the bid and ask prices under monopoly are no longer constant (as was the case in the uniform-uniform model of section 3), since both the market power and asymmetric components now vary with beliefs, μ_t. The normality of ω means that learning no longer takes the stark form as in Section 3. These differences, however, mean that we no longer have analytical tractability and have to resort to numerical computations.

Next, we assign values to parameters. This is not intended to be a full-fledged calibration: our goal is to show that the forces we have identified are significant present for an empirically reasonable parameterization. We apply the model to a widely studied over-the-counter market, that for US corporate bonds. We interpret differences in quality as stemming from changes in credit ratings. Consider the a bond that is rated AAA. Conditional on not being downgraded, the expected payoff of the bond upon maturity is c_h. If it is downgraded (to say, AA), however, the expected payoff drops to c_l. Since bid-ask spreads and beliefs are only a function of the *relative* payoff in the two states (i.e. $c_h - c_l$), we can normalize $c_l = 0$. The relative payoff is then mapped to the drop in the market price of the bond in the event of a downgrade. Feldhütter (2012) reports that average spreads on AA bonds are about 26 bps higher than on AAA-rated ones.[17]. For a 5-year par bond (face value $100) with a coupon of 3%, this spread difference translates into a price change of $1.16. The initial belief μ_0 is chosen to match the unconditional rating transition probabilities: according to the 2016 Annual Global Corporate Default Study and Rating Transitions published by Standard & Poor's, the likelihood of a AAA-rated US corporate bond retaining that rating over a 5-year horizon is 0.50. Accordingly, we set $\mu_0 = 0.50$.

To pin down the sizes of liquidity shocks and meeting probabilities, we use estimates from Feldhutter (2012). That paper estimates the parameters of a continuous time model of over-the-counter trading along

[17]See Table 1 of that paper. The average sell spread on a AAA bond is 9 bps (on trades greater than USD 100,000 in size), while the corresponding spread for AA bonds is 35 bps.

the lines of Duffie et al. (2005) using data on secondary market transactions in US corporate bonds. We map his estimate of holding costs, the sole source of gains from trade in his environment, to the magnitude of liquidity shocks in our model.[18] Assuming that the relative magnitude of aggregate and idiosyncratic components are of equal magnitude, this procedure leads to $\sigma_\omega^2 = \sigma_\varepsilon^2 = 0.16$.

Feldhutter (2012) also provides estimates for the arrival rate of meetings with dealers from the perspective of traders in the market. His point estimate, an annualized rate of 40, can directly be mapped into the parameters governing our meeting technology[19] : the probability of meeting at least 1 dealer (π) as well as the conditional probability of meeting more than 1 dealer (α_c). Interpreting a period in our model as a week (5 business days) or 0.02 years (a year is assumed to have 250 business days) yields our baselines values $(\pi, \alpha_c) = (0.55, 0.25)$.

This leaves δ. As discussed in Section 2.1, this parameter can be interpreted in multiple ways – either as the likelihood of the asset not maturing in any given period or as the probability of an individual trader remaining active in the market. We adopt a baseline value of $\delta = 0.9$, which implies that, conditional on not trading, a trader remains in the market for about 10 weeks.

Figure 1 plots key equilibrium objects for two different values of π, namely 0.25 and 0.75. The top left panel plots the difference between traders' reservation values in the high and low state. It shows that $R_h(\mu) - R_l(\mu)$ decreases with π for all μ. The intuition behind this result is the same as in the uniform-uniform model: more frequent trading opportunities increases the weight of the net option value $\Omega_j(\mu')$ in the reservation value $R_j(\mu)$. Since the difference $\Omega_h(\mu') - \Omega_l(\mu') < 0$, this has the effect of bringing the reservation values closer to each other.

The remaining panels plot spreads – the bottom ones plot spreads quoted by the market-maker under monopoly and competition, while the top right one plots the average (computed using α_c as the weight). They show that spreads are non-monotonic in μ: this occurs because of the asymmetric information component, which is largest when uncertainty is highest, i.e. when μ is in an intermediate range. More interestingly, the figure also reveals that spreads are decreasing in π for any given μ, i.e., that more frequent trading opportunities reduces spreads for any fixed set of beliefs. Intuitively, increasing π reduces $R_h - R_l$ for any given μ causing traders to act more alike in both states of the world. As a result,

[18] Specifically, he estimates a flow holding cost of 2.91, which lasts for an average of 0.31 years or equivalently, a total cost of $(2.91)(0.31) = 0.90$. We interpret this as the average difference in valuations between agents receiving positive liquidity shocks and those that receive negative shocks. Under the assumption of normally distributed liquidity shocks, this transaltes into a variance of 0.32.

[19] Feldhutter (2012) estimates this separately for different trade sizes. We use the one corresponding to the smallest trade size.

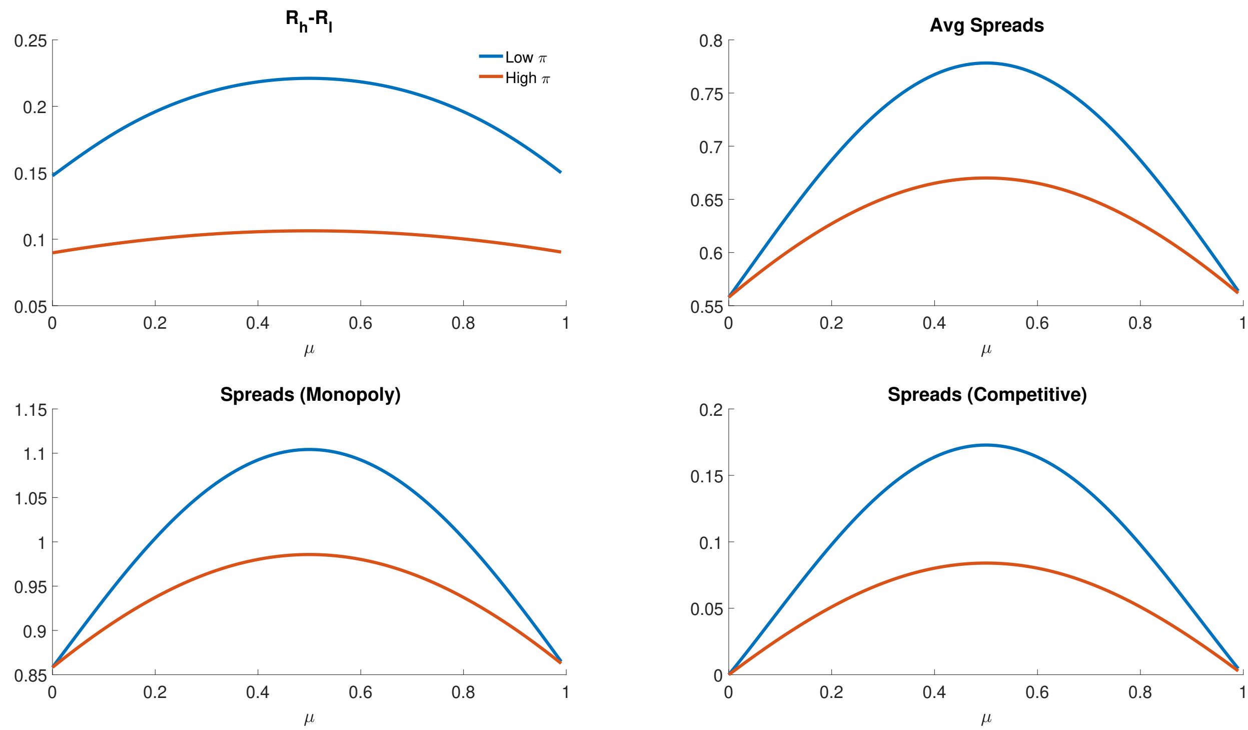

Figure 1: Effect of π on prices and spreads.

dealers' optimal prices require a smaller adjustment for asymmetric information. Consider the monopolist's quotes. Recall from equations (8) and (9) that the asymmetric information component of the dealers' optimal prices depend on the density of traders around the liquidity shock thresholds, which are functions of the reservation values (and prices). Higher π pushes the reservation values closer, and through them, the thresholds. As a result, the asymmetric information component of the spread shrinks. To put it differently, when trading decisions are driven to a greater extent by liquidity shocks, as opposed to asset quality, there is less adverse selection from the dealers' perspective, leading to tighter spreads.

However, this effect on spreads obtains with beliefs held fixed, i.e. it is essentially a static effect. There is another, dynamic effect of the decline in the difference between reservation values – it also slows down learning, exactly as in the special case. As a result, dealers remain relatively more uncertain about the true asset quality. This force reduces the rate at which the asymmetric information component of spreads shrinks over time, keeping spreads higher.

Figure 2 illustrates this dynamic effect. It plots the evolution of beliefs and spreads over time for two different values of π, under the assumption that the true quality of the asset is h. The panels plot the average values across 10,000 sample paths – the realized path of dealer beliefs and, therefore, prices depends on the realized sequence of ω_t. Dealer beliefs (top left panel) start at $\mu_0 = 0.5$ in both cases

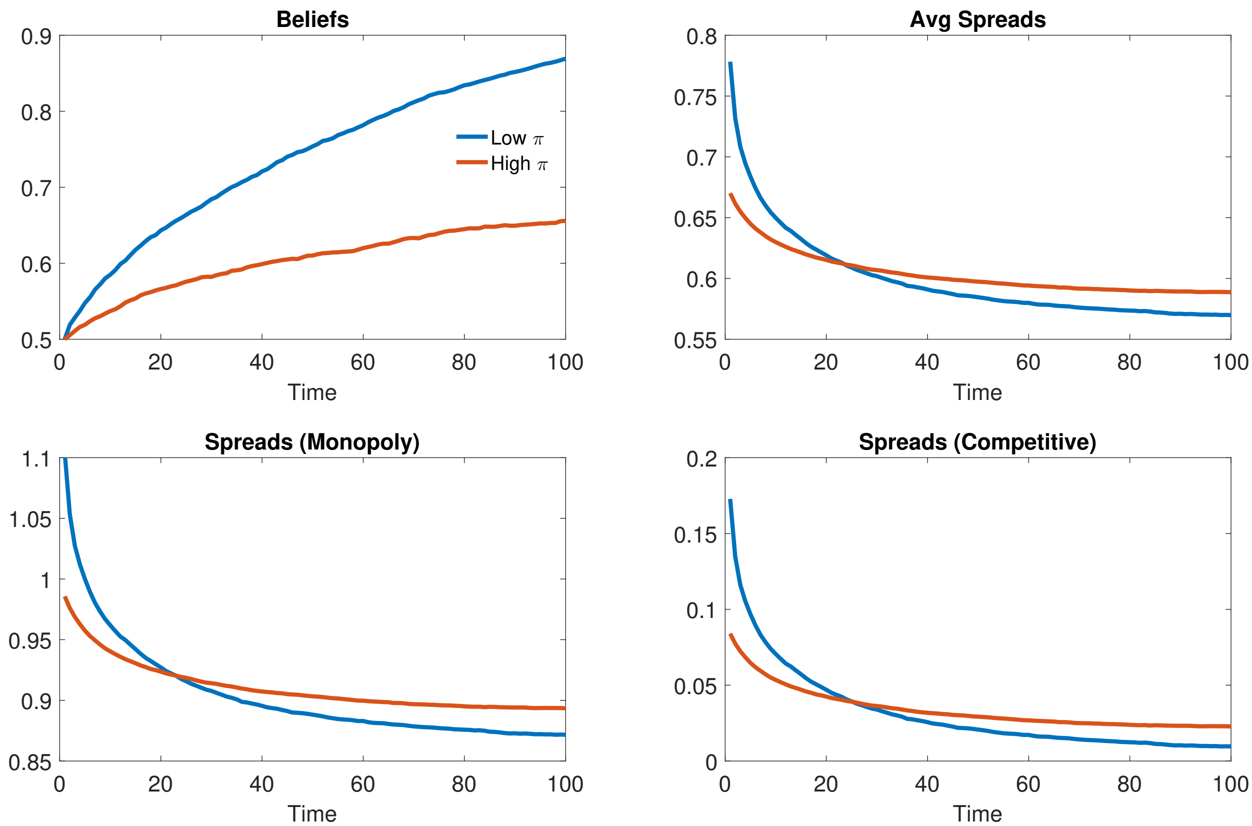

Figure 2: Effect of π over time.

and drift upwards: since the true state is h, both lines eventually converge to 1. However, the pace is slower when π is higher; more frequent trading opportunities leads to slower learning. Spreads are tighter initially in the high π case, but eventually end up wider. This is because, initially, beliefs are very similar in both cases (since they both start at the same level by assumption), so the static effect dominates and spreads narrow with higher π. Over time, however, the differential pace of learning kicks in, keeping uncertainty high and spreads wide in the high π scenario, relative to the low π case. This is true both under monopoly and competitive spreads, as the bottom two panels show.

4.3 A Stationary Version

In this subsection, we modify our baseline framework to allow the asset quality to change over time. Specifically, in each period, it switches with probability ρ. This means that, in the long run, the market will enter a stochastic steady state, characterized by a distribution over beliefs (and therefore, prices and allocations). We point out the main changes from the baseline non-stationary version below.

Dealer beliefs about quality at the end of a given period are still given by equation (26), but their beliefs in the following period are adjusted to account for the possibility that quality has changed. Let μ denote the beliefs at the end of the period. This can be mapped into beliefs at the time of pricing in the following period by

$$\mathcal{M}(\mu) = \mu(1-\rho) + (1-\mu)\rho \tag{43}$$

Next, the reservation values reflect the possibility of quality changes in the following period and are given by:

$$\begin{aligned}
R_j(\mu) \quad = \quad & (1-\delta)c_j + \delta(1-\pi)(1-\rho)\int_\omega R_j\left(\tilde{\mu}_j^+(\mathcal{M}(\mu),\omega)\right)dF(\omega) + \delta(1-\pi)\rho\int_\omega R_{-j}\left(\tilde{\mu}_{-j}^+(\mathcal{M}(\mu),\omega)\right)dF(\omega) \\
+ \quad & \delta\pi(1-\rho)\int_\omega \left\{ \sum_{k\in\{m,c\}} \alpha^k \left[B^k(\mathcal{M}(\mu))G\left(\underline{\varepsilon}_j^k(\mathcal{M}(\mu),\omega)\right)\right.\right. \\
+ \quad & \left.\left.\int_{\underline{\varepsilon}_j^k(\mathcal{M}(\mu),\omega)}^{\overline{\varepsilon}_j^k(\mathcal{M}(\mu),\omega)} \left[\omega + \varepsilon + R_j\left(\tilde{\mu}_j^+(\mathcal{M}(\mu),\omega)\right)\right] dG(\varepsilon) + A^k(\mathcal{M}(\mu))\left[1 - G\left(\overline{\varepsilon}_j^k(\mathcal{M}(\mu),\omega)\right)\right]\right]\right\} dF(\omega). \\
+ \quad & \delta\pi\rho\int_\omega \left\{ \sum_{k\in\{m,c\}} \alpha^k \left[B^k(\mathcal{M}(\mu))G\left(\underline{\varepsilon}_{-j}^k(\mathcal{M}(\mu),\omega)\right)\right.\right. \\
+ \quad & \left.\left.\int_{\underline{\varepsilon}_{-j}^k(\mathcal{M}(\mu),\omega)}^{\overline{\varepsilon}_{-j}^k(\mathcal{M}(\mu),\omega)} \left[\omega + \varepsilon + R_{-j}\left(\tilde{\mu}_{-j}^+(\mathcal{M}(\mu),\omega)\right)\right] dG(\varepsilon) + A^k(\mathcal{M}(\mu))\left[1 - G\left(\overline{\varepsilon}_{-j}^k(\mathcal{M}(\mu),\omega)\right)\right]\right]\right\} dF(\omega).
\end{aligned}$$

The dealers' valuation of the asset in the two states $(\tilde{v}_h, \tilde{v}_l)$ solve the following linear system

$$\begin{aligned}
\tilde{v}_h &= (1-\delta)v_h + \delta(1-\rho)\tilde{v}_h + \delta\rho\tilde{v}_l \\
\tilde{v}_l &= (1-\delta)v_l + \delta(1-\rho)\tilde{v}_l + \delta\rho\tilde{v}_l
\end{aligned}$$

The pricing equations take the same form as the baseline non-stationary version with $(\tilde{v}_h, \tilde{v}_l)$ replacing (v_h, v_l). This version of the model is also easy to solve numerically.

In order to highlight the effect of π, we solve the model with the switching probability ρ set to 0.5% for 3 different values of π (all other parameters held fixed at their baseline values). The results are presented in Figure 3. The top left panel plots the distribution of beliefs in the long run. It shows that higher π shifts the mass of the distribution towards the center, i.e. where dealers are more uncertain about asset quality. Intuitively, this occurs for the same reason as in the baseline non-stationary version – more frequent trading opportunities bring reservation values of the traders in the two states closer to each, making trading less informative and reducing the speed of learning. The effects on average spreads (in the remaining 3 panels) are more complicated. On the one hand, higher π reduces adverse selection (and therefore, pushes spreads lower) for a given level of beliefs μ. But, it also pushes the long-run distribution towards intermediate μ's where spreads are higher. Whether average spreads are wider or narrower depends on which of these effects dominates. For example, when spreads rise from 0.55 (our baseline value) to 0.75, the change in the distribution is stronger than the effect on spreads with fixed beliefs, leading to wider spreads on average. The opposite happens when π rises further from 0.75 to 0.95.

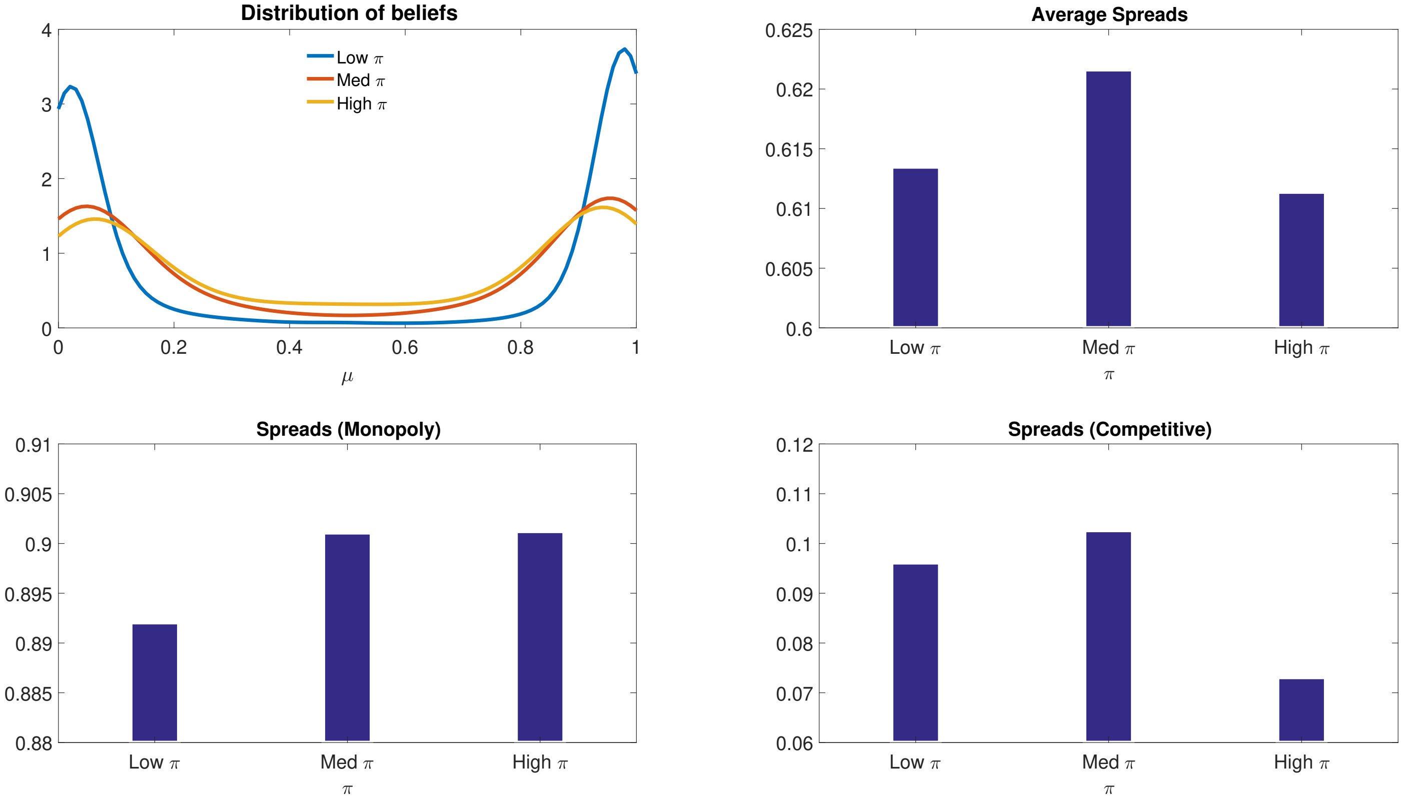

Figure 3: Effect of π in stationary version

Note: Low $\pi = 0.55$, Med $\pi = 0.75$, High $\pi = 0.95$

5 Conclusion

In the previous sections, we have laid out a unified framework for analyzing asset markets characterized by both informational and search frictions. Our results uncover novel interactions between these two fundamental imperfections, overturning conventional wisdom based on studying them in isolation. They help reconcile a puzzling feature of many OTC markets in recent years – even as technological and policy-induced changes have reduced trading frictions and increased transparency, measures of liquidity (e.g. bid-ask spreads) have largely remained unchanged, or even widened. Our analysis suggests that this may not be particularly surprising in markets where both types of frictions are present.

There are many directions for future research. A key challenge in many financial markets is to quantify the severity of different types of frictions. This is particularly relevant for regulators, who can then devise appropriate policies. In ongoing work, we use the theoretical framework developed in this paper to inform an empirical strategy that can disentangle information and search frictions from observable transaction data. Another interesting direction would be add other sources of market illiquidity (e.g. inventory costs).

References

ACHARYA, V. V. AND T. C. JOHNSON (2007): "Insider trading in credit derivatives," *Journal of Financial Economics*, 84, 110–141.

AFONSO, G. (2011): "Liquidity and congestion," *Journal of Financial Intermediation*, 20, 324–360.

BABUS, A. AND P. KONDOR (2016): "Trading and Information Diffusion in Over&the&Counter Markets," *working paper*.

BETHUNE, Z., B. SULTANUM, AND N. TRACHTER (2016): "Private Information in Over-the-Counter Markets," .

BIAIS, B., T. FOUCAULT, AND S. MOINAS (2015): "Equilibrium fast trading," *Journal of Financial Economics*, 116, 292–313.

BLOUIN, M. R. AND R. SERRANO (2001): "A decentralized market with common values uncertainty: Non-steady states," *The Review of Economic Studies*, 68, 323–346.

BURDETT, K. AND K. L. JUDD (1983): "Equilibrium Price Dispersion," *Econometrica*, 51, 955–969.

CAMARGO, B. AND B. LESTER (2014): "Trading dynamics in decentralized markets with adverse selection," *Journal of Economic Theory*, 153, 534–568.

CHIU, J. AND T. V. KOEPPL (2016): "Trading dynamics with adverse selection and search: Market freeze, intervention and recovery," *The Review of Economic Studies*, 83, 969–1000.

CHOI, M. (2016): "Imperfect Information Transmission and Adverse Selection in Asset Markets," *Unpublished paper, University of Iowa.*

COPELAND, T. E. AND D. GALAI (1983): "Information effects on the bid-ask spread," *the Journal of Finance*, 38, 1457–1469.

DU, S. AND H. ZHU (2017): "What is the optimal trading frequency in financial markets?" *The Review of Economic Studies*, 84, 1606–1651.

DUFFIE, D., N. GÂRLEANU, AND L. H. PEDERSEN (2005): "Over-the-Counter Markets," *Econometrica*, 73, 1815–1847.

DUFFIE, D., S. MALAMUD, AND G. MANSO (2009): "Information percolation with equilibrium search dynamics," *Econometrica*, 77, 1513–1574.

DUFFIE, D. AND G. MANSO (2007): "Information percolation in large markets," *The American economic review*, 97, 203–209.

FELDHUTTER, P. (2012): "The Same Bond at Different Prices: Identifying Search Frictions and Selling Pressures," *The Review of Financial Studies*, 25, 1155–1206.

FELDHÜTTER, P. (2012): "Short-term corporate bond yield spreads," .

FUCHS, W. AND A. SKRZYPACZ (2015): "Government interventions in a dynamic market with adverse selection," *Journal of Economic Theory*, 158, 371–406.

GEHRIG, T. (1993): "Intermediation in search markets," *Journal of Economics & Management Strategy*, 2, 97–120.

GLOSTEN, L. R. AND P. R. MILGROM (1985): "Bid, ask and transaction prices in a specialist market with heterogeneously informed traders," *Journal of financial economics*, 14, 71–100.

GOLOSOV, M., G. LORENZONI, AND A. TSYVINSKI (2014): "Decentralized trading with private information," *Econometrica*, 82, 1055–1091.

GROSSMAN, S. J. AND J. E. STIGLITZ (1980): "On the impossibility of informationally efficient markets," *The American economic review*, 70, 393–408.

GUERRIERI, V. AND R. SHIMER (2014): "Dynamic Adverse Selection: A Theory of Illiquidity, Fire Sales, and Flight to Quality," *American Economic Review*, 104, 1875–1908.

HELLWIG, M. F. (1980): "On the aggregation of information in competitive markets," *Journal of economic theory*, 22, 477–498.

HENDERSHOTT, T. AND P. C. MOULTON (2011): "Automation, speed, and stock market quality: The NYSE's hybrid," *Journal of Financial Markets*, 14, 568–604.

HUANG, R. D. AND H. R. STOLL (1997): "The components of the bid-ask spread: A general approach," *The Review of Financial Studies*, 10, 995–1034.

HUGONNIER, J., B. LESTER, AND P.-O. WEILL (2014): "Heterogeneity in decentralized asset markets," Tech. rep., National Bureau of Economic Research.

KAYA, A. AND K. KIM (2015): "Trading dynamics with private buyer signals in the market for lemons," .

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

KYLE, A. S. (1985): "Continuous auctions and insider trading," *Econometrica: Journal of the Econometric Society*, 1315–1335.

LAGOS, R. AND G. ROCHETEAU (2009): "Liquidity in Asset Markets with Search Frictions," *Econometrica*, 77, 403–426.

LAUERMANN, S. AND A. WOLINSKY (2016): "Search with adverse selection," *Econometrica*, 84, 243–315.

LESTER, B., G. ROCHETEAU, AND P.-O. WEILL (2015): "Competing for order flow in OTC markets," *Journal of Money, Credit and Banking*, 47, 77–126.

MADHAVAN, A., D. PORTER, AND D. WEAVER (2005): "Should securities markets be transparent?" *Journal of Financial Markets*, 8, 265–287.

MENKVELD, A. J. AND M. A. ZOICAN (2017): "Need for speed? Exchange latency and liquidity," *The Review of Financial Studies*, 30, 1188–1228.

PAGNOTTA, E. AND T. PHILIPPON (2015): "Competing on speed," *Econometrica*, forthcoming.

SPULBER, D. F. (1996): "Market making by price-setting firms," *The Review of Economic Studies*, 63, 559–580.

WOLINSKY, A. (1990): "Information revelation in a market with pairwise meetings," *Econometrica: Journal of the Econometric Society*, 1–23.

Appendix

A Proofs

A.1 Proof of Proposition 1

Proof. For ease of notation, we supress depends of R_j, r_j, p to μ. From (41), we have

$$\begin{aligned} R_h - R_l &= (1-\delta)(v_h - v_l) + \delta(r_h - r_l) - \delta\pi \sum_k \alpha_k \frac{B^k - A^k + 2e}{2e}(r_h - r_l) \\ &= (1-\delta)(v_h - v_l) + \delta(r_h - r_l) - \delta\pi\alpha_c \frac{r_h - r_l}{2} \\ &\quad - \delta\pi\alpha_m \frac{\sqrt{e^2 - 4Cov_j(r_j, v_j)} + e}{2e}(r_h - r_l) \end{aligned} \tag{44}$$

Note that $r_j = (1-p) R_j + p v_j$. Therefore,

$$\begin{aligned} Cov_j(r_j, v_j) &= (1-p) Cov_j(R_j, v_j) + p Var(v_j) \\ &= (1-p) Cov_j(R_j, v_j) + p\mu(1-\mu)(v_h - v_l)^2 \end{aligned}$$

Moreover,

$$\begin{aligned} \mathbb{E}_j R_j = \mathbb{E}_j v_j &\to \mu R_h + (1-\mu) R_l = \mu v_h + (1-\mu) v_l \\ R_l = \frac{\mathbb{E}_j v_j - \mu R_h}{1-\mu} &\to R_h - R_l = \frac{R_h - \mathbb{E}_j v_j}{1-\mu} \end{aligned}$$

which then it implies that

$$\begin{aligned} Cov(R_j, v_j) &= \mu R_h v_h + (1-\mu)\left[v_l - \frac{\mu}{1-\mu}(R_h - v_h)\right] v_l - (\mathbb{E}v_j)^2 \\ &= \mu R_h (v_h - v_l) + v_l \mathbb{E}v_j - (\mathbb{E}v_j)^2 \\ &= \mu(v_h - v_l)(R_h - \mathbb{E}v_j) \\ &= \mu(1-\mu)(v_h - v_l)(R_h - R_l) \end{aligned}$$

Finally, we realize that $r_h - r_l = (1-p)(R_h - R_l) + p(v_h - v_l)$ and that from Bayesian updating, $p = \frac{R_h - R_l}{2m}$. The above expressions allow us to write (44) as an equation in p given by

$$\begin{aligned} 2mp &= (1-\delta)(v_h - v_l) + \delta[(1-p)2mp + (v_h - v_l)p] \\ &\quad - \delta\pi\alpha_c[(1-p)2mp + (v_h - v_l)p] \\ &\quad - \delta\pi\alpha_m \frac{\sqrt{e^2 - 4\mu(1-\mu)(v_h - v_l)[2mp(1-p) + p(v_h - v_l)]} + e}{2e}[(1-p)2mp + (v_h - v_l)p] \end{aligned}$$

Since $\alpha_c + \alpha_m = 1$, we can write the above as

$$\begin{aligned} 0 =& -2mp + (1-\delta)(\nu_h - \nu_l) + \delta\left(1 - \frac{\pi}{2}\right)[(1-p)2mp + (\nu_h - \nu_l)p] \\ & - \frac{\delta\pi\alpha_m}{2}\sqrt{1 - \frac{4}{e^2}\mu(1-\mu)(\nu_h-\nu_l)[(1-p)2mp + (\nu_h-\nu_l)p]}[(1-p)2mp + (\nu_h-\nu_l)p] \end{aligned}$$

which is the same as (42). We refer to the right hand side of the above as $Z(p,\mu;\Xi)$. We show that $Z(p,\mu;\Xi)$ has a unique solution. To see this note that

$$\begin{aligned} Z(0,\mu;\Xi) &= (1-\delta)(\nu_h - \nu_l) > 0 \\ Z\left(\frac{\nu_h - \nu_l}{2m}, \mu; \Xi\right) &= -(\nu_h - \nu_l) + (1-\delta)(\nu_h-\nu_l) + \delta\left(1-\frac{\pi}{2}\right)(\nu_h - \nu_l) \\ & - \frac{\delta\pi\alpha_m}{2}\sqrt{1 - \frac{4}{e^2}\mu(1-\mu)(\nu_h-\nu_l)^2}(\nu_h - \nu_l) \\ &= -\frac{\pi}{2}\delta(\nu_h-\nu_l)\left[1 + \alpha_m\sqrt{1-\frac{4}{e^2}\mu(1-\mu)(\nu_h-\nu_l)^2}\right] < 0 \end{aligned}$$

This implies that there exists a $p^* \in \left[0, \frac{\nu_h - \nu_l}{2m}\right]$ that solves our equation. In addition, if we define $g(p) = p(2m(1-p) + \nu_h - \nu_l)$, then

$$\begin{aligned} Z_p =& -2m + \delta\left(1-\frac{\pi}{2}\right)g'(p) \\ & - \frac{\delta\pi\alpha_c}{2}\sqrt{1-\frac{4}{e^2}\mu(1-\mu)(\nu_h-\nu_l)g(p)}g'(p) \\ & + \frac{\delta\pi\alpha_c}{4}\frac{g(p)\frac{4}{e^2}\mu(1-\mu)(\nu_h-\nu_l)g'(p)}{\sqrt{1-\frac{4}{e^2}\mu(1-\mu)(\nu_h-\nu_l)g(p)}} \\ =& -2m + \delta\left(1-\frac{\pi}{2}\right)g'(p) \\ & - \delta\pi\alpha_c\frac{2g'(p)\left[1-\frac{4}{e^2}\mu(1-\mu)(\nu_h-\nu_l)g(p)\right] - g(p)\frac{4}{e^2}\mu(1-\mu)(\nu_h-\nu_l)g'(p)}{4\sqrt{1-\frac{4}{e^2}\mu(1-\mu)(\nu_h-\nu_l)g(p)}} \\ =& -2m + \delta\left(1-\frac{\pi}{2}\right)g'(p) \\ & - \delta\pi\alpha_c\frac{g'(p)\left[2-\frac{12}{e^2}\mu(1-\mu)(\nu_h-\nu_l)g(p)\right]}{4\sqrt{1-\frac{4}{e^2}\mu(1-\mu)(\nu_h-\nu_l)g(p)}} \end{aligned} \tag{45}$$

Note that

$$g'(p) = 2m(1-2p) + \nu_h - \nu_l$$

by the Assumption (2)

$$\begin{aligned} (\nu_h - \nu_l + 2m)\delta\left(1-\frac{\pi}{2}\right) - 2m < 0 \rightarrow g'(0)\delta\left(1-\frac{\pi}{2}\right) - 2m < 0 \\ \forall p \in \left[0, \frac{\nu_h-\nu_l}{2w}\right]\delta\left(1-\frac{\pi}{2}\right)g'(p) - 2m < 0 \end{aligned}$$

where the second inequality follows since $g'(p)$ is decreasing in p. In addition,

$$\frac{12\mu(1-\mu)}{e^2}(\nu_h-\nu_l)\,g(p) \leqslant \frac{3}{e^2}(\nu_h-\nu_l)^2 < 2$$

Finally, $\forall p \in \left[0, \frac{\nu_h-\nu_l}{2m}\right], g'(p) > 0$ given (2). This implies that the last expression in 45 is negative and therefore, $Z_p(p,\mu;\Xi) < 0$ for all $p \in \left[0, \frac{\nu_h-\nu_l}{2m}\right]$. Hence, $Z(p^*,\mu;\Xi) = 0$ has a unique solution in $\left[0, \frac{\nu_h-\nu_l}{2m}\right]$. Note that since $R_j(\mu)$'s are increasing in μ, $R_h(\mu) \leqslant \nu_h$ and $R_l(\mu) \geqslant \nu_l$. This implies that $R_h(\mu) - R_l(\mu) \leqslant \nu_h - \nu_l$ which then implies that $p \leqslant \frac{\nu_h-\nu_l}{2m}$. ■

A.2 Proof of Lemma 1

Proof. Recall the equation that describes p:

$$\begin{aligned}0 =& -2mp + (1-\delta)(\nu_h-\nu_l) + \delta\left(1-\frac{\pi}{2}\right)[(1-p)\,2mp + (\nu_h-\nu_l)\,p] \\ &- \frac{\delta\pi\alpha_m}{2}\sqrt{1-\frac{4}{e^2}\mu(1-\mu)(\nu_h-\nu_l)[(1-p)\,2mp+(\nu_h-\nu_l)\,p]}\,[(1-p)\,2mp+(\nu_h-\nu_l)\,p]\end{aligned}$$

The above function, $Z(p,\mu;\Xi)$, depends on μ only through $\mu(1-\mu)$. Moreover, $Z(p,\mu;\Xi)$ is increasing in $\mu(1-\mu)$. Therefore, $Z_\mu > 0$ when $\mu < 1/2$ and $Z_\mu < 0$ when $\mu > 1/2$. Note that we have

$$Z_p\,dp^* + Z_\mu\,d\mu = 0 \rightarrow \frac{dp^*}{d\mu} = -\frac{Z_p}{Z_\mu}$$

From proof of proposition 1 we know that $Z_p < 0$. Therefore, $\frac{dp^*}{d\mu} > 0$ when $\mu < 1/2$ and vice versa.

Additionally, bid-ask spreads are given by

$$A^c - B^c = e - \sqrt{e^2 - 4\mu(1-\mu)(\nu_h-\nu_l)[(1-p^*)\,2mp^* + (\nu_h-\nu_l)\,p^*]}$$

The above is increasing in $\mu(1-\mu)$ and p^*. Since both of these are hump-shaped in μ with maximum at $\mu = 1/2$, so is the bid-ask spread.

Finally, recall that

$$R_h - R_l = 2mp^*$$

and hence $R_h - R_l$ is also hump-shaped in μ and maxed at $\mu = 1/2$. This concludes the proof. ■

A.3 Proof of Proposition 2

Proof. We have

$$\begin{aligned}Z_\pi =& -\frac{\delta}{2}[(1-p)\,2mp + (\nu_h-\nu_l)] \\ &- \frac{\delta}{2}\sqrt{1-\frac{4}{e^2}\mu(1-\mu)(\nu_h-\nu_l)[(1-p)\,2mp+(\nu_h-\nu_l)\,p]}\,[(1-p)\,2mp+(\nu_h-\nu_l)\,p]\end{aligned}$$

and therefore, $Z_\pi < 0$. Hence,

$$\frac{dp^*}{d\pi} = -\frac{Z_p}{Z_\pi} < 0.$$

■

A.4 Proof of Proposition 3

Proof. From the text we have

$$\begin{aligned}
\frac{\partial}{\partial \pi}\sigma_{t,j} =& -t\alpha_c\left(1-p^*\right)^{t-1}\frac{\partial p^*}{\partial \pi}\left(1-\sqrt{1-\frac{4}{e^2}\left(\nu_h-\nu_l\right)\mu\left(1-\mu\right)g\left(p^*\right)}\right)\\
&+\alpha_c\left(1-p^*\right)^t\frac{\frac{4}{e^2}\left(\nu_h-\nu_l\right)\mu\left(1-\mu\right)g'\left(p^*\right)\frac{\partial p^*}{\partial \pi}}{\sqrt{1-\frac{4}{e^2}\left(\nu_h-\nu_l\right)\mu\left(1-\mu\right)g\left(p^*\right)}}\\
=&\alpha_c\left(1-p^*\right)^t\frac{\partial p^*}{\partial \pi}\left[\frac{\frac{4}{e^2}\left(\nu_h-\nu_l\right)\mu\left(1-\mu\right)g'\left(p^*\right)}{\sqrt{1-\frac{4}{e^2}\left(\nu_h-\nu_l\right)\mu\left(1-\mu\right)g\left(p^*\right)}}-\frac{t}{1-p^*}\left(1-\sqrt{1-\frac{4}{e^2}\left(\nu_h-\nu_l\right)\mu\left(1-\mu\right)g\left(p^*\right)}\right)\right]
\end{aligned}$$

In the above, $\frac{\partial p^*}{\partial \pi} < 0$ and for large enough t the expression in the brackets is negative. Hence, when t is large enough $\frac{\partial}{\partial \pi}\sigma_{t,j} > 0$. ■

A.5 Proof of Proposition 4

Proof. Note that the full information model is described by the following relationships

$$\begin{aligned}
\delta R &= \delta c + \pi\left[\int_{B-R}^{A-R}\varepsilon dG\left(\varepsilon\right) + \left(B-R\right)G\left(B-R\right) + \left(A-R\right)\left(1-G\left(A-R\right)\right)\right]\\
A &= \nu + \frac{1-G\left(A-R\right)}{g\left(A-R\right)}\\
B &= \nu - \frac{G\left(B-R\right)}{g\left(B-R\right)}
\end{aligned}$$

where

$$\xi_a\left(A-R\right) = \frac{1-G\left(A-R\right)}{g\left(A-R\right)}, \xi_b\left(B-R\right) = \frac{G\left(B-R\right)}{g\left(B-R\right)}$$

are the semi-elasticities of demand (supply) with respect to prices and are therefore equal to the markups charged by the dealer.

Note. We have

$$\begin{aligned}
\xi_a''\left(\varepsilon\right) = \frac{d^2}{d\varepsilon^2}\frac{1-G\left(\varepsilon\right)}{g\left(\varepsilon\right)} &= \frac{d}{d\varepsilon}\left(1-\frac{\left(1-G\left(\varepsilon\right)\right)g'\left(\varepsilon\right)}{g\left(\varepsilon\right)^2}\right)\\
&= -\frac{d}{d\varepsilon}\frac{\left(1-G\left(\varepsilon\right)\right)g'\left(\varepsilon\right)}{g\left(\varepsilon\right)^2}\\
&= \frac{g'}{g} - \frac{\left(1-G\right)g''}{g^2} + 2\frac{\left(1-G\right)\left(g'\right)^2}{g^3}
\end{aligned}$$

Note that

$$G(-\varepsilon) = 1 - G(\varepsilon)$$
$$g(-\varepsilon) = g(\varepsilon)$$
$$-g'(-\varepsilon) = g'(\varepsilon)$$
$$g''(-\varepsilon) = g''(\varepsilon)$$

Thus

$$\begin{aligned}\xi_a''(\varepsilon) &= -\frac{g'(-\varepsilon)}{g(-\varepsilon)} - \frac{G(-\varepsilon)\,g''(-\varepsilon)}{g(-\varepsilon)^2} + 2\frac{G(-\varepsilon)\left(g'(-\varepsilon)\right)^2}{g(-\varepsilon)^3} \\ &= \xi_b''(-\varepsilon) > 0\end{aligned}$$

We show the claim first by assuming that $v > c$ - the case where $v < c$ follows a similar argument is thus ommited. The following Lemma proves useful in showing the main result:

Lemma 1. *Suppose Condition 1 holds and that $v > c$, then we have*

$$R > c, A - R > R - B.$$

Proof. Suppose that $R \leqslant c$, then

$$\int_{B-R}^{A-R} \varepsilon dG(\varepsilon) + (B - R)\,G(B - R) + (A - R)\,(1 - G(A - R)) \leqslant 0$$

We can write the above as function $\Omega(A - R, B - R)$. Note that

$$\Omega(x, -x) = 0$$
$$\Omega_y(x, y) = G(y) > 0$$
$$\Omega_x(x, y) = 1 - G(x) > 0$$

This implies that $\Omega(A - R, B - R) \leqslant 0$ must lead to $A - R \leqslant R - B$, i.e., there is more buyers than sellers!

Hence by Condition 1

$$\frac{1 - G(A - R)}{g(A - R)} \geqslant \frac{1 - G(R - B)}{g(R - B)} = \frac{G(B - R)}{g(B - R)}$$

Thus

$$A + B = 2v + \frac{1 - G(A - R)}{g(A - R)} - \frac{G(B - R)}{g(B - R)} \geqslant 2v$$

Since

$$A - R \leqslant R - B \to A + B \leqslant 2R \to 2v \leqslant 2R \leqslant 2c$$

which is a contradiction. Thus the claim follows. ■

The above lemma implies that when there is gains from trade, the trader's capture some of these gains and that there are more sellers than buyers.

To prove the main claim, we first note that

$$\begin{aligned}\frac{dA}{dR} &= \frac{-\xi_a'(A-R)}{1-\xi_a'(A-R)} \in [0,1] \\ \frac{dB}{dR} &= \frac{\xi_b'(B-R)}{1+\xi_b'(B-R)} \in [0,1]\end{aligned}$$

Thus

$$\frac{d(A-B)}{dR} = \frac{-\xi_a'(A-R)}{1-\xi_a'(A-R)} - \frac{\xi_b'(B-R)}{1+\xi_b'(B-R)}$$

where

$$\begin{aligned}-\xi_a'(\varepsilon) &= 1 + \frac{(1-G(\varepsilon))\,g'(\varepsilon)}{g(\varepsilon)^2} \\ \xi_b'(\varepsilon) &= 1 - \frac{G(\varepsilon)\,g'(\varepsilon)}{g(\varepsilon)^2} \\ \xi_b'(-\varepsilon) &= 1 - \frac{G(-\varepsilon)\,g'(-\varepsilon)}{g(-\varepsilon)^2} = 1 + \frac{(1-G(\varepsilon))\,g'(\varepsilon)}{g(\varepsilon)^2} \\ &= -\xi_a'(\varepsilon)\end{aligned}$$

From the previous lemma we have $A-R > R-B$ and therefore, from assumption 1, we must have

$$-\xi_a'(A-R) < -\xi_a'(R-B) = \xi_b'(B-R)$$

We, therefore, have

$$\begin{aligned}\frac{d(A-B)}{dR} &= \frac{-\xi_a'(A-R)}{1-\xi_a'(A-R)} - \frac{\xi_b'(B-R)}{1+\xi_b'(B-R)} \\ &= \frac{-\xi_a'(A-R) - \xi_b'(B-R)}{(1-\xi_a'(A-R))\left(1+\xi_b'(B-R)\right)} < 0\end{aligned}$$

It is thus left to show that $\frac{dR}{d\pi} > 0$. We have

$$\begin{aligned}\delta dR &= \pi\left[\Omega_x(dA-dR) + \Omega_y(dB-dR)\right] + d\pi\Omega(A-R,B-R) \\ \delta dR &= \pi\left[(1-G(A-R))(dA-dR) + G(B-R)(dB-dR)\right] + d\pi\Omega(A-R,B-R) \\ &= \pi\left[-(1-G(A-R))\frac{dR}{1-\xi_a'} - G(B-R)\frac{dR}{1+\xi_b'}\right] + d\pi\Omega(A-R,B-R) \\ \frac{dR}{d\pi} &= \frac{\Omega(A-R,B-R)}{\delta + \pi\left[\frac{(1-G(A-R))}{1-\xi_a'} + \frac{G(B-R)}{1+\xi_b'}\right]} > 0\end{aligned}$$

This concludes the proof. ■

B Additional Results for the Special Case

B.1 Dealers have no incentive to experiment with prices

Here we establish that dealers have no incentive to set statically sub-optimal prices that might reveal to them the state of the world. To see why, note that the set of bids that could potentially reveal the state of the world lie in the set $\Xi_1 = (R_l(\mu) - m - e, R_h(\mu) - m - e)$ or $\Xi_2 = (R_l(\mu) + m + e, R_h(\mu) + m + e)$.[20] For any bid in the first interval, observing a trader with an asset accept the offer would reveal that the state is l. For any bid in the second interval, observing a trader with an asset reject the offer would reveal that the state is h.

Now, suppose the dealer sets a bid $\hat{B} \in \Xi_2$; the argument for a bid in Ξ_1 is symmetric. An optimal offer would never generate zero trades in both states of the world. Hence, after observing the volume of sells, there are three possibilities for the corresponding signal S:

1. $S \in \Sigma_l(\mu) \equiv [-m + R_l(0), -m + R_h(\mu))$. In this case, the state of the world was revealed anyway, so there are no benefits to experimentation.

2. $S \in \Sigma_h(\mu) \equiv (m + R_l(\mu), m + R_h(1)]$. Again, in this case the state of the world was revealed anyway, so there are no benefits to experimentation.

3. $S \in \Sigma_b(\mu) \equiv [-m + R_h(\mu), m + R_l(\mu)]$. In this case, all traders accept the offer $\hat{B}$, and the state of the world is not revealed to the dealer.

B.2 Valuations of traders and dealers are equal in expectation

Here we establish that (34) which states that in expectation - given dealers' information - valuation of dealers and traders are equal.

We first show this when $\mu = 0, 1$, i.e., when dealers are fully informed about the j. Note that when $\mu = 1$,

$$\begin{aligned}
B^c(1) &= \frac{r_h(1) + v_h - e}{2} + \frac{1}{2}\sqrt{(e + v_h - r_h(1))^2} = v_h \\
A^c(1) &= \frac{r_h(1) + v_h + e}{2} - \frac{1}{2}\sqrt{(e + v_h - r_h(1))^2} = v_h \\
B^m(1) &= \frac{r_h(1) + v_h - e}{2} \\
A^m(1) &= \frac{r_h(1) + v_h + e}{2}
\end{aligned}$$

where the above holds because with full information, $\mathrm{Cov}_j\left(r_j, v_j\right) = 0$.

[20] The argument for the ask price is symmetric.

Note that $r_h(1) = R_h(1)$ and value functions can be written as

$$r_h(1) = (1-\delta)\nu_h + \delta r_h(1) + \delta\pi \sum_{k=c,m} \alpha_k \frac{(B^k - A^k + 2e)}{2e} \left(\frac{A^k + B^k}{2} - r_h(1) \right)$$

$$(1-\delta) r_h(1) = (1-\delta)\nu_h + \delta\pi\alpha_c(\nu_h - r_h(1)) + \delta\pi\alpha_m \frac{1}{2}\frac{\nu_h - r_h(1)}{2}$$

Obviously the unique solution of the above equation is $r_h(1) = \nu_h$. Similarly, we can show that $r_l(0) = \nu_l$. This implies that

$$r_j(\mu) = p(\mu) R_j(\mu) + (1 - p(\mu))\nu_j$$

As a result

$$\mathbb{E}_j r_j(\mu) = p(\mu)\mathbb{E}_j R_j(\mu) + (1-p(\mu))\mathbb{E}_j \nu_j$$

Suppose to the contrary that $\mathbb{E}_j r_j > \mathbb{E}_j \nu_j$. Given the above expression, this implies that $\mathbb{E}_j R_j(\mu) > \mathbb{E}_j \nu_j$. Note that the Bellman equations for traders net-option values are given by

$$\begin{aligned} R_j(\mu) = &(1-\delta)\nu_j + \delta r_j(\mu) \\ &+ \delta\pi \sum_{k=c,m} \alpha_k \frac{(B^k - A^k + 2e)}{2e} \left(\frac{A^k + B^k}{2} - r_j(\mu) \right) \end{aligned}$$

Using the formulas for prices provided in the text, we have

$$A^c(\mu) = A^m(\mu) - \Psi(\mu)$$

$$B^c(\mu) = B^m(\mu) + \Phi(\mu)$$

with

$$\Phi(\mu) = \frac{1}{2}\sqrt{\left(\mathbb{E}_j(\nu_j - r_j) + e\right)^2 - 4Cov_j(r_j, \nu_j)}$$
$$\Psi(\mu) = \frac{1}{2}\sqrt{\left(\mathbb{E}_j(r_j - \nu_j) + e\right)^2 - 4Cov_j(r_j, \nu_j)}$$

Note that when $\mathbb{E}_j r_j > \mathbb{E}_j \nu_j$, then $\Psi(\mu) > \Phi(\mu)$

$$r_j(\mu) = (1-p(\mu)) R_j(\mu) + p(\mu)\nu_j$$

$$\mathbb{E}_j r_j = (1-p)\mathbb{E}_j R_j + p\mathbb{E}_j \nu_j$$

Using the above, we can write the Bellman equations as

$$\begin{aligned} R_j(\mu) = &(1-\delta)\nu_j + \delta r_j(\mu) \\ &+ \delta\pi\alpha_m \frac{1}{2}\frac{\nu_j - r_j(\mu)}{2} \\ &+ \delta\pi\alpha_c \frac{\Phi(\mu) + \Psi(\mu) + e}{2} \left(\frac{\nu_j - r_j(\mu) + \Phi(\mu) - \Psi(\mu)}{2} \right) \end{aligned}$$

Taking expectation with respect to j in the above equation, we have

$$\begin{aligned}\mathbb{E}_j R_j =& (1-\delta)\,\mathbb{E}_j \nu_j + \delta \mathbb{E}_j r_j \\ &+ \delta\pi\alpha_m \frac{\mathbb{E}_j \nu_j - \mathbb{E}_j r_j}{4} \\ &+ \delta\pi\alpha_c \frac{\Phi(\mu) + \Psi(\mu) + e}{2} \frac{\mathbb{E}_j \nu_j - \mathbb{E}_j r_j + \Phi(\mu) - \Psi(\mu)}{2}\end{aligned}$$

Since $\Phi(\mu) < \Psi(\mu)$ and $\mathbb{E}_j \nu_j < \mathbb{E}_j r_j$, the last two terms in the right hand side of the above expression are negative. Hence,

$$\mathbb{E}_j R_j < (1-\delta)\,\mathbb{E}_j \nu_j + \delta \mathbb{E}_j r_j = (1-\delta+\delta(1-p(\mu)))\,\mathbb{E}_j \nu_j + \delta p(\mu)\,\mathbb{E}_j R_j$$

Therefore,

$$(1-\delta p(\mu))\,\mathbb{E}_j R_j(\mu) < (1-\delta p(\mu))\,\mathbb{E}_j \nu_j$$

The above inequality is in constradiction to the initial assumption where $\mathbb{E}_j r_j > \mathbb{E}_j \nu_j$. This is the required contradiction that proves our intial assumption wrong. Similarly, we can show that $\mathbb{E}_j r_j < \mathbb{E}_j \nu_j$ cannot hold. Hence, $\mathbb{E}_j r_j = \mathbb{E}_j \nu_j$.

B.3 Pricing under Monopoly and Competition

B.3.1 Monopoly Pricing

Suppose that for all realizations of ω, $B^k - R_j(\mu) - \omega \in [-e, e]$. This means that if updated beliefs are given byμ', then the probability of a sale by a trader is given by

$$\frac{B^k - \omega - R_j(\mu') + e}{2e}$$

This implies that for a realization of ω, the probability of sale by a trader is given by

$$\frac{B^k - \omega - r_j(\mu) + e}{2e}$$

where $r_j(\mu) = (1-p(\mu))\,R_j(\mu) + p(\mu)\,R_j(\mathbf{1}[j=h])$. Therefore profits of a dealer who buys is given by

$$\sum_j \mu_j \int_{-m}^{m} \frac{B - \omega - r_j + e}{2e} \left(\nu_j - B\right) \frac{d\omega}{2w}$$

The derivative of the above objective with respect to B is given by

$$\frac{\mathbb{E}_j \nu_j - B}{2e} - \mathbb{E}_j \frac{B - r_j + e}{2e} = 0 \to B^m = \frac{\mathbb{E}_j \nu_j + \mathbb{E}_j r_j - e}{2}$$

Similarly, if we assume that $R_j(\mu) + \omega - A^k \in [-e, e]$, then the profits for a selling monopolist dealer is given by

$$\sum_j \mu_j \int \frac{e - A + \omega + r_j}{2e} \left(A - \nu_j\right) \frac{d\omega}{2w}$$

The first order condition with respect to A is given by

$$-\frac{A-\mathbb{E}_j v_j}{2e}+\mathbb{E}_j\frac{e-A+r_j}{2e}=0\to A^m=\frac{\mathbb{E}_j v_j+\mathbb{E}_j r_j+e}{2}$$

B.3.2 Pricing under Competition

In equilibrium and under competition, profits of a buying trader must be zero and we must have that

$$-B^2+\left(\mathbb{E}_j r_j+\mathbb{E}_j v_j-e\right)B-\left(\mathbb{E}_j v_j\left(r_j-e\right)\right)=0$$

The above equation has two solutions. By Bertrand competition, equilibrium must be the higher solution. To see this, suppose that $B_1 < B_2$ are the roots of this equation. Suppsoe, further, that equilibrium bid is B_1. Then one of the dealers can deviate to $B_2-\varepsilon$ for a small but positive value of ε and achieve positive profits. The reason is that if $B > B_2$ profits must be negative - since ultimately they are negative and the equation has only two roots. Therefore, just below B_2, they must be positive. Hence, this is a profitable deviation. Therefore, equilibrium bid must be given by

$$\begin{aligned}B^c&=\frac{\mathbb{E}_j r_j+\mathbb{E}_j v_j-e+\sqrt{\left(\mathbb{E}_j r_j+\mathbb{E}_j v_j-e\right)^2-4\mathbb{E}_j v_j\left(r_j-e\right)}}{2}\\&=\frac{\mathbb{E}_j r_j+\mathbb{E}_j v_j-e+\sqrt{\left(\mathbb{E}_j v_j\right)^2+\left(\mathbb{E}_j r_j-e\right)^2+2\mathbb{E}v_j\mathbb{E}\left(r_j-e\right)-4\mathbb{E}_j v_j\left(r_j-e\right)}}{2}\end{aligned}$$

The discriminant in the above can be written as

$$\begin{aligned}\left(\mathbb{E}_j v_j\right)^2+\left(\mathbb{E}_j r_j\right)^2-2e\mathbb{E}_j r_j+e^2+2\mathbb{E}_j v_j\mathbb{E}_j r_j-2e\mathbb{E}_j v_j+4e\mathbb{E}_j v_j-4\mathbb{E}_j v_j r_j&=\\\left(\mathbb{E}_j v_j\right)^2+\left(\mathbb{E}_j r_j\right)^2-2e\mathbb{E}_j r_j+e^2+2\mathbb{E}_j v_j\mathbb{E}_j r_j+2e\mathbb{E}_j v_j-4\mathbb{E}_j v_j r_j+4\mathbb{E}_j v_j\mathbb{E}_j r_j-4\mathbb{E}_j v_j\mathbb{E}_j r_j&=\\\left(\mathbb{E}_j v_j\right)^2+\left(\mathbb{E}_j r_j\right)^2-2e\mathbb{E}_j r_j+e^2-2\mathbb{E}_j v_j\mathbb{E}_j r_j+2e\mathbb{E}_j v_j-4\mathrm{Cov}_j\left(v_j,r_j\right)&=\\\left(\mathbb{E}_j v_j-\mathbb{E}_j r_j+e\right)^2-4\mathrm{Cov}\left(v_j,r_j\right)&\end{aligned}$$

Therefore, we can write the above a

$$B^c=\frac{\mathbb{E}_j\left(v_j+r_j\right)-e+\sqrt{\left(\mathbb{E}_j\left(v_j-r_j\right)+e\right)^2-4\mathrm{Cov}_j\left(v_j,r_j\right)}}{2}$$

For the above to be a valid solution, we must have that

$$e+\mathbb{E}_j\left(v_j-r_j\right)\geqslant 2\sqrt{\mathrm{Cov}\left(v_j,r_j\right)}$$

In other words, e must be large enough. We will verify that these conditions are indeed satisfied under Assumptions 1-3.

As for the ask, we have that profits are given by

$$\sum_j \mu_j \int \frac{e - A + \omega + r_j}{2e} \left(A - \nu_j\right) \frac{d\omega}{2w}$$

Therefore, the zero-profit condition must be given by

$$-A^2 + \left(\mathbb{E}_j r_j + e + \mathbb{E}_j \nu_j\right) A - \mathbb{E}_j \nu_j \left(e + r_j\right) = 0$$

Then, a similar Bertrand type argument as above implies that the equilibrium ask price is the lower of the two roots of the above. Thus, we have

$$\begin{aligned}
A^c &= \frac{\mathbb{E}\left(r_j + \nu_j\right) + e - \sqrt{\left(\mathbb{E}\left(r_j + \nu_j\right) + e\right)^2 - 4\mathbb{E}\nu_j \left(e + r_j\right)}}{2} \\
&= \frac{\mathbb{E}\left(r_j + \nu_j\right) + e - \sqrt{\left(\mathbb{E}r_j\right)^2 + \left(\mathbb{E}\nu_j\right)^2 + e^2 + 2e\mathbb{E}\nu_j + 2e\mathbb{E}r_j + 2\mathbb{E}r_j\mathbb{E}\nu_j - 4\mathbb{E}\nu_j \left(e + r_j\right)}}{2} \\
&= \frac{\mathbb{E}\left(r_j + \nu_j\right) + e - \sqrt{\left(\mathbb{E}r_j\right)^2 + \left(\mathbb{E}\nu_j\right)^2 + e^2 - 2e\mathbb{E}\nu_j + 2e\mathbb{E}r_j - 2\mathbb{E}r_j\mathbb{E}\nu_j - 4\mathrm{Cov}\left(\nu_j, r_j\right)}}{2} \\
&= \frac{\mathbb{E}\left(r_j + \nu_j\right) + e - \sqrt{\left(\mathbb{E}\left(r_j - \nu_j\right) + e\right)^2 - 4\mathrm{Cov}\left(\nu_j, r_j\right)}}{2}
\end{aligned}$$

For this to exist, we must have that

$$e + \mathbb{E}_j \left(r_j - \nu_j\right) \geqslant 2\sqrt{\mathrm{Cov}\left(\nu_j, r_j\right)}$$