

Market-making with Search and Information Frictions[‡]

Benjamin Lester
Federal Reserve Bank of Philadelphia

Venky Venkateswaran
NYU Stern School of Business

Ali Shourideh
Carnegie Mellon University

Ariel Zetlin-Jones
Carnegie Mellon University

May 17, 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 have reduced trading frictions in over-the-counter markets. Our main result is that reducing trading frictions can lead to less liquidity, as measured by bid-ask spreads. The key insight is that more frequent trading—or more competition among dealers—makes traders' behavior less dependent on asset quality. As a result, dealers learn about asset quality more slowly and set wider bid-ask spreads to compensate for this increase in uncertainty.

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

*We thank Marzena Rostek, Giuseppe Moscarini, Nicolae Garleanu, Edouard Schaal, Dimitri Vayanos, Mike Golosov, Burton Hollifield, Aleh Tsyvinski, Pierre-Olivier Weill, and participants at various conferences and seminar audiences at Berkeley-Haas, Boston College, Carnegie Mellon, CREI, European University Institute, the New York Fed, Princeton, and the University of Wisconsin for useful comments and discussions. We are especially grateful to Michael Choi for an insightful discussion.

[†]The views expressed here are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

1 Introduction

Many assets trade in decentralized, dealer-intermediated, over-the-counter (OTC) markets, such as corporate bonds, municipal bonds, and various types of derivatives (such as interest rate and credit default swaps). These markets have undergone significant changes in recent years as a result of both technological innovations and regulatory initiatives. One of the primary consequences of these changes has been a reduction in *trading frictions*: investors have gained the opportunity to trade more quickly with a wider set of dealers. For example, in the corporate bond market, there has been a sizable shift away from voice-based trading—where investors contact dealers sequentially via telephone to get a quote—to electronic trading platforms—where investors can instantaneously submit a request for quotes that is received by many dealers.¹ Similar developments have taken place in the markets for swaps, where regulatory mandates have forced trading activity from decentralized, OTC markets onto more centralized exchanges.² As these changes continue—and, in some cases, accelerate—a natural question arises: What is the effect of reducing trading frictions on market liquidity?

The literature that uses search and matching models of trade to study OTC markets offers a stark prediction: enabling investors to contact dealers more easily erodes the dealers' market power, causing bid-ask spreads to fall and thus market liquidity to rise.³ However, this literature typically abstracts from a second, canonical source of illiquidity: *asymmetric information*. Since the seminal work of [Glosten and Milgrom \(1985\)](#), the contribution of information asymmetries to market liquidity has been studied extensively and documented empirically in a variety of financial markets.⁴ Recognizing the potentially important role played by asymmetric information raises additional questions: Do changes in trading frictions mitigate or exacerbate the effects of information frictions? Do the stark predictions that come from models with only trading frictions remain true in settings with both trading and information frictions?

The goal of this paper is to provide some answers to these questions. To do so, we develop a unified framework that incorporates both trading and information frictions. We characterize equilibrium prices,

¹For instance, according to The Securities Industry and Financial Markets Association, the fraction of total volume (for investment grade corporate bonds) traded on electronic exchanges increased from 8% in 2013 to 20% in 2015, as the total number of electronic platforms more than doubled. For a detailed description of the transition of several markets from dealer-based platforms to electronic platforms, and ever-increasing execution speeds, see Appendix A of [Pagnotta and Philippon \(2015\)](#).

²In 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. In addition to moving trading from an OTC market to a centralized setting, it also requires an investor's request for price quotes to be circulated to at least three dealers before the trade can be executed. Similar regulatory requirements have been implemented in European markets as a consequence of MiFID II.

³See, for example, the large body of work following [Duffie et al. \(2005\)](#).

⁴For example, several papers, including [Goldstein et al. \(2006\)](#) and [Edwards et al. \(2007\)](#), analyze the role of adverse selection in the corporate bond market. The effects of asymmetric information on liquidity and spreads have also been studied in the over-the-counter markets for futures ([Ma et al., 1992](#)), foreign exchange swaps ([Bollerslev and Melvin, 1994](#)), and municipal bonds ([Harris and Piwowar, 2006](#); [Brancaccio et al., 2017](#)), to name just a few.

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.

Our main result is that the stark predictions that arise from models with only trading frictions are no longer necessarily true when both trading and information frictions are present. In particular, we show that reducing trading frictions has two, opposing effects. On the one hand, as in traditional models, it generates more competition among dealers, leading to tighter spreads and more liquid markets. On the other hand, however, we show that it also slows down the process by which dealers learn about the quality of the assets being traded. Over time, this exacerbates the effects of asymmetric information and leads to *wider* bid-ask spreads. Hence, we show that a one-time decrease in trading frictions typically leads to an immediate fall in spreads, relative to the baseline, but wider spreads in the long run.

To understand the intuition behind our results, 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 she can't trade; the trader might be matched with a single dealer, in which case the dealer has some degree of market power; or the trader 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. At the end of each period, 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.

Our main result then follows from a sequence of intermediate results. First, we show that dealers learn more quickly when traders' behavior—which is summarized by their reservation values—is more distinct in different states of the world (i.e., under different asset qualities). Second, and most important, we show that reducing trading frictions implies that traders' reservation values are more similar across asset qualities. Lastly, we show that dealers set wider bid-ask spreads when they are more uncertain about asset

quality. Taken together, these results imply that reducing trading frictions slows down dealers' learning, and can actually reduce market liquidity (as measured by the bid-ask spread) in the long run.

But why do traders' reservation values become less dependent on asset quality as trading frictions become less severe? Reservation values can be decomposed into two pieces: one that depends on the fundamental value of the asset and another that depends on the option value of trading the asset in the future. When traders' contact dealers more easily, their reservation values depend less on the former and more on the latter. This is intuitive, since the option value of trading an asset obviously relies on the frequency with which a trader has the *opportunity* to trade it when his preferences dictate a desire to do so. Moreover, the option value of trading is larger when the asset quality is low than it is when the asset quality is high. This is because a trader is more likely to want to exercise the option of selling a low quality asset, relative to a high quality asset.⁵ Hence, reducing trading frictions puts less weight on the component that makes reservation values more distinct, and more weight on the component that diminishes the difference between reservation values in the high and low quality states of the world.

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, despite the complex nature of the (dynamic, non-stationary) equilibrium in this environment, we are able to identify a special case of our model that allows for a full analytical characterization. Exploiting the tractability of this special case reveals a number of novel interactions between the three key ingredients described above, and the resulting implications for observable outcomes. Finally, given these implications, our model offers a framework to interpret existing empirical results that are hard to rationalize in models with only one of these frictions, such as a positive relationship between spreads and trading speed (as documented in, e.g., [Hendershott and Moulton, 2011](#)) or a positive relationship between spreads and trading volume (as documented in, e.g., [Lin et al., 1995](#); [Chordia et al., 2001](#)). Moreover, if our model can help us understand the affects of trading frictions and information frictions in the past, it may also prove helpful for anticipating the effects of future changes to OTC markets, as a number of current regulatory proposals are aimed at either eliminating trading frictions or reducing information asymmetries in financial markets.

⁵By a similar logic, a trader is more likely to want to exercise the option to buy a high quality asset, relative to a low quality asset, which similarly reduces the option value of acquiring a high quality asset relative to a low quality asset.

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 flexible specification of the model, assign parameter values that are roughly consistent with those in the existing literature and/or the data, and solve the model numerically. This section is intended to serve two main purposes. First, it provides some assurance that the forces we identify using the special case of our model are not a byproduct of the specific assumptions we impose in Section 3, and are relevant in a plausible region of the parameter space. Second, it allows us to derive additional insights and results that cannot be derived analytically. In Section 5, we consider a stationary version of the model and discuss the welfare ramifications of reducing trading frictions. Interestingly, we show that a deterioration in market liquidity—as measured by a widening bid-ask spread—need not correspond to a decline in welfare. Section 6 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 bid-ask spreads in full information settings.⁶ Gehrig (1993), Spulber (1996), and, more recently, Lester et al. (2015) analyze pricing under asymmetric information about preferences, i.e., about traders' *private* values of holding the asset.⁷ In our paper, the traders are privately informed both about their preferences *and* about a common value component of the asset, which leads to adverse selection. Moreover, since the common component is an aggregate one, there is a role for learning over time by the uninformed market-makers.⁸

Second, our paper contributes to several strands of the voluminous literature that focuses on the effects of asymmetric information in settings without trading and/or search frictions. Perhaps the most obvious is the strand that focuses on the effects of asymmetric information on the bid-ask spread, following the seminal contributions of Copeland and Galai (1983), Glosten and Milgrom (1985), and Kyle

⁶A few papers in this literature also find that easing trading frictions can have counter-intuitive effects on liquidity, even under full information. For example, in Lagos and Rocheteau (2009), changing trading frictions affects the endogenous distribution of asset holdings, which leads to non-monotonic effects on spreads. In Afonso (2011), the presence of congestion externalities plays a crucial role. In our paper, the counter-intuitive effects stem from the interaction with asymmetric information and learning—channels which are not present in these papers.

⁷Another, more recent, example is Bethune et al. (2016).

⁸This 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 Koepl (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 involving other assets.

(1985). In addition, our focus on the informational content of endogenous market signals is shared by the strand of this literature that studies information aggregation in rational expectations equilibrium (REE) models, pioneered by [Grossman and Stiglitz \(1980\)](#) and [Hellwig \(1980\)](#). Our analysis highlights novel interactions between asymmetric information and search frictions, and shows how they lead to surprising and counter-intuitive implications for liquidity and prices.⁹

The combination of trading frictions, adverse selection, and learning 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. \(2014\)](#), [Lauermann and Wolinsky \(2016\)](#), and [Lauermann et al. \(2017\)](#). 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.¹⁰

Finally, by analyzing the effects of reducing trading frictions, our analysis also makes 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 learning over time in our framework.

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: a measure N of “traders” and a mass D of “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., they can take on arbitrarily large 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 $c_l < c_h$.¹¹ If the asset

⁹A related point appears in [Rostek and Wernetka \(2015\)](#), who show that increasing trader participation—and hence making markets larger—can lead to a decrease in liquidity. [Asriyan et al. \(2017\)](#) study information aggregation in a dynamic setting, and find that sellers’ incentives to delay trade could lead to failure of information aggregation, even with a large number of traders. In contrast, in our paper, trading frictions affect the extent to which trading decisions—and, therefore, the informational content of the endogenous signal—depend on the fundamental (common) value of the asset.

¹⁰A related literature also studies learning and information diffusion in network settings; see, e.g., [Babus and Kondor \(2016\)](#).

¹¹There 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 $t' > t$.

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 mean zero and full support over the real line.¹²

Dealers receive a payoff v_j when the asset matures, with $v_h > v_l$, 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 v_j and $-v_j$, 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 perfectly informed about the quality of 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 liquidity shocks are also privately observed by the traders.

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_n denote the probability that a trader meets $n \in \{0, 1, \dots\}$ dealers. As we describe below, any meeting with $n \geq 2$ dealers will have the same outcome. Hence, the trading frictions can be succinctly summarized by two sufficient 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 $n \geq 1$ dealers,

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

We define the probability of meeting with $n \geq 2$ dealers (a “competitive” meeting), conditional on meeting with at least one dealer, by $\alpha_c = 1 - \alpha_m$.¹³

¹²This assumption is made only for simplicity and is not essential for our results. In Section 3, we analyze a version where these distributions have bounded support.

¹³This convenient transformation allows us to separate 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).

After meetings occur, each dealer quotes a bid and an ask price, i.e., prices at which he'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.¹⁴ 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 the expected profits from trade are zero. We denote the prices quoted by a monopolist dealer 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 that the asset is of quality h at the beginning of trading at time t .

2.2 Traders' Optimal Behavior

Let $W_{j,t}^q$ 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 a trader who does not own the asset, we have

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

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_t^k and future payoffs $W_{j,t+1}^q$ —can be calculated using the information available to a trader at time t , which includes the true quality of the asset as well as the current beliefs of dealers. We describe in detail below how the trader uses this information to formulate beliefs.

In words, the first expression inside the brackets 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

¹⁴This is in contrast to the assumption in [Burdett and Judd \(1983\)](#) or, more recently, [Lester et al. \(2017\)](#), where offers are made under uncertainty about the number of other offers being received

of the asset at price A_t^k 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,

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

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 $\varepsilon_{i,t} \geq \bar{\varepsilon}_{j,t}^k$, while a trader who owns the asset will sell if $\varepsilon_{i,t} \leq \underline{\varepsilon}_{j,t}^k$, where these cutoffs satisfy

$$\begin{aligned} -A_t^k + \omega_t + \bar{\varepsilon}_{j,t}^k + W_{j,t+1}^1 &= W_{j,t+1}^0 \\ \omega_t + \underline{\varepsilon}_{j,t}^k + W_{j,t+1}^1 &= B_t^k + W_{j,t+1}^0. \end{aligned}$$

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

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

Then, the optimal behavior of traders is succinctly summarized by the cutoffs, for $k \in \{m, c\}$,

$$\underline{\varepsilon}_{j,t}^k = B_t^k - \omega_t - R_{j,t+1}, \quad (2)$$

$$\bar{\varepsilon}_{j,t}^k = A_t^k - \omega_t - R_{j,t+1}, \quad (3)$$

along with reservation values

$$R_{j,t} = (1 - \delta)c_j + \delta \mathbb{E}_{\omega, \varepsilon} [R_{j,t+1}] + \delta \pi \sum_{k=c,m} \alpha_k \Omega_{j,t}^k, \quad (4)$$

where the term

$$\begin{aligned} \Omega_{j,t} &= \mathbb{E}_{\omega, \varepsilon} \left[\max (B_t^k - R_{j,t+1} - \omega_t - \varepsilon_{i,t}, 0) - \max (-A_t^k + \omega_t + \varepsilon_{i,t} + R_{j,t+1}, 0) \right] \\ &= \mathbb{E}_{\omega} \left\{ \left[B_t^k G(\underline{\varepsilon}_{j,t}^k) + \int_{\underline{\varepsilon}_{j,t}^k}^{\bar{\varepsilon}_{j,t}^k} [\omega_t + \varepsilon_{i,t} + R_{j,t+1}] dG(\varepsilon_{i,t}) + A_t^k [1 - G(\bar{\varepsilon}_{j,t}^k)] \right] - R_{j,t+1} \right\} \end{aligned} \quad (5)$$

is the *net option value* of holding an asset of quality j in a meeting of type k . In words, by acquiring a unit of the asset, a trader gains the option of selling it at a later date but, since holdings are restricted to $\{0, 1\}$,

she gives up the option of buying it later at a different (potentially better) price. The trader's reservation value, $R_{j,t}$, includes the difference between the expected value of these two options, multiplied by the probability of a trading opportunity, $\delta\pi$. The representation in (5) is obtained by using the cutoff rules described above. The expectation 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_t^q denote the measure of traders who have asset holdings $q \in \{0, 1\}$ at time t . When the asset is of quality $j \in \{l, h\}$, we have

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

Naturally, the measure of traders that own an asset in period $t + 1$ is equal to the measures of traders that owned an asset in period t and did not sell, plus the measure of traders that did not own an asset but chose to buy. The law of motion for the measure of traders that don't own an asset follows the same logic. We assume that the initial distribution of owners and non-owners, $(N_0^0, N_0^1 = N - N_0^0)$, is common knowledge. Hence, as we describe below, dealers will know (N_t^0, N_t^1) 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 he 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

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

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:

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

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

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

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

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

$$\mu_t (1 - \mu_t) (v_h - v_l) \frac{\mathbb{E}_\omega [g(\bar{\varepsilon}_{h,t}^m) - g(\bar{\varepsilon}_{l,t}^m)]}{\mathbb{E}_{j,\omega} [g(\bar{\varepsilon}_{j,t}^m)]} = \frac{\text{Cov}_j (\mathbb{E}_\omega [g(\bar{\varepsilon}_{j,t}^m)], v_j)}{\mathbb{E}_{j,\omega} [g(\bar{\varepsilon}_{j,t}^m)]}.$$

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.¹⁵ 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 $v_h = v_l$.

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:

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

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

¹⁵Note that this asymmetric information component can be positive or negative.

value of the asset conditional on a trader selling (buying) at that price. Formally, this zero profit condition can be written

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

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

Re-arranging yields

$$A_t^c = \mu_t v_h + (1 - \mu_t) v_l + \mu_t (1 - \mu_t) (v_h - v_l) \frac{\mathbb{E}_\omega \left[G(\bar{\varepsilon}_{l,t}^c) - G(\bar{\varepsilon}_{h,t}^c) \right]}{\mathbb{E}_{j,\omega} \left[1 - G(\bar{\varepsilon}_{j,t}^c) \right]} \quad (10)$$

$$B_t^c = \mu_t v_h + (1 - \mu_t) v_l - \mu_t (1 - \mu_t) (v_h - v_l) \frac{\mathbb{E}_\omega \left[G(\underline{\varepsilon}_{l,t}^c) - G(\underline{\varepsilon}_{h,t}^c) \right]}{\mathbb{E}_{j,\omega} \left[G(\underline{\varepsilon}_{j,t}^c) \right]}. \quad (11)$$

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

$$\mu_t (1 - \mu_t) (v_h - v_l) \frac{\mathbb{E}_\omega \left[G(\bar{\varepsilon}_{l,t}^c) - G(\bar{\varepsilon}_{h,t}^c) \right]}{\mathbb{E}_{j,\omega} \left[1 - G(\bar{\varepsilon}_{j,t}^c) \right]} = \frac{\text{Cov}_j \left(\mathbb{E}_\omega \left[1 - G(\bar{\varepsilon}_{j,t}^c) \right], v_j \right)}{\mathbb{E}_{j,\omega} \left[1 - G(\bar{\varepsilon}_{j,t}^c) \right]}.$$

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 (10)-(11) reveals that prices under competition and monopoly are similar, 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 (10)-(11) 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 traders 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 $(\bar{\varepsilon}_{j,t}^m, \bar{\varepsilon}_{j,t}^c, \underline{\varepsilon}_{j,t}^m, \underline{\varepsilon}_{j,t}^c)$. Moreover, since dealers know which prices have been offered in equilibrium, each of these thresholds ultimately contains the same information. Consider, for example, $\bar{\varepsilon}_{j,t}^m$, which depends on the ask price A_t^m , 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

$$S_t \equiv R_{t+1} + \omega_t, \quad (12)$$

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 traders' 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 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

$$\omega_{\iota,t}^* = S_t - R_{\iota,t+1}(\mu_{t+1}), \quad (13)$$

where we've made explicit that traders' reservation values at $t + 1$ depend on the evolution of dealers' beliefs. In words, $\omega_{\iota,t}^*$ is the value of ω_t consistent with the signal S_t and traders' optimal behavior—summarized by reservation values—when the asset is of quality $\iota \in \{l, h\}$. Since both $\omega_{l,t}^*$ and $\omega_{h,t}^*$ are consistent with the signal S_t , *by construction*, both $R_{l,t+1}$ and $R_{h,t+1}$ in (13) are calculated under the same information sets, i.e., they both correspond to the same future beliefs, μ_{t+1} .

Figure 1 illustrates the dealers' learning process graphically. The left panel plots the density over signals in the two states of the world, and illustrates how a dealer uses the signal to infer the two possible aggregate shocks. The right panel plots the density $f(\omega)$, which the dealer uses to update his beliefs.

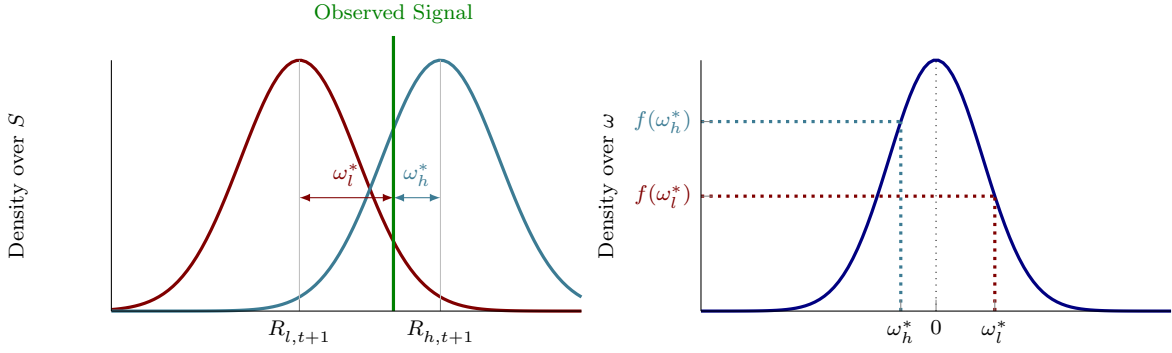


Figure 1: Dealer learning.

In particular, by Bayes Rule, the dealers' updated beliefs are $\mu_{t+1} = \frac{\mu_t f(\omega_{h,t}^*)}{\mu_t f(\omega_{h,t}^*) + (1-\mu_t) f(\omega_{l,t}^*)}$. Since $\omega_{l,t}^*$ and $\omega_{h,t}^*$ themselves depend on future beliefs, the law of motion for μ_t is a function $\mu_{t+1}(\mu_t, S_t)$ that solves the fixed point problem:

$$\mu_{t+1} = \frac{\mu_t}{\mu_t + (1-\mu_t) \frac{f(\omega_{l,t}^*)}{f(\omega_{h,t}^*)}} = \frac{\mu_t}{\mu_t + (1-\mu_t) \frac{f(S_t - R_{l,t+1}(\mu_{t+1}))}{f(S_t - R_{h,t+1}(\mu_{t+1}))}}. \quad (14)$$

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

$$\mu_{t+1} = \frac{\mu_t}{\mu_t + (1-\mu_t) \frac{f(\omega_t + R_{j,t+1}(\mu_{t+1}) - R_{l,t+1}(\mu_{t+1}))}{f(\omega_t + R_{j,t+1}(\mu_{t+1}) - R_{h,t+1}(\mu_{t+1}))}}. \quad (15)$$

In words, given current beliefs μ_t , the true quality of the asset $j \in \{l, h\}$, 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 decisions for traders—without affecting future beliefs.

This verifies the conjecture that the dealers' pricing problem is a static one. In particular, dealers do not have an incentive to deviate from the optimal prices derived above 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 $\omega_t + \varepsilon_{i,t}$. Therefore, at any quoted price, her action can reveal at most $R_{t+1} + \omega_t + \varepsilon_{i,t}$. For example, if the dealer quotes a bid B' and the trader sells, the dealer learns that $R_{t+1} + \omega_t + \varepsilon_{i,t} \leq B'$. However, at the end of the period, the dealer perfectly learns $R_{t+1} + \omega_t$ by observing market-wide volume.¹⁶ Since the idiosyncratic liquidity shock $\varepsilon_{i,t}$ is independent from the quality of the asset, j , 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 , and realizations of the aggregate liquidity shock, ω_t . Such an equilibrium can be represented recursively as a collection of functions $\{\underline{\varepsilon}_j^k, \bar{\varepsilon}_j^k, R_j, A^k, B^k, \mu^+, \tilde{\mu}_j^+, N_j^{0,+}, N_j^{1,+}\}$ for $j \in \{l, h\}$ and $k \in \{m, c\}$ such that:

1. Taking as given the pricing and belief updating functions, traders' decisions to buy or sell are determined by:

$$\underline{\varepsilon}_j^k(\mu, \omega) = B^k(\mu) - \omega - R_j(\tilde{\mu}_j^+(\mu, \omega)) \quad (16)$$

$$\bar{\varepsilon}_j^k(\mu, \omega) = A^k(\mu) - \omega - R_j(\tilde{\mu}_j^+(\mu, \omega)) \quad (17)$$

$$R_j(\mu) = (1 - \delta)c_j + \delta(1 - \pi) \int_{\omega} R_j(\tilde{\mu}_j^+(\mu, \omega)) dF(\omega) + \delta\pi \int_{\omega} \left\{ \sum_{k \in \{m, c\}} \alpha^k [B^k(\mu) G(\underline{\varepsilon}_j^k(\mu, \omega)) + \int_{\underline{\varepsilon}_j^k(\mu, \omega)}^{\bar{\varepsilon}_j^k(\mu, \omega)} [\omega + \varepsilon + R_j(\tilde{\mu}_j^+(\mu, \omega))] dG(\varepsilon) + A^k(\mu) [1 - G(\bar{\varepsilon}_j^k(\mu, \omega))]] \right\} dF(\omega). \quad (18)$$

2. Given traders' 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:

$$0 = \sum_{j \in \{l, h\}} \mu_j \int_{\omega} [1 - G(\bar{\varepsilon}_j^m(\mu, \omega)) - g(\bar{\varepsilon}_j^m(\mu, \omega)) (A^m - v_j)] dF(\omega) \quad (19)$$

$$0 = \sum_{j \in \{l, h\}} \mu_j \int_{\omega} [-G(\underline{\varepsilon}_j^m(\mu, \omega)) + g(\underline{\varepsilon}_j^m(\mu, \omega)) (v_j - B^m)] dF(\omega) \quad (20)$$

¹⁶This follows from the assumption that liquidity shocks have full support, which guarantees non-zero volumes for every state/price. However, as we will show in Section 3, even if shocks are drawn from a finite support, dealers do not have an incentive to experiment.

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:

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

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

3. For any $S \in \mathbb{R}$, dealers' beliefs evolve according to $\mu^+(\mu, S)$, which is a solution to:

$$\mu^+ = \frac{\mu}{\mu + (1 - \mu) \frac{f(S - R_l(\mu^+))}{f(S - R_h(\mu^+))}}. \quad (23)$$

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

$$\mu^+ = \frac{\mu}{\mu + (1 - \mu) \frac{f(\omega + R_j(\mu^+) - R_l(\mu^+))}{f(\omega + R_j(\mu^+) - R_h(\mu^+))}}. \quad (24)$$

Moreover, traders' expectations are consistent with the evolution of dealers' beliefs conditional on observing signal $S = R_j \left(\tilde{\mu}_j^+(\mu, \omega) \right) + \omega$, i.e.,

$$\tilde{\mu}_j^+(\mu, \omega) = \mu^+ \left(\mu, R_j \left(\tilde{\mu}_j^+(\mu, \omega) \right) + \omega \right) \quad \text{for } j \in \{l, h\}. \quad (25)$$

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

$$N_j^{1,+}(\mu, \omega) = N_j^1 \left[1 - \pi + \pi \left(1 - \sum_{k \in \{m, c\}} G \left(\underline{\varepsilon}_j^k(\mu, \omega) \right) \right) \right] + N_j^0 \pi \left(1 - \sum_{k \in \{m, c\}} G \left(\bar{\varepsilon}_j^k(\mu, \omega) \right) \right) \quad (26)$$

$$N_j^{0,+}(\mu, \omega) = N_j^1 \pi \sum_{k \in \{m, c\}} G \left(\underline{\varepsilon}_j^k(\mu, \omega) \right) + N_j^0 \left[1 - \pi + \pi \sum_{k \in \{m, c\}} G \left(\bar{\varepsilon}_j^k(\mu, \omega) \right) \right]. \quad (27)$$

Note that the laws of motion for N_j^1 and N_j^0 depend only on the thresholds $\{\underline{\varepsilon}_j^k, \bar{\varepsilon}_j^k\}$, 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.¹⁷

¹⁷Intuitively, by construction, ω_l^* and ω_h^* 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 ω_l^* are identical to the evolution of these variables when the asset quality is h and the aggregate shock is ω_h^* .

3 Frictions, Learning, and Prices: A Tractable Case

From the definition above, 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. In this section, we make a few parametric assumptions that allow for a full characterization of the equilibrium.¹⁸

We then use this characterization to explore how trading 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 traders 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.

3.1 Parametric Assumptions

We make three key assumptions.

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 simplifies 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.

¹⁸In the next section, we consider a more flexible specification, describe how to solve the model numerically, and confirm that our key results are preserved in a plausible region of the parameter space.

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

$$m \geq \frac{1}{2} (v_h - v_l) \max \left\{ 1, \frac{\delta}{1 - \delta} \right\} \quad \text{and} \quad \frac{e}{2} \geq v_h - v_l + m.$$

This second assumption ensures that, for all prices offered in equilibrium and all realizations of ω , the thresholds $\bar{\varepsilon}_{t,j}^k, \underline{\varepsilon}_{t,j}^k$ 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., $v_j = c_j$ for $j \in \{l, h\}$.*

This last assumption allows for a more direct comparison with many existing models (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 (14) 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 $\omega \in [-m, m]$ and $f(\omega) = 0$ for $\omega \notin [-m, m]$. 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

$$\Sigma_l(\mu) = [-m + R_l(0), -m + R_h(\mu)] \tag{28}$$

$$\Sigma_b(\mu) = [-m + R_h(\mu), m + R_l(\mu)] \tag{29}$$

$$\Sigma_h(\mu) = [m + R_l(\mu), m + R_h(1)]. \tag{30}$$

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 $R_h(\mu) \leq R_h(1)$. 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 (29)–(30). 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

$$p(\mu) = \frac{R_h(\mu) - R_l(\mu)}{2m}. \quad (31)$$

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 traders 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 traders' 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 :

$$r_j(\mu) = \mathbb{E}_\omega [R_j(\mu^+)] = (1 - p(\mu)) R_j(\mu) + p(\mu) R_j(\mathbf{1}[j = h]). \quad (32)$$

Given this notation, it is straightforward to establish (see Appendix A.1) 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, with uniformly distributed ε , the density of marginal traders is the same in both states of the world, so that the covariance between, e.g., $g(\bar{\varepsilon}_j)$ or $g(\underline{\varepsilon}_j)$ and asset quality v_j is zero. From the optimal pricing equations, (8) and (9), this implies that the adverse selection term disappears when a dealer is a monopolist. 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

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

In the Appendix, we prove that, under Assumptions 1-3,

$$\mathbb{E}_j r_j(\mu) = \mathbb{E}_j v_j. \quad (33)$$

Using this property, the expressions for bid and ask prices simplify to:¹⁹

$$B^m = \mathbb{E}_j v_j - \frac{e}{2} \quad (34)$$

$$A^m = \mathbb{E}_j v_j + \frac{e}{2} \quad (35)$$

$$B^c = \mathbb{E}_j v_j - \frac{e}{2} + \sqrt{\left(\frac{e}{2}\right)^2 - \text{Cov}_j(r_j, v_j)} \quad (36)$$

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

Equations (34)–(37) illustrate that the model with uniformly distributed shocks and two types of meetings (monopolist and competitive) is tractable enough to admit analytical 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 reflects that dealers earn rents 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 asymmetric information, as in, e.g., [Glosten and Milgrom \(1985\)](#).

¹⁹This property also implies that $R_j(1 | j = h) = v_j$, i.e., that reservation values under full information are equal to the value of owning the asset and not trading it.

3.4 Reservation Values

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

$$R_j(\mu) = (1 - \delta)v_j + \delta r_j(\mu) + \delta\pi \sum_{k=c,m} \alpha_k \Omega_j^k(\mu), \quad (38)$$

where $r_j(\mu)$ is as defined in (32) and $\Omega_j^k(\mu)$ is the net option value of holding a quality j asset in a type k meeting, i.e., the option value of selling the asset at a later date less the option value of buying it. Under Assumptions 1–3, this reduces to

$$\Omega_j^k(\mu) = \frac{B^k - A^k + 2e}{2e} \left(\frac{B^k - r_j(\mu) - (r_j(\mu) - A^k)}{2} \right) = \frac{B^k - A^k + 2e}{2e} (\mathbb{E}_j v_j - r_j(\mu)). \quad (39)$$

The first term in this expression is the ex ante probability (before ε and ω are realized) that the trader will optimally choose to trade in a type k meeting, given prices A^k and B^k . The second term is the expected difference between the surplus the trader 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 (conditional on beliefs μ), the net option value is decreasing in the expected continuation value, r_j . Moreover, one can show that $r_h(\mu) > r_l(\mu)$. Hence, equation (39) directly implies that the net option value is smaller when the asset is of quality h . 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 v_j , so that $\Omega_h^k(\mu) < \Omega_l^k(\mu)$ for all $k \in \{m, c\}$ and $\mu \in [0, 1]$.*

Intuitively, the net option value is decreasing in v_j because the high quality asset is less likely to be sold and more likely to be bought in equilibrium, while the low quality asset is more likely to be sold and less likely to be bought. Hence, the option to buy (sell) is worth relatively more when the asset quality is h (l).

3.5 Equilibrium Characterization

To characterize the equilibrium, we can use (38)–(39) to write

$$R_h(\mu) - R_l(\mu) = (1 - \delta)(v_h - v_l) + \delta(r_h(\mu) - r_l(\mu)) + \delta\pi \sum_{k=c,m} \alpha_k [\Omega_h^k(\mu) - \Omega_l^k(\mu)]. \quad (40)$$

Then, using (31)–(32) and (34)–(37), we can rewrite (40) as an equation in a single unknown, p , for any given belief μ .²⁰

$$2mp = (1 - \delta)(v_h - v_l) + \delta(1 - \pi)(p(v_h - v_l) + 2mp(1 - p)) - \frac{\delta\pi\alpha_c}{2} \sqrt{1 - \frac{4}{e^2}(v_h - v_l)\mu(1 - \mu)[2m(1 - p)p + p(v_h - v_l)](p(v_h - v_l) + 2mp(1 - p))}. \quad (41)$$

Proposition 1. *Under assumptions 1–3, there exists a unique $p^*(\mu)$ that solves (41).*

Importantly, solving for $p^*(\mu)$ is sufficient for a full characterization of the model: one can use it to construct reservation values $\{R_j(\mu)\}_{j \in \{l, h\}}$ and equilibrium prices $\{A^k, B^k\}_{k \in \{m, c\}}$ using (34)–(39). In the next subsection, we exploit this characterization to understand how the key frictions in the model affect equilibrium outcomes.

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).

The Effect of Beliefs

Recall that spreads in monopolist meetings are constant, $A^m - B^m = e$ for all μ , so beliefs only affect bid-ask spreads in competitive meetings. In the Appendix, we establish the following results.

Lemma 1. *The following objects are all hump-shaped in μ with a maximum at $\mu = 0.5$: (i) the bid-ask spread in competitive meetings, $A^c - B^c$; (ii) the difference in reservation values, $R_h - R_l$; and (iii) the probability that the true asset quality is revealed, p .*

Recall from the pricing equations (36)–(37) that the bid-ask spread in competitive meetings is increasing in the covariance between reservation values and asset quality. At the core of the result in Lemma 1 is the observation that this covariance is maximized when uncertainty is maximal (i.e., when $\mu = 0.5$). To

²⁰We provide a more detailed derivation in the proof of Proposition 1, in the Appendix.

see why, we establish in the Appendix that

$$\text{Cov}_j(r_j, v_j) = \mu(1 - \mu)(v_h - v_l)^2 \left[1 - (1 - p(\mu)) \left(1 - \frac{R_h(\mu) - R_l(\mu)}{v_h - v_l} \right) \right]. \quad (42)$$

The first term on the right-hand side of (42) captures the direct effect of μ on the bid-ask spread: this term is simply the prior variance about the asset's quality, and is clearly maximized at $\mu = 0.5$. The second term captures the indirect effects of μ on the bid-ask spread. Understanding this term requires understanding the equilibrium interaction between spreads, reservation values, and learning.²¹ In particular, when dealers are uncertain ($\mu \approx 0.5$), the wider spreads from the direct effect makes traders less likely to trade, as a larger fraction of idiosyncratic liquidity shocks lie in the "inaction region." As the likelihood of trade falls, the difference in the net option values (across the two states) narrows. Intuitively, a decline in the probability of trading 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 , causing $\Omega_h - \Omega_l$ to be less negative. This leads to an increase in $R_h - R_l$ and $p = (R_h - R_l)/2m$, i.e., the gap between reservation values widens and thus learning speeds up. As a result, $\text{Cov}(r_j, v_j)$ increases, leading to a further increase in spreads.

The Effect of Trading frictions

Now, consider a decrease in the severity of search frictions. In what follows, we focus on the effect of changing the frequency of trading opportunities, π , and show later that the effects of changing the fraction of competitive meetings, α_c , are similar. Our first result states that an increase in π unambiguously slows down learning.

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

To understand the intuition, recall from (31) that p is proportional to $R_h - R_l$. From (40), we have

$$\frac{d(R_h - R_l)}{d\pi} = \delta \left[\sum_k \alpha_k (\Omega_h^k - \Omega_l^k) + \frac{d(r_h - r_l)}{d\pi} + \pi \sum_k \alpha_k \frac{d(\Omega_h^k - \Omega_l^k)}{d\pi} \right] \quad (43)$$

Again, the first term in this expression represents the direct effect of increasing π : it places more weight on the difference in net option values, $\Omega_h^k - \Omega_l^k$, which are negative. As a result, the direct effect of raising π is to attenuate the difference in reservation values, $R_h - R_l$. The remaining two terms in (43) capture the indirect effects of a change in trading frictions, which operate through the expected reservation values r_j .

These effects are similar to the feedback channel described in the discussion of Lemma 1: an increase in

²¹Indeed, if the traders did not account for the option value of future trading—i.e., if $R_j = v_j$ —then the second term would drop out and beliefs would only affect spreads through the typical, direct effect captured in the first term.

π raises the likelihood of trading, making the difference in net option values more negative. Hence, these indirect effects further reduce the difference in reservation values.

We now turn to the ultimate effect of trading frictions on bid-ask spreads. There are two, opposing effects of increasing π . 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, which implies a decrease in $\text{Cov}(r_j, v_j)$. Recall from (34)–(37), the bid-ask spread in competitive meetings is given by

$$A^c - B^c = e - \sqrt{e^2 - 4\text{Cov}_j(r_j, v_j)}.$$

Hence, holding beliefs fixed, a decline in trading frictions leads to a lower spread. Intuitively, since the increase in π makes traders behave more similarly in the two states of the world—i.e., the likelihood of a trader 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 it also influences how beliefs change over time. We refer to this as the *dynamic effect*: an increase in π implies that dealers will remain uncertain about the true asset quality for longer (Lemma 2). Since bid-ask spreads are larger when dealers are more uncertain (Lemma 1), this dynamic effect implies that an increase in π puts upward pressure on spreads. We will show that this dynamic effect eventually dominates, implying that more frequent trading opportunities leads to wider spreads in the long run.

To state this formally, we let $Y_{j,t}^c$ denote the expected competitive spread in period t when the asset is of quality $j \in \{l, h\}$. Formally,

$$Y_{t,j}^c \equiv \mathbb{E}_{\omega^t} [(A_t^c - B_t^c) \mid j].$$

Plugging in the expressions for equilibrium bid and ask prices yields

$$Y_{t,j}^c = \underbrace{(1 - p(\mu_0))^t}_{\text{Prob(No revelation)}} \left[e - \sqrt{e^2 - 4\text{Cov}_j(r_j, v_j)} \right],$$

where μ_0 is the dealers' initial beliefs. Thus, $Y_{j,t}^c$ is the product of probability that the state has not yet been revealed by time t and the bid-ask spread when dealers remain uncertain of the asset quality (note that the competitive spread after quality is revealed falls to 0). The former is determined by the speed of

learning, $p(\cdot)$. Differentiating with respect to π yields

$$\frac{\partial Y_{t,j}^c}{\partial \pi} = \underbrace{\frac{2(1-p(\mu_0))^t}{\sqrt{e^2 - 4\text{Cov}_j(r_j, v_j)}}}_{>0} \underbrace{\frac{\partial \text{Cov}_j(r_j, v_j)}{\partial \pi}}_{<0} + \underbrace{-t(1-p(\mu_0))^{t-1}}_{<0} \underbrace{\frac{\partial p}{\partial \pi}}_{<0} \underbrace{\left[e - \sqrt{e^2 - 4\text{Cov}_j(r_j, v_j)} \right]}_{>0} \quad (44)$$

The first term in (44) reflects the static effect: holding beliefs fixed, an increase in π reduces the covariance between the quality of the asset and reservation values, reducing adverse selection and, therefore, spreads. The second term captures the dynamic effect: an increase in π reduces the probability that the state is revealed ($\frac{\partial p}{\partial \pi} < 0$), which puts upward pressure on the spread. As t grows and $(1-p(\mu))^t \rightarrow 0$, the first term shrinks relative to the second term, and thus the dynamic effect eventually dominates. That is, an increase in π pushes the spread wider for sufficiently large t .

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

Since the monopoly spread is constant for all μ , the *average* quoted spread is also eventually larger when π goes up. Defining the expected average spread, conditional on state j , by

$$Y_{t,j} \equiv \mathbb{E}_{\omega^t} \left[\sum_{k=c,m} \alpha_k (A_t^k - B_t^k) \mid j \right],$$

the next result follows immediately.

Corollary 1. *There exists a $\tau < \infty$ such that $\frac{\partial Y_{t,j}}{\partial \pi} > 0$ for all $t \geq \tau$.*

The Effect of Competition

Finally, we note that the effects of changing the fraction of competitive trading opportunities (α_c) are similar to those from changing the frequency of overall trading opportunities (π). First, the following result establishes that a rise in α_c unambiguously slows down the learning process.

Proposition 4. *For any $\mu \in (0, 1)$, $\frac{\partial p^*(\mu)}{\partial \alpha_c} < 0$.*

Intuitively, since spreads are smaller in competitive meetings (compared to monopolist meetings), an increase in α_c implies a relatively higher likelihood of trading. As before, this amplifies the difference in the net option values, i.e., $\Omega_h^c - \Omega_l^c < \Omega_h^m - \Omega_l^m < 0$.²² As a result, a rise in the fraction of competitive meetings attenuates $R_h - R_l$, which slows down learning. The effect on bid-ask spreads then follows from the same logic as Proposition 3: the dynamic effect from slower learning eventually causes spreads (in competitive meetings) to widen.

²²This follows from (39) and the fact that $A^c - B^c < A^m - B^m$.

Proposition 5. *There exists a $\tau < \infty$ such that $\frac{\partial Y_{t,j}^c}{\partial \alpha_c} > 0$ for all $t \geq \tau$.*

However, increasing α_c has slightly different effects on the *average* quoted spread than increasing π . In particular, since spreads in competitive meetings are smaller than those in monopoly meetings, an increase in α_c has an additional, compositional effect that puts downward pressure on the average spread. As a result, though it's still possible that the average spread is increasing in α_c for some t , one can also show that this compositional effect dominates as $t \rightarrow \infty$.

4 Numerical Analysis

In Section 3, we established a sequence of analytical results under several assumptions, most notably that aggregate and idiosyncratic liquidity shocks were drawn from uniform distributions. In this section, we relax these assumptions. This complicates the analysis considerably, as the interaction between optimal bid and ask prices, reservation values, and future prices and beliefs introduce additional effects that were absent from our analysis with uniform distributions. However, while these complications make analytical results harder to derive, it is straightforward to solve the model numerically. In what follows, we parameterize the model, solve for the equilibrium numerically, and confirm that the key results established in the model with uniform shocks—which afforded us tractability—are relevant in an empirically plausible region of the parameter space in the model with a more flexible specification.

4.1 Solving the model

Before specifying our parametric assumptions, we briefly describe the iterative process by which one can solve the general model. First, given a grid for the state variables μ and ω , we start with an initial guess for the reservation values, $R_j(\mu)$, for $j \in \{l, h\}$. Then, given $R_j(\mu)$, traders' expectations of dealer beliefs μ^+ for each ω are obtained by solving (15). Using these updating equations, we compute optimal prices for any beliefs μ . The processes for beliefs and prices can then be combined to yield an updated guess for the reservation value functions $R_j(\mu)$, using (4). We use this to update the original guess, and repeat until convergence.

4.2 Parameterization

We assume that both aggregate and idiosyncratic liquidity shocks are drawn from a mean-zero normal distribution, i.e., that $\omega \sim N(0, \sigma_\omega^2)$ and $\varepsilon \sim N(0, \sigma_\varepsilon^2)$. The normality of idiosyncratic liquidity shocks implies that bid and ask prices in monopoly meetings are no longer constant (as in Section 3), since both

the market power and asymmetric components described in equations (8) and (9) now depend on beliefs, μ_t . Moreover, the normality of aggregate liquidity shocks implies that learning occurs gradually, unlike the stark form it took under a uniform distribution.

To assign parameter values, we choose a widely studied over-the-counter market as a guide: the market for US corporate bonds. To start, we interpret differences in asset quality as stemming from changes in credit ratings. Consider a bond that is rated AAA at some initial point in time. Conditional on not being downgraded, we interpret v_h as the expected payoff of the bond upon maturity. If it is downgraded (most likely, to AA), we assume that the expected payoff drops to v_l . Since bid-ask spreads and beliefs are only a function of the *relative* payoff in the two states (i.e., $v_h - v_l$), we can normalize $v_l = 0$. Then, we map the relative payoff 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 20 bps higher than spreads on AAA-rated bonds.²³ For a 5-year par bond with face value \$100 and a coupon of 2%, this difference in spread translates into a price change of \$0.95. Hence, we set $v_h = 0.95$.

Next, we choose the initial belief μ_0 to match the unconditional probability of a rating transition. 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 parameters governing the distribution of liquidity shocks and the meeting probabilities, we again rely on estimates from [Feldhütter \(2012\)](#). He estimates the parameters of a continuous time model of over-the-counter trading along 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.²⁴ Assuming that aggregate and idiosyncratic components are of equal magnitude, this procedure yields $\sigma_\omega^2 = \sigma_\varepsilon^2 = 0.16$.²⁵

[Feldhütter \(2012\)](#) also provides estimates for the arrival rate of meetings with dealers from the perspective of traders in the market. We map his estimate—which is an annualized rate of 40—into the parameters governing our meeting technology: the probability of meeting at least 1 dealer (π) and the conditional probability of meeting more than 1 dealer (α_c).²⁶ Interpreting a period in our model as 5

²³See Table 1 of that paper. The average spread on a AAA bond is 5 bps (for trades greater than USD 100,000 in size), while the corresponding spread for AA bonds is 25 bps.

²⁴Specifically, he estimates a flow holding cost of 2.91, which lasts for an average of 0.31 years or equivalently, a total expected 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 receiving negative shocks. With normally distributed shocks, this translates into a variance of 0.32.

²⁵In the Appendix, we repeat the analysis under alternative assumptions and show that they produce similar results.

²⁶[Feldhütter \(2012\)](#) estimates arrival rates separately for different trade sizes. We use the one for the smallest trade size.

business days or 0.02 years (since a year is assumed to have 250 business days) yields our baseline values $(\pi, \alpha_c) = (0.55, 0.25)$.²⁷

This leaves δ . As we discussed in Section 2.1, this parameter can be interpreted in several ways. In particular, one could interpret it as the likelihood of the asset not maturing in each period, so that $1/(1 - \delta)$ represents an asset’s expected maturity. Alternatively, one could interpret δ as the probability that an individual trader remains active in the market, so that $1/(1 - \delta)$ represents the expected amount of time a trader devotes to actively buying or selling assets, before holding his current position until maturity. We adopt the latter interpretation and adopt a baseline value of $\delta = 0.9$, which implies that a trader remains in the market for about 10 weeks.

Parameter	Value	Description
Payoffs:		
v_h	0.95	Payoff of h-quality asset
v_l	0	Payoff of l-quality asset
σ_ω^2	0.16	Variance of aggregate liquidity shock
σ_ε^2	0.16	Variance of idiosyncratic liquidity shock
Meeting technology:		
$1 - \delta$	0.10	Pr(asset maturing)
π	0.55	Pr(meeting at least 1 dealer)
α_c	0.35	(Conditional) Pr(meeting > 1 dealer)
Beliefs:		
μ_0	0.50	Initial dealer belief

Table 1: **Parameters**

Table 1 summarizes our parameter choices. Figure 2 plots the two key equilibrium objects—reservation values and spreads—as a function of dealers’ beliefs, μ . The red solid line plots these objects under our baseline parameterization with $\pi = 0.55$. To illustrate the effects of trading frictions, we repeat the plots when frictions become more severe ($\pi = 0.35$, the dashed line) and less severe ($\pi = 0.75$, the dotted line).

The plots confirm our key findings from Section 3. The top left panel, which plots $R_h(\mu) - R_l(\mu)$, confirms that reservation values in the two states become closer to each other as π increases. The remaining panels plot spreads in monopoly meetings, competitive meetings, and on average. These plots confirm that spreads are hump-shaped in μ , as the effects of asymmetric information are most pronounced when uncertainty is highest (i.e., for intermediate values of μ). The figure also confirms 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, as noted above, increasing π reduces $R_h - R_l$ for any given μ . As a result, traders behave more similarly in the two states of the world, so that dealers’ optimal prices require a smaller

²⁷Note that this frequency of trade is also consistent with the estimates in He and Milbradt (2014).

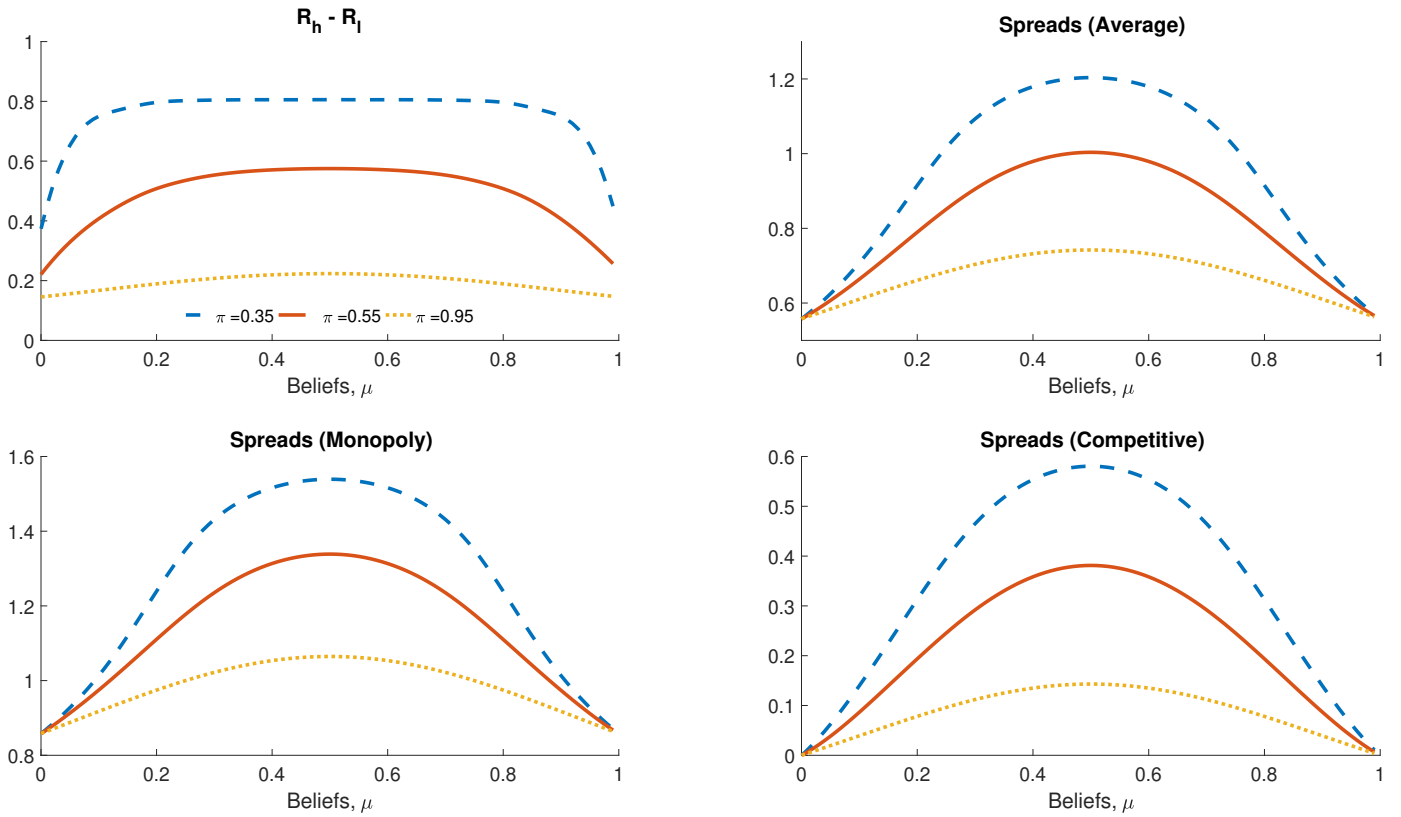


Figure 2: Effect of π on reservation values and spreads.

adjustment for asymmetric information.

However, note that the plots in Figure 2 illustrate only the static effects of varying π , i.e., holding beliefs fixed. Figure 3 illustrates the dynamic effects by plotting the evolution of beliefs and spreads over time for different values of π , given the true quality of the asset is h . The panels plot the average values across 40,000 sample paths—recall that the realized path of beliefs (and thus prices) depend on the realized sequence of ω_t . Dealers' beliefs (top left panel) start at $\mu_0 = 0.5$ and drift upwards. Since the true state is h , beliefs in all three cases eventually converge to 1. However, the pace is clearly slower when π is higher; more frequent trading opportunities leads to slower learning, exactly as in the model in Section 3. Spreads are initially tighter 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. This is true both for monopoly and competitive spreads, as the bottom two panels show.

So far, we have focused on bid-ask spreads. In Figure 4, we plot two other metrics for liquidity that are commonly used in the literature. The left panel plots average trading volume, defined as the sum of

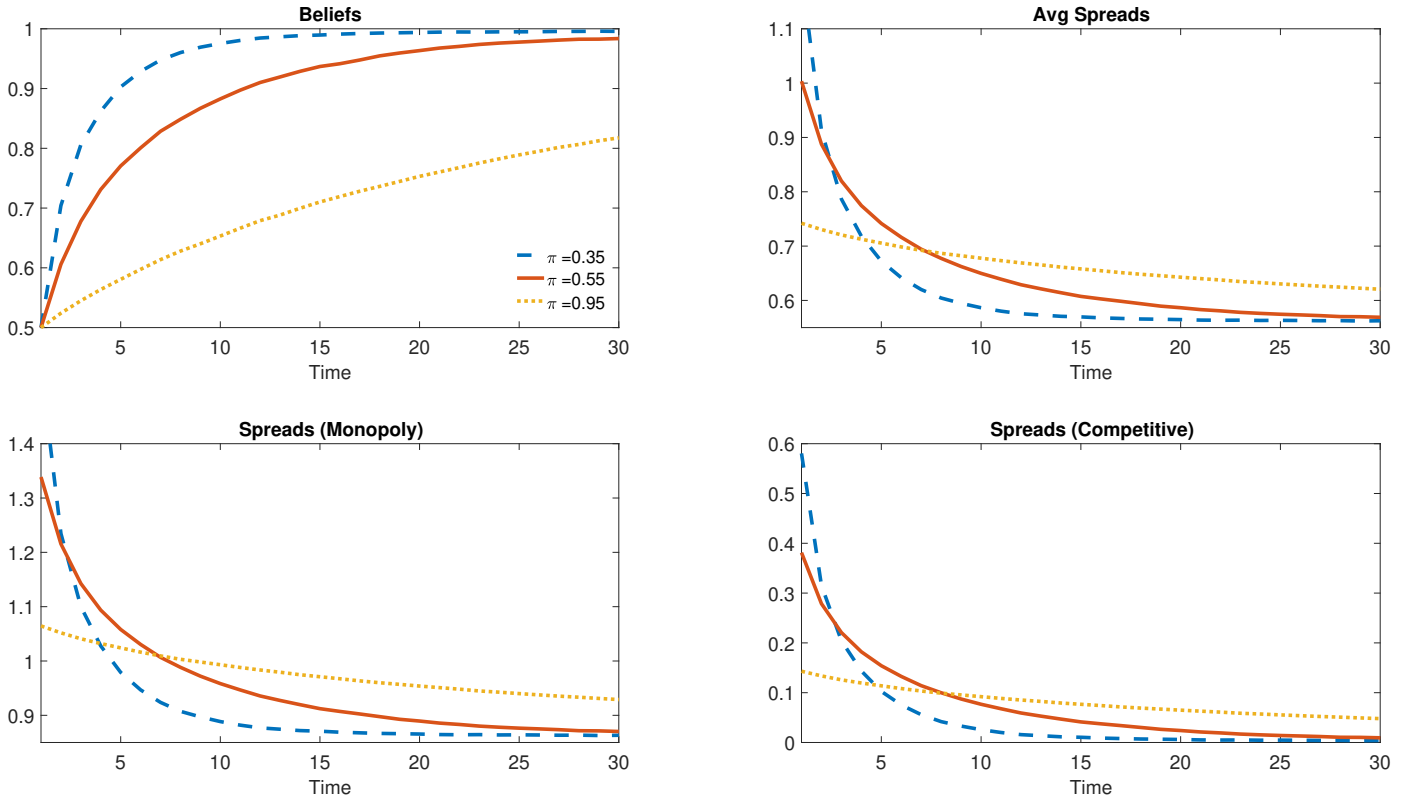


Figure 3: Effect of π on beliefs and spreads over time.

buys and sells by traders, averaged across meetings and sample paths. The second plots price impact, defined as the elasticity of the change in prices with respect to net volume.²⁸ The figure highlights that different measures of liquidity need not move together in markets where both information and trading frictions are present. In particular, the left panel illustrates that trading volume increases monotonically with the frequency of trading opportunities; even when spreads widen, the direct effect of more meetings dominates the indirect effect of fewer trades per meeting. Hence, unlike many traditional models, our framework can rationalize a positive relationship between spreads and trading volume, as documented in, e.g., [Lin et al. \(1995\)](#) and [Chordia et al. \(2001\)](#).²⁹ Price impact, on the other hand, displays a pattern similar to that of bid-ask spreads: higher values of π lead to smaller price impact initially, but larger price impact in later periods. Intuitively, when π is high, dealers learn more slowly. For fixed beliefs,

²⁸Formally, price changes and net volume are defined as follows:

$$\text{PrcChg}_t = \log \sum_{k=c,m} \alpha_k \left(\frac{A_t^k + B_t^k}{2} \right) - \log \sum_{k=c,m} \alpha_k \left(\frac{A_{t-1}^k + B_{t-1}^k}{2} \right)$$

$$\text{NetVol}_{t-1} = \log(\text{Buys}_{t-1}) - \log(\text{Sells}_{t-1}) = \log \left(N_{t-1}^0 \pi \sum_{k=c,m} \alpha_k (1 - G(\bar{\varepsilon}_{t-1}^k)) \right) - \log \left(N_{t-1}^1 \pi \sum_{k=c,m} \alpha_k G(\bar{\varepsilon}_{t-1}^k) \right)$$

We regress (separately for each t) PrcChg_t on NetVol_{t-1} and a constant. The graph plots the coefficient on NetVol_{t-1} .

²⁹Also see [Liu and Wang \(2016\)](#).

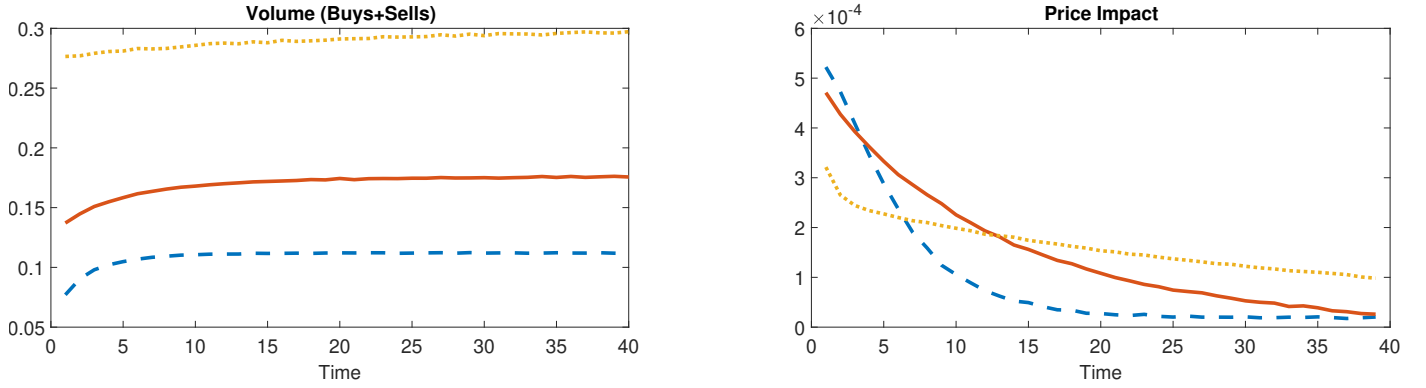


Figure 4: Effect of π on volume and price impact over time.

this manifests itself as a reduced impact of trading activity on prices. However, this also implies that uncertainty is greater in later rounds, which increases the scope for learning from volume. As with bid-ask spreads, this dynamic effect eventually dominates and so, price impact goes up in later rounds.

On a final note, our results point to a more general insight: the conventional predictions from models with only one friction do not necessarily extend to setting with multiple frictions. While we have focused exclusively on the effect of reducing trading frictions, our framework can also be used to investigate the implications of mitigating asymmetric information. For example, it is possible to show that, for certain parameter values, bid-ask spreads are convex in dealer beliefs, μ . This implies that additional information, which induces a mean-preserving spread in beliefs, can lead to lower liquidity in the form of wider bid-ask spreads. We demonstrate this using a numerical example in Appendix C.2, and leave a more formal analysis for future work.

5 A Stationary Version

In this section, we modify our baseline framework to allow the asset quality to change over time. This implies that the equilibrium, conditional on the asset not maturing, can be characterized in terms of prices, allocations, and the distribution of beliefs that prevail in the stochastic steady state.

5.1 Characterizing the Stochastic Steady State Equilibrium

We assume that the asset quality changes in each period with probability ρ . Otherwise, the model remains unchanged. The equations characterizing the equilibrium change only slightly relative to the baseline, non-stationary version. Dealers' beliefs about asset quality at the end of each period are still given by equation (24). However, to obtain beliefs in the following period, we need to adjust for the possibility

that asset quality has changed. In particular, if dealers' beliefs at the end of a period are given by μ , then beliefs in the following period (before pricing decisions are made) can be written

$$\mathcal{M}(\mu) = \mu(1 - \rho) + (1 - \mu)\rho. \quad (45)$$

Similarly, the expressions for traders' reservation values also must be adjusted to reflect the possibility of the asset quality changing. Given asset quality $j \in \{l, h\}$, let $\rho_j \equiv \rho$ and $\rho_{-j} \equiv 1 - \rho$. Then we have

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

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

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

The pricing equations take the same form as in the baseline, non-stationary version, with $(\tilde{v}_h, \tilde{v}_l)$ replacing (v_h, v_l) . The iterative algorithm described in Section 4 can then be applied to this modified system of equations to solve for the laws of motion for beliefs, reservation values, and prices. The stochastic steady state is then obtained by simulation.

5.2 Comparative Statics

We now turn to the effects of varying trading frictions on the stochastic steady state. We set the switching probability ρ to 0.5% and hold all other parameters fixed at their baseline values. Figure 5 shows the results for three different values of π . The top left panel plots the distribution of beliefs in the long run (again, conditional on the asset not maturing). It shows that higher π shifts the mass of the distribution towards the center, where dealers are more uncertain about asset quality. As in the non-stationary model, more frequent trading opportunities reduces the speed of learning and keeps uncertainty high.

The effect on average spreads—illustrated in the remaining 3 panels—is more complicated. As in the non-stationary version, there are two opposing forces. On the one hand, increasing π reduces adverse selection and, therefore, pushes spreads down for any level of beliefs, μ . However, increasing π also pushes the long-run distribution of beliefs towards intermediate values of μ , which are associated with

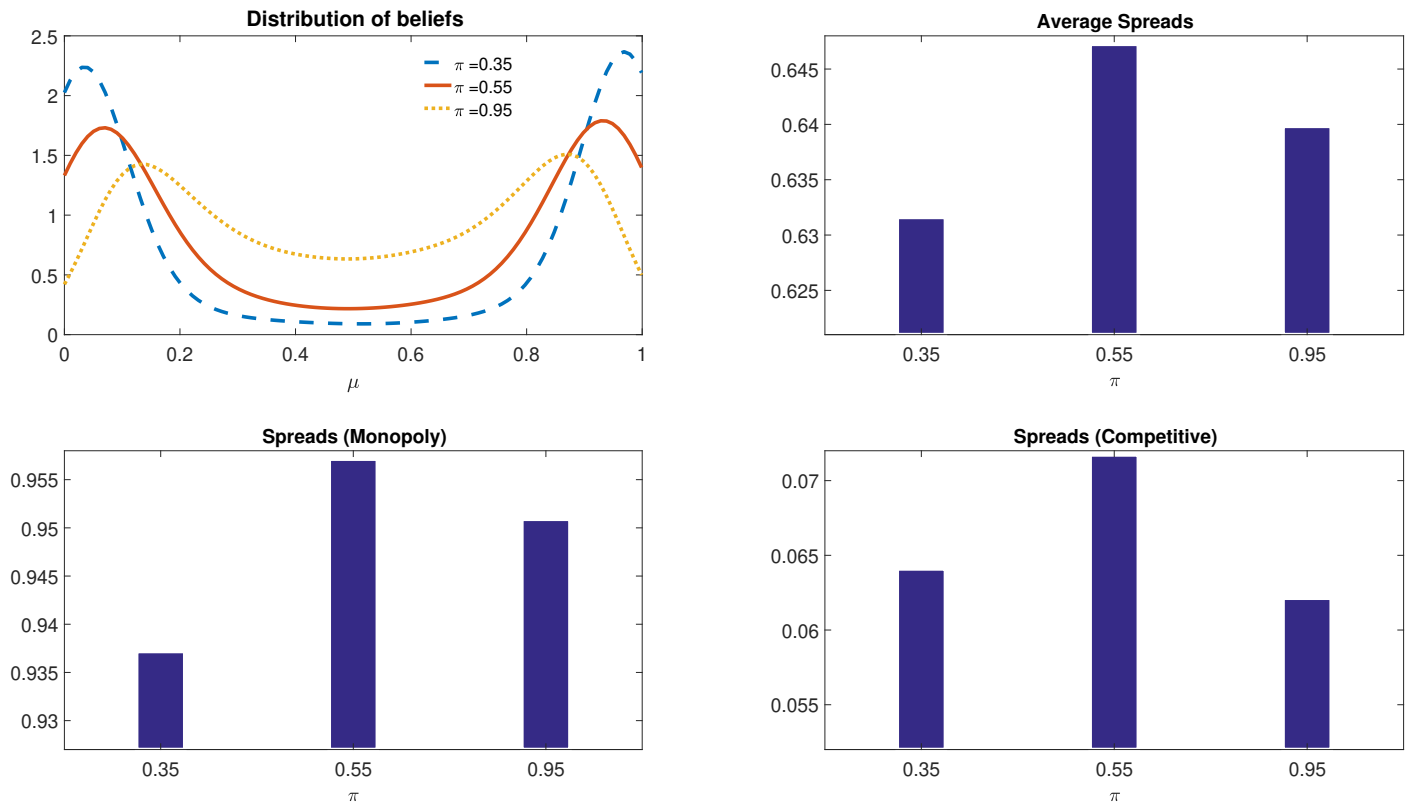


Figure 5: Effect of π in stationary version.

higher spreads. Whether average spreads are wider or narrower in the stochastic steady state depends on which of these two effects dominates. For example, when spreads rise from 0.35 to 0.55, the latter effect dominates and average spreads widen. However, when π rises, further from 0.55 to 0.95, the former effect dominates and spreads narrow.

Welfare Next, we explore implications for welfare, defined as the sum of the payoffs of all agents (traders and dealers) in the economy. As payments between traders and dealers cancel out, welfare only depends on the extent to which gains from trade (arising from the aggregate and idiosyncratic liquidity shocks) are exploited. This, in turn, depends on the thresholds that govern trading. Formally, let

$$Q_{jt} \equiv N_t^1 Q_j^1(\mu_t, \omega_t) + N_t^0 Q_j^0(\mu_t, \omega_t),$$

where

$$\begin{aligned}
Q_j^1(\mu, \omega) &= (1 - \pi) \left[\int_{-\infty}^{\infty} (\omega + \varepsilon) dG(\varepsilon) \right] + \pi \sum_{k=c,m} \alpha_k \left[\int_{\underline{\varepsilon}_j^k}^{\infty} (\omega + \varepsilon) dG(\varepsilon) \right] \\
&= (1 - \pi)\omega + \pi \left[\omega \sum_{k=c,m} \alpha_k (1 - G(\underline{\varepsilon}_j^k)) + \sum_{k=c,m} \alpha_k \int_{\underline{\varepsilon}_j^k}^{\infty} \varepsilon dG(\varepsilon) \right], \\
Q_j^0(\mu, \omega) &= \pi \sum_{k=c,m} \alpha_k \int_{\bar{\varepsilon}_j^k}^{\infty} (\omega + \varepsilon) dG(\varepsilon) \\
&= \pi \left[\omega \sum_{k=c,m} \alpha_k (1 - G(\bar{\varepsilon}_j^k)) + \sum_{k=c,m} \alpha_k \int_{\bar{\varepsilon}_j^k}^{\infty} \varepsilon dG(\varepsilon) \right].
\end{aligned}$$

In words, Q_{jt} represents the gains from trade that are realized in period t when the asset is of quality j . The first component of Q_{jt} is the gains from trade that are realized by those traders who enter period t owning a unit of the asset: there are N_t^1 such agents, and the average gains from trade that accrue to these agents, Q_j^1 , is the average liquidity shock, adjusted for the fact that those with idiosyncratic realizations below $\underline{\varepsilon}_j$ sell their asset if they have the opportunity to do so. The second component is the gains from trade that are realized by those traders who enter period t not owning a unit of the asset.

The expressions above illustrate the two channels through which reducing trading frictions affects welfare. The first channel is standard: *more meetings per period* imply more opportunities for exploiting gains from trade. The second channel acts through the thresholds $\bar{\varepsilon}_j$ and $\underline{\varepsilon}_j$, which depend on prices and, therefore, beliefs. Again, an increase in π slows down learning, causing spreads to widen and thus *fewer trades per meeting*.

In the left panel of Figure 6, we plot the average value of Q_{jt} in the stochastic steady for the three values of π analyzed earlier. It shows that, under our calibration, the first channel dominates, i.e., welfare increases with π , even though bid-ask spreads are hump-shaped (as shown in the right panel, which is reproduced from Figure 5).

6 Conclusion

Many assets have traditionally traded in markets with frictions: the process of finding a counterparty, receiving price quotes, and executing trades was costly and/or took significant amounts of time. As a result of both technological innovations and regulatory mandates, these frictions have been reduced in recent years. What are the expected effects of these changes on observable outcomes like prices, spreads, and trading volume?

In order to answer to this question, we develop a framework that incorporates two canonical sources

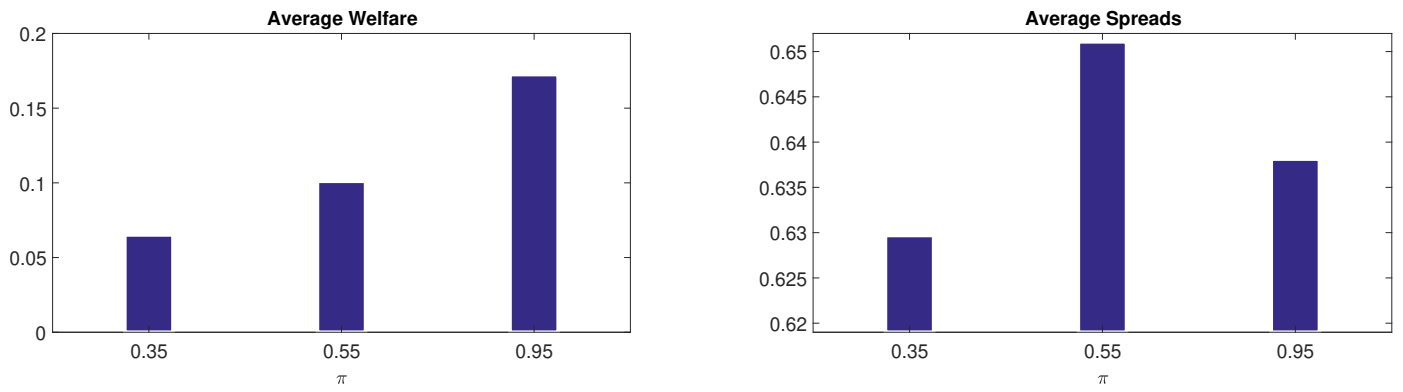


Figure 6: Effect of π on welfare (Q_{jt}) and spreads.

of illiquidity: search frictions and information frictions. Our analysis reveals novel interactions between these two frictions, which can overturn the conventional wisdom that derives from studying them in isolation. In particular, we show that reducing search frictions can slow down the process of information revelation, thus exacerbating the effects of asymmetric information and causing a deterioration in (standard measures of) market liquidity. These results can help us understand existing empirical evidence, and evaluate the effects of future reforms. In addition, our analysis highlights that various measures of liquidity need not respond to shocks in the same direction, nor are these measures of liquidity necessarily informative about welfare.

Our framework opens up a variety of opportunities for future research. For one, our results suggest that reducing information frictions—say, through regulations that promote transparency or require more disclosure—can also have counter-intuitive effects on liquidity and welfare. While we touch on this in the current paper (see Appendix C.2), a more careful exploration is left for future work. Second, our equilibrium analysis offers a structural framework to disentangle the effects of search and information frictions on empirical measures of liquidity such as bid-ask spreads, price impact, or trading volume that are reported in recent, transaction-level data sets from OTC markets. Lastly, our model is amenable to a variety of extensions, including other sources of market illiquidity (e.g., inventory costs).

References

- AFONSO, G. (2011): "Liquidity and congestion," *Journal of Financial Intermediation*, 20, 324–360. 5
- ASRIYAN, V., W. FUCHS, AND B. GREEN (2017): "Information Aggregation in Dynamic Markets with Adverse Selection," . 6
- BABUS, A. AND P. KONDOR (2016): "Trading and Information Diffusion in Over-the-Counter Markets," *working paper*. 6
- BETHUNE, Z., B. SULTANUM, AND N. TRACHTER (2016): "Private Information in Over-the-Counter Markets," . 5
- BIAIS, B., T. FOUCAULT, AND S. MOINAS (2015): "Equilibrium fast trading," *Journal of Financial Economics*, 116, 292–313. 6
- 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. 6
- BOLLERSLEV, T. AND M. MELVIN (1994): "Bid-ask spreads and volatility in the foreign exchange market: An empirical analysis," *Journal of International Economics*, 36, 355–372. 2
- BRANCACCIO, G., D. LI, AND N. SCHÜRHOFF (2017): "Learning by Trading: The Case of the US Market for Municipal Bonds," . 2
- BURDETT, K. AND K. L. JUDD (1983): "Equilibrium Price Dispersion," *Econometrica*, 51, 955–969. 8
- CAMARGO, B. AND B. LESTER (2014): "Trading dynamics in decentralized markets with adverse selection," *Journal of Economic Theory*, 153, 534–568. 5
- CHIU, J. AND T. V. KOEPL (2016): "Trading dynamics with adverse selection and search: Market freeze, intervention and recovery," *The Review of Economic Studies*, 83, 969–1000. 5
- CHOI, M. (2016): "Imperfect Information Transmission and Adverse Selection in Asset Markets," *Unpublished paper, University of Iowa*. 5
- CHORDIA, T., R. ROLL, AND A. SUBRAHMANYAM (2001): "Market liquidity and trading activity," *The Journal of Finance*, 56, 501–530. 4, 30

- COPELAND, T. E. AND D. GALAI (1983): "Information effects on the bid-ask spread," *the Journal of Finance*, 38, 1457–1469. 5
- DU, S. AND H. ZHU (2017): "What is the optimal trading frequency in financial markets?" *The Review of Economic Studies*, 84, 1606–1651. 6
- DUFFIE, D., N. GÂRLEANU, AND L. H. PEDERSEN (2005): "Over-the-Counter Markets," *Econometrica*, 73, 1815–1847. 2, 5, 20, 27
- DUFFIE, D., S. MALAMUD, AND G. MANSO (2009): "Information percolation with equilibrium search dynamics," *Econometrica*, 77, 1513–1574. 6
- DUFFIE, D. AND G. MANSO (2007): "Information percolation in large markets," *The American economic review*, 97, 203–209. 6
- EDWARDS, A. K., L. E. HARRIS, AND M. S. PIWOWAR (2007): "Corporate bond market transaction costs and transparency," *The Journal of Finance*, 62, 1421–1451. 2
- FELDHUTTER, P. (2012): "The Same Bond at Different Prices: Identifying Search Frictions and Selling Pressures," *The Review of Financial Studies*, 25, 1155–1206. 27
- FELDHÜTTER, P. (2012): "Short-term corporate bond yield spreads," . 27
- FUCHS, W. AND A. SKRZYPACZ (2015): "Government interventions in a dynamic market with adverse selection," *Journal of Economic Theory*, 158, 371–406. 5
- GEHRIG, T. (1993): "Intermediation in search markets," *Journal of Economics & Management Strategy*, 2, 97–120. 5
- 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. 2, 5, 11, 12, 18, 20
- GOLDSTEIN, M. A., E. S. HOTCHKISS, AND E. R. SIRRI (2006): "Transparency and liquidity: A controlled experiment on corporate bonds," *The review of financial studies*, 20, 235–273. 2
- GOLOSOV, M., G. LORENZONI, AND A. TSYVINSKI (2014): "Decentralized trading with private information," *Econometrica*, 82, 1055–1091. 6
- GROSSMAN, S. J. AND J. E. STIGLITZ (1980): "On the impossibility of informationally efficient markets," *The American economic review*, 70, 393–408. 6

- 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. 5
- HARRIS, L. E. AND M. S. PIOWAR (2006): “Secondary trading costs in the municipal bond market,” *The Journal of Finance*, 61, 1361–1397. 2
- HE, Z. AND K. MILBRADT (2014): “Endogenous liquidity and defaultable bonds,” *Econometrica*, 82, 1443–1508. 28
- HELLWIG, M. F. (1980): “On the aggregation of information in competitive markets,” *Journal of economic theory*, 22, 477–498. 6
- HENDERSHOTT, T. AND P. C. MOULTON (2011): “Automation, speed, and stock market quality: The NYSE’s hybrid,” *Journal of Financial Markets*, 14, 568–604. 4
- HUGONNIER, J., B. LESTER, AND P.-O. WEILL (2014): “Heterogeneity in decentralized asset markets,” Tech. rep., National Bureau of Economic Research. 5
- KAYA, A. AND K. KIM (2015): “Trading dynamics with private buyer signals in the market for lemons,” . 5
- KIM, K. (2017): “Information about sellers’ past behavior in the market for lemons,” *Journal of Economic Theory*, 169, 365–399. 5
- KYLE, A. S. (1985): “Continuous auctions and insider trading,” *Econometrica: Journal of the Econometric Society*, 1315–1335. 5
- LAGOS, R. AND G. ROCHETEAU (2009): “Liquidity in Asset Markets with Search Frictions,” *Econometrica*, 77, 403–426. 5
- LAUERMANN, S., W. MERZYN, AND G. VIRÁG (2017): “Learning and price discovery in a search market,” *The Review of Economic Studies*. 6
- LAUERMANN, S. AND A. WOLINSKY (2016): “Search with adverse selection,” *Econometrica*, 84, 243–315. 6
- 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. 5
- LESTER, B. R., A. SHOURIDEH, V. VENKATESWARAN, AND A. ZETLIN-JONES (2017): “Screening and Adverse Selection in Frictional Markets,” *Journal of Political Economy*. 8

- LIN, J.-C., G. C. SANGER, AND G. G. BOOTH (1995): "Trade size and components of the bid-ask spread," *The Review of Financial Studies*, 8, 1153–1183. 4, 30
- LIU, H. AND Y. WANG (2016): "Market making with asymmetric information and inventory risk," *Journal of Economic Theory*, 163, 73–109. 30
- MA, C. K., R. L. PETERSON, AND R. S. SEARS (1992): "Trading noise, adverse selection, and intraday bid-ask spreads in futures markets," *Journal of Futures Markets*, 12, 519–538. 2
- MENKVELD, A. J. AND M. A. ZOICAN (2017): "Need for speed? Exchange latency and liquidity," *The Review of Financial Studies*, 30, 1188–1228. 6
- PAGNOTTA, E. AND T. PHILIPPON (2015): "Competing on speed," *Econometrica*, forthcoming. 2, 6
- ROSTEK, M. AND M. WERETKA (2015): "Information and strategic behavior," *Journal of Economic Theory*, 158, 536–557. 6
- SPULBER, D. F. (1996): "Market making by price-setting firms," *The Review of Economic Studies*, 63, 559–580. 5
- WOLINSKY, A. (1990): "Information revelation in a market with pairwise meetings," *Econometrica: Journal of the Econometric Society*, 1–23. 6

Appendix

A Omitted Proofs from Section 3

A.1 Optimal Prices

In this section, we derive the optimal prices set by dealers in type $k \in \{m, c\}$ meetings given current beliefs $\mu \in (0, 1)$.³⁰ In what follows, we will conjecture—and later confirm—that the trading thresholds are interior, i.e., that $\{\underline{\varepsilon}^k, \bar{\varepsilon}^k\}$ always lie in the interior of the support $[-e, e]$.

A.1.1 Monopoly Meetings

Consider the problem of a dealer with beliefs $\mu \in (0, 1)$ in a monopoly meeting. Given the uniform distribution, the expected profits of the dealer can be written as

$$\sum_{j \in \{l, h\}} \mu_j \left[\int_{-m}^m \frac{B - \omega - R_j(\tilde{\mu}_j^+(\mu, \omega)) + e}{2e} \frac{d\omega}{2m} \right] (v_j - B),$$

where

$$\tilde{\mu}_l^+(\mu, \omega) = \begin{cases} 0 & \text{if } \omega < -m + R_h(\mu) - R_l(\mu) \\ \mu & \text{else} \end{cases}$$

and

$$\tilde{\mu}_h^+(\mu, \omega) = \begin{cases} 1 & \text{if } \omega > m - [R_h(\mu) - R_l(\mu)] \\ \mu & \text{else.} \end{cases}$$

In words, the expected profits are the product of the probability of selling at price B (the first term in brackets) and the payoff conditional on selling (the second term). Since $p(\mu) = [R_h(\mu) - R_l(\mu)]/2m$ and $\mathbb{E}[\omega] = 0$, straightforward algebra reveals that expected profits can be re-written as

$$\sum_{j \in \{l, h\}} \mu_j \frac{B - r_j(\mu) + e}{2e} (v_j - B),$$

where $\mu_h \equiv \mu$, $\mu_l \equiv 1 - \mu$, and

$$r_j(\mu) = (1 - p(\mu)) R_j(\mu) + p(\mu) R_j(\mathbf{1}[j = h]).$$

The first-order condition of the dealer's expected profit with respect to the bid price B is then given by

$$\frac{\mathbb{E}_j v_j - B}{2e} - \mathbb{E}_j \frac{B - r_j + e}{2e} = 0,$$

where we've suppressed the argument of r_j for convenience. This yields the optimal bid price in a monopolist meeting,

$$B^m = \frac{\mathbb{E}_j v_j + \mathbb{E}_j r_j - e}{2}.$$

The optimal ask price in a monopolist meeting is constructed in an identical manner.

³⁰Optimal prices under full information are standard, and hence we omit the explicit derivation here.

A.1.2 Competitive Meetings

In competitive meetings, dealers' expected profits must equal zero, so that the bid price must satisfy

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

While this equation has two solutions, only the larger solution is consistent with Bertrand competition. To see why, let $B_1 < B_2$ denote the roots of this equation and suppose that the equilibrium bid price is B_1 . We claim that a dealer could deviate to $B_2 - \varepsilon$, for some $\varepsilon > 0$, and achieve strictly positive profits. To see why, note that profits are negative for all $B > B_2$. Since the equation above has only two roots, profits must be positive for some value of $B \in (B_1, B_2)$. Therefore, the equilibrium bid in competitive meetings must be the larger of the two roots, i.e.,

$$B^c = \frac{\mathbb{E}_j r_j + \mathbb{E}_j v_j - e + \sqrt{(\mathbb{E}_j r_j + \mathbb{E}_j v_j - e)^2 - 4\mathbb{E}_j v_j (r_j - e)}}{2}. \quad (46)$$

Using straightforward algebra, (46) can be re-written as

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

In order for B^c to be a valid (i.e., real) solution, we require the discriminant to be positive, or $e + \mathbb{E}_j (v_j - r_j) \geq 2\sqrt{\text{Cov}_j (v_j, r_j)}$. We will verify below that, in the equilibrium we characterize, this condition is indeed satisfied under Assumptions 1-3.

A similar derivation reveals that the ask price in competitive meetings must be

$$A^c = \frac{\mathbb{E}_j (r_j + v_j) + e - \sqrt{(\mathbb{E}_j (r_j - v_j) + e)^2 - 4\text{Cov}_j (v_j, r_j)}}{2}.$$

For this solution to be valid, we require

$$e + \mathbb{E}_j (r_j - v_j) \geq 2\sqrt{\text{Cov}_j (v_j, r_j)}$$

which we also verify is satisfied in the equilibrium we characterize under Assumptions 1-3.

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

Here we establish (33) from the main text, which states that in expectation—given dealers' beliefs—the valuations of dealers and traders are equal, i.e. $\mathbb{E}_j r_j = \mathbb{E}_j v_j$. We first show this when dealers are fully

informed about asset quality. For example, when $\mu = 1$, $\text{Cov}_j (r_j, v_j) = 0$ and thus

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

Since $r_h(1) = R_h(1)$, the expected reservation values may be written as

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

or, equivalently,

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

The unique solution is clearly $r_h(1) = v_h$. It is equally straightforward to show that $r_l(0) = v_l$.

As a result, for any μ ,

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

and, in expectation,

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

Now, suppose by way of contradiction that $\mathbb{E}_j r_j > \mathbb{E}_j v_j$, so that $\mathbb{E}_j R_j(\mu) > \mathbb{E}_j v_j$. Recall that the traders' reservation values are given by

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

Using the expressions for prices, we have

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

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

where

$$\begin{aligned} \Psi(\mu) &= \frac{1}{2} \sqrt{(\mathbb{E}_j (r_j - v_j) + e)^2 - 4 \text{Cov}_j (r_j, v_j)} \\ \Phi(\mu) &= \frac{1}{2} \sqrt{(\mathbb{E}_j (v_j - r_j) + e)^2 - 4 \text{Cov}_j (r_j, v_j)}. \end{aligned}$$

Note that $\mathbb{E}_j r_j > \mathbb{E}_j v_j$ implies $\Psi(\mu) > \Phi(\mu)$. Using the relationships between monopoly and competitive prices, we can write traders' reservation values as

$$R_j(\mu) = (1 - \delta) v_j + \delta r_j(\mu) + \delta \pi \alpha_m \frac{1}{2} \frac{v_j - r_j(\mu)}{2} + \delta \pi \alpha_c \frac{\Phi(\mu) + \Psi(\mu) + e}{2} \left(\frac{v_j - r_j(\mu) + \Phi(\mu) - \Psi(\mu)}{2} \right).$$

Taking expectation with respect to j yields have

$$\mathbb{E}_j R_j = (1 - \delta) \mathbb{E}_j v_j + \delta \mathbb{E}_j r_j + \delta \pi \alpha_m \frac{\mathbb{E}_j v_j - \mathbb{E}_j r_j}{4} + \delta \pi \alpha_c \left(\frac{\Phi(\mu) + \Psi(\mu) + e}{2} \right) \frac{\mathbb{E}_j v_j - \mathbb{E}_j r_j + \Phi(\mu) - \Psi(\mu)}{2}.$$

Since $\Phi(\mu) < \Psi(\mu)$ and $\mathbb{E}_j v_j < \mathbb{E}_j r_j$, the last two terms on the right-hand side are negative. Hence,

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

which implies

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

This inequality contradicts the initial assumption that $\mathbb{E}_j r_j > \mathbb{E}_j v_j$. A symmetric argument can be used to show that $\mathbb{E}_j r_j < \mathbb{E}_j v_j$ cannot hold. Hence, we must have $\mathbb{E}_j r_j = \mathbb{E}_j v_j$ and $\mathbb{E}_j R_j = \mathbb{E}_j v_j$.

A.3 Proof of Proposition 1

To prove that the equilibrium is characterized by a unique value, $p^*(\mu)$, consider first the difference in traders' reservations values. From (39)–(40), the difference in reservation values can be written

$$R_h - R_l = (1 - \delta) (v_h - v_l) + \delta (r_h - r_l) - \delta \pi \sum_k \alpha_k \frac{B^k - A^k + 2e}{2e} (r_h - r_l) \quad (47)$$

where, for ease of notation, we have suppressed the dependence of R_j , r_j , and p on μ . Next, we will show that the right-hand side of (47) can be expressed as a function of p , $R_h - R_l$, and the exogenous primitives, which we denote Ξ . Since $p = (R_h - R_l)/2m$, this is sufficient to establish that (47) is an equation in a single variable, p . Finally, we prove that for any μ , there is a unique $p^*(\mu)$ which satisfies (47).

From (34)–(37), we can write the difference in traders' reservation values as

$$R_h - R_l = (1 - \delta) (v_h - v_l) + \delta (r_h - r_l) - \delta \pi \alpha_m \frac{r_h - r_l}{2} - \delta \pi \alpha_c \frac{r_h - r_l}{2} \frac{e + \sqrt{e^2 - 4\text{Cov}(r_j, v_j)}}{e}. \quad (48)$$

Since $r_j = (1 - p) R_j + p v_j$,

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

Moreover, since $\mathbb{E}_j R_j = \mathbb{E}_j v_j$, tedious but straightforward algebra can be used to show that

$$\text{Cov}_j (R_j, v_j) = \mu (1 - \mu) (v_h - v_l) (R_h - R_l).$$

Substituting into the earlier expression for $\text{Cov}_j (r_j, v_j)$ yields

$$\begin{aligned} \text{Cov}_j (r_j, v_j) &= (1 - p) \mu (1 - \mu) (v_h - v_l) (R_h - R_l) + p \mu (1 - \mu) (v_h - v_l)^2 \\ &= \mu (1 - \mu) (v_h - v_l)^2 \left[(1 - p) \frac{(R_h - R_l)}{(v_h - v_l)} + p \right]. \end{aligned}$$

Re-arranging yields (42) in the main text.

Finally, we note that $r_h - r_l = (1 - p) (R_h - R_l) + p (v_h - v_l)$. The above expressions allow us to write (47) as an equation in p given by

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

Substituting $\alpha_c + \alpha_m = 1$ and re-arranging yields (41) in the main text.

Let $Z(p, \mu; \Xi)$ denote the right-hand side of (41). We now show that $Z(p, \mu; \Xi)$ has a unique solution on the interval $p \in [0, \frac{v_h - v_l}{2m}]$.³¹ To see this, note that

$$Z(0, \mu; \Xi) = (1 - \delta) (v_h - v_l) > 0$$

and

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

Hence, there exists a $p^* \in [0, \frac{v_h - v_l}{2m}]$ that solves (41). Defining the auxiliary function

$$g(p) = p (2m (1 - p) + v_h - v_l), \tag{49}$$

³¹Since $p = (R_h - R_l)/2m$, a solution to $Z(p, \mu; \Xi) = 0$ with $p \geq (v_h - v_l)/2m$ would imply $R_h - R_l \geq v_h - v_l$ which is a contradiction.

the derivative of Z with respect to p is

$$\begin{aligned}
Z_p &= -2m + \delta \left(1 - \frac{\pi}{2}\right) g'(p) \\
&\quad - \frac{\delta\pi\alpha_c}{2} \sqrt{1 - \frac{4}{e^2}\mu(1-\mu)(v_h - v_l)g(p)} g'(p) + \frac{\delta\pi\alpha_c}{4} \frac{g(p) \frac{4}{e^2}\mu(1-\mu)(v_h - v_l)g'(p)}{\sqrt{1 - \frac{4}{e^2}\mu(1-\mu)(v_h - v_l)g(p)}} \\
&= -2m + \delta \left(1 - \frac{\pi}{2}\right) g'(p) \\
&\quad - \delta\pi\alpha_c \frac{2g'(p) \left[1 - \frac{4}{e^2}\mu(1-\mu)(v_h - v_l)g(p)\right] - g(p) \frac{4}{e^2}\mu(1-\mu)(v_h - v_l)g'(p)}{4\sqrt{1 - \frac{4}{e^2}\mu(1-\mu)(v_h - v_l)g(p)}} \\
&= -2m + g'(p) \delta \left(1 - \frac{\pi}{2}\right) - \delta\pi\alpha_c \frac{g'(p) \left[2 - \frac{12}{e^2}\mu(1-\mu)(v_h - v_l)g(p)\right]}{4\sqrt{1 - \frac{4}{e^2}\mu(1-\mu)(v_h - v_l)g(p)}} \tag{50}
\end{aligned}$$

Note that

$$g'(p) = 2m(1 - 2p) + v_h - v_l.$$

From Assumption 2, we know that

$$\begin{aligned}
-2m + (v_h - v_l) \frac{\delta}{1 - \delta} < 0 &\Rightarrow -2m + (v_h - v_l) \frac{\delta \left(1 - \frac{\pi}{2}\right)}{1 - \delta \left(1 - \frac{\pi}{2}\right)} < 0 \\
&\Rightarrow -2m + (v_h - v_l + 2m) \delta \left(1 - \frac{\pi}{2}\right) < 0
\end{aligned}$$

Since $v_h - v_l + 2m = g'(0)$ and $g''(p) < 0$, it follows that, for all $p \in [0, \frac{v_h - v_l}{2m}]$,

$$-2m + g'(p) \delta \left(1 - \frac{\pi}{2}\right) < 0.$$

Assumption 2 also implies³²

$$\frac{12\mu(1-\mu)}{e^2} (v_h - v_l) g(p) \leq \frac{3}{e^2} (v_h - v_l)^2 < 2$$

and $g'(p) > 0$ for all $p \in [0, \frac{v_h - v_l}{2m}]$. Hence, the last term in (50) is negative and, therefore, $Z_p(p, \mu; \Xi) < 0$ for all $p \in [0, \frac{v_h - v_l}{2m}]$. We conclude that $Z(p^*, \mu; \Xi) = 0$ has a unique solution in $[0, \frac{v_h - v_l}{2m}]$. Note that since $R_j(\mu)$ is increasing in μ for $j \in \{l, h\}$, $R_h(\mu) \leq v_h$ and $R_l(\mu) \geq v_l$. This implies that $R_h(\mu) - R_l(\mu) \leq v_h - v_l$, which then implies that $p \leq \frac{v_h - v_l}{2m}$.

A.4 Verification of Equilibrium

First, we verify that the solutions to the quadratic equation that defines $p^*(\mu)$ are indeed real numbers.

Notice,

$$e^2 - 4\text{Cov}_j(r_j, v_j) = e^2 - 4\mu(1-\mu)(v_h - v_l)[2mp(1-p) + p(v_h - v_l)].$$

³²Notice, since $e \geq 2(v_h - v_l) + 2m$ and $m \geq (v_h - v_l)/2$, Assumption 2 implies $e \geq 3(v_h - v_l)$. Hence, $e^2 \geq (3/2)(v_h - v_l)^2$.

Since the negative expression is maximized at $\mu = 1/2$, $e^2 - 4\text{Cov}_j(r_j, v_j) \geq 0$ as long as

$$e^2 - (v_h - v_l)[2mp(1-p) + p(v_h - v_l)] \geq 0.$$

Since $p \leq (v_h - v_l)/2m$, this inequality holds as long as

$$e^2 - (v_h - v_l)^2 \geq 0$$

which is satisfied by Assumption 2.

Next, we verify our conjecture that traders' thresholds implied by the equilibrium bid and ask prices are interior, which fully verifies our earlier conjectures about the nature of the equilibrium. That is, we show that, for all $j \in \{l, h\}$, $k \in \{m, c\}$, and $\omega \in [-m, m]$,

$$A^k - \omega - R_j(\mu_j'(\mu, \omega)) \in [-e, e]$$

and

$$B^k - \omega - R_j(\mu_j'(\mu, \omega)) \in [-e, e].$$

Consider first the ask price and suppose $j = l$. Given the definition of $\tilde{\mu}_l^+(\mu, \omega)$, above, it is sufficient to prove that $A^k - \omega - R_l(0) \in [-e, e]$ for all $\omega \leq -m + R_h(\mu) - R_l(\mu)$, and $A^k - \omega - R_l(\mu) \in [-e, e]$ for all $\omega \geq -m + R_h(\mu) - R_l(\mu)$. Since $R_l(0) \leq R_l(\mu)$, it suffices to prove

$$A^k - \omega - R_l(0) \leq e \tag{51}$$

and

$$-e \leq A^k - \omega - R_l(\mu). \tag{52}$$

We first verify these inequalities for the monopoly price. Since $\mathbb{E}_j R_j = \mathbb{E}_j v_j$, we can write

$$A^m = \mathbb{E}_j R_j + \frac{e}{2}.$$

Since A^m is increasing in μ and $A^m = R_h(1) + e/2$, (51) holds if

$$R_h(1) + \frac{e}{2} + m - R_l(0) \leq e.$$

Since $R_h(1) = v_h$ and $R_l(0) = v_l$, this condition reduces to

$$v_h - v_l + m \leq \frac{e}{2},$$

which is guaranteed by Assumption 2.

To see that (52) holds, we first re-write this inequality using the ask price:

$$\mathbb{E}_j R_j + \frac{e}{2} - m - R_l(\mu) = \mu(R_h(\mu) - R_l(\mu)) - m + \frac{e}{2} \geq -e.$$

Hence, (52) holds if

$$\mu(R_h(\mu) - R_l(\mu)) \geq m - 3\frac{e}{2}.$$

Under Assumption 2,

$$m - 3\frac{e}{2} \leq -2m - 3(v_h - v_l) \leq 0.$$

Since $R_h(\mu) \geq R_l(\mu)$, Assumption 2 then also ensures (52) is satisfied.

We now verify (51) and (52) for the competitive price. Recall that the competitive price can be written as

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

where $\Psi(\mu) = \frac{1}{2}\sqrt{e^2 - 4\text{Cov}_j(r_j, v_j)}$. Since $A^c < A^m$, (51) necessarily holds. Substituting for the ask price in inequality (52) yields

$$\mu(R_h(\mu) - R_l(\mu)) - \Psi(\mu) \geq m - 3\frac{e}{2}.$$

Since $\Psi(\mu) \leq e/2$, it suffices to prove that $\mu(R_h(\mu) - R_l(\mu)) \geq m - e$. This result follows since $m - e < m - e/2$ and, under Assumption 2, $m - e/2 < 0$.

Consider next the ask price when $j = h$. Much like when $j = l$, to verify that traders' thresholds are interior in equilibrium, it suffices to prove that

$$A^k - \omega - R_h(\mu) \leq e \tag{53}$$

and

$$-e \leq A^k - \omega - R_h(1). \tag{54}$$

Substituting the monopoly ask price, (53) requires

$$-(1 - \mu)(R_h(\mu) - R_l(\mu)) + \frac{e}{2} - m \leq e$$

or

$$-\frac{e}{2} - m \leq (1 - \mu)(R_h(\mu) - R_l(\mu))$$

which is necessarily satisfied since $R_h(\mu) \geq R_l(\mu)$. Since A^m is smallest when $\mu = 0$, to verify (54) it suffices to prove that

$$v_l + \frac{e}{2} - \omega - v_h \geq -e$$

or

$$3\frac{e}{2} \geq v_h - v_l + m,$$

which is implied by Assumption 2. As for the competitive price, (53) is satisfied because $A^c < A^m$. Since the competitive price is minimized at $\mu = 0$ and equal to v_l , to verify (54) it suffices to prove that

$$v_l - m - v_h \geq -e$$

which is also implied by Assumption 2. A symmetric argument applies to the bid price, so we omit the proof.

A.5 Proof of Lemma 1

From (41), one can see that $Z(p, \mu; \Xi)$ 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 $\mu > 1/2$.

From the proof of Proposition 1 we know that $Z_p < 0$. Therefore, $\frac{dp^*}{d\mu} = -\frac{Z_\mu}{Z_p} > 0$ when $\mu < 1/2$ and < 0 when $\mu > 1/2$. Moreover, bid-ask spreads are given by

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

The right-hand side is increasing in $\mu(1 - \mu)$ and p^* . Since both of these are hump-shaped in μ with maximum at $\mu = 1/2$, so too 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.6 Proof of Proposition 2

We have

$$\begin{aligned} Z_\pi &= -\frac{\delta}{2} [(1 - p)2mp + (v_h - v_l)] \\ &\quad - \frac{\delta\alpha_c}{2} \sqrt{1 - \frac{4}{e^2}\mu(1 - \mu)(v_h - v_l)[(1 - p)2mp + (v_h - v_l)p] [(1 - p)2mp + (v_h - v_l)p]} \\ &< 0. \end{aligned}$$

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

A.7 Proof of Proposition 3

From the text, we have

$$\begin{aligned} \frac{\partial}{\partial \pi} Y_{t,j}^c &= -t(1-p^*)^{t-1} \frac{\partial p^*}{\partial \pi} \left(1 - \sqrt{1 - \frac{4}{e^2} (v_h - v_l) \mu (1 - \mu) g(p^*)} \right) \\ &\quad + (1-p^*)^t \frac{\frac{4}{e^2} (v_h - v_l) \mu (1 - \mu) g'(p^*) \frac{\partial p^*}{\partial \pi}}{\sqrt{1 - \frac{4}{e^2} (v_h - v_l) \mu (1 - \mu) g(p^*)}} \\ &= (1-p^*)^t \frac{\partial p^*}{\partial \pi} \left[\frac{\frac{4}{e^2} (v_h - v_l) \mu (1 - \mu) g'(p^*)}{\sqrt{1 - \frac{4}{e^2} (v_h - v_l) \mu (1 - \mu) g(p^*)}} - \frac{t}{1-p^*} \left(1 - \sqrt{1 - \frac{4}{e^2} (v_h - v_l) \mu (1 - \mu) g(p^*)} \right) \right] \end{aligned}$$

The expression in the brackets is negative for sufficiently large values of t . Hence, since $\frac{\partial p^*}{\partial \pi} < 0$, $\frac{\partial}{\partial \pi} Y_{t,j}^c > 0$, is also negative for sufficiently large values of t .

A.8 The Effect of Changes in α_c .

A.8.1 Proof of Proposition 4

Since $Z_p < 0$ and

$$Z_{\alpha_c} = -\frac{\delta \pi}{2} \sqrt{1 - \frac{4}{e^2} \mu (1 - \mu) (v_h - v_l) [(1-p) 2mp + (v_h - v_l) p] [(1-p) 2mp + (v_h - v_l) p]} < 0,$$

it follows that $\frac{dp^*}{d\alpha_c} = -\frac{Z_{\alpha_c}}{Z_p} < 0$.

A.8.2 Proof of Proposition 5

Using the results derived above, along with the function $g(p)$ introduced in (49), we can write the competitive spread as

$$Y_{t,j}^c = (1-p(\mu_0))^t \left[e - \alpha_c \sqrt{e^2 - 4(v_h - v_l) \mu_0 (1 - \mu_0) g(p(\mu_0))} \right].$$

Differentiating with respect to α_c , we have

$$\begin{aligned} \frac{\partial}{\partial \alpha_c} Y_{t,j}^c &= -t(1-p(\mu_0))^{t-1} \frac{dp(\mu_0)}{d\alpha_c} \left[e - \alpha_c \sqrt{e^2 - 4(v_h - v_l) \mu_0 (1 - \mu_0) g(p(\mu_0))} \right] \\ &\quad - (1-p(\mu_0))^t \left[\sqrt{e^2 - 4(v_h - v_l) \mu_0 (1 - \mu_0) g(p(\mu_0))} + \frac{1}{2} \alpha_c \frac{-4(v_h - v_l) \mu_0 (1 - \mu_0) g'(p(\mu_0)) \frac{dp(\mu_0)}{d\alpha_c}}{\sqrt{e^2 - 4(v_h - v_l) \mu_0 (1 - \mu_0) g(p(\mu_0))}} \right]. \end{aligned}$$

Since $dp(\mu_0)/d\alpha_c < 0$, the first term is strictly positive. Thus, as in the proof of Proposition 3, the first term dominates and ensures $\partial Y_{t,j}^c / \partial \pi > 0$ for sufficiently large t .

A.8.3 The Effect of α_c on Average Spreads

Before proceeding, we consider the effect of an increase in α_c on the average quoted spread. This spread satisfies

$$Y_{t,j} \equiv e(1 - \alpha_c) + \alpha_c (1 - p(\mu_0))^t \left[e - \alpha_c \sqrt{e^2 - 4\text{Cov}(r_j, v_j)} \right].$$

Differentiating with respect to α_c yields

$$\begin{aligned} \frac{\partial Y_{t,j}}{\partial \alpha_c} &= -e + (1 - p(\mu_0))^t \left[e - \alpha_c \sqrt{e^2 - 4\text{Cov}(r_j, v_j)} \right] \\ &\quad - \alpha_c t (1 - p(\mu_0))^{t-1} \frac{\partial p}{\partial \alpha_c} \left[e - \alpha_c \sqrt{e^2 - 4\text{Cov}(r_j, v_j)} \right] \\ &\quad - \alpha_c (1 - p(\mu_0))^t \sqrt{e^2 - 4\text{Cov}(r_j, v_j)} \\ &\quad + 2\alpha_c^2 (1 - p(\mu_0))^t \left[e^2 - 4\text{Cov}(r_j, v_j) \right]^{-\frac{1}{2}} \frac{\partial \text{Cov}(r_j, v_j)}{\partial \alpha_c}. \end{aligned}$$

Since $t(1 - p(\mu_0))^{t-1} \rightarrow 0$ as $t \rightarrow \infty$, spreads are necessarily decreasing as $t \rightarrow \infty$. We then have an opposite result of Corollary 1 with respect to α_c . That is, there exists a $\tau < \infty$ such that $dY_{t,j}/d\alpha_c < 0$ for all $t \geq \tau$.

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 $\Xi_2 = (R_l(\mu) + m + e, R_h(\mu) + m + e)$.³³ 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.

³³The argument for the ask price is symmetric.

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.

C Additional Numerical Results

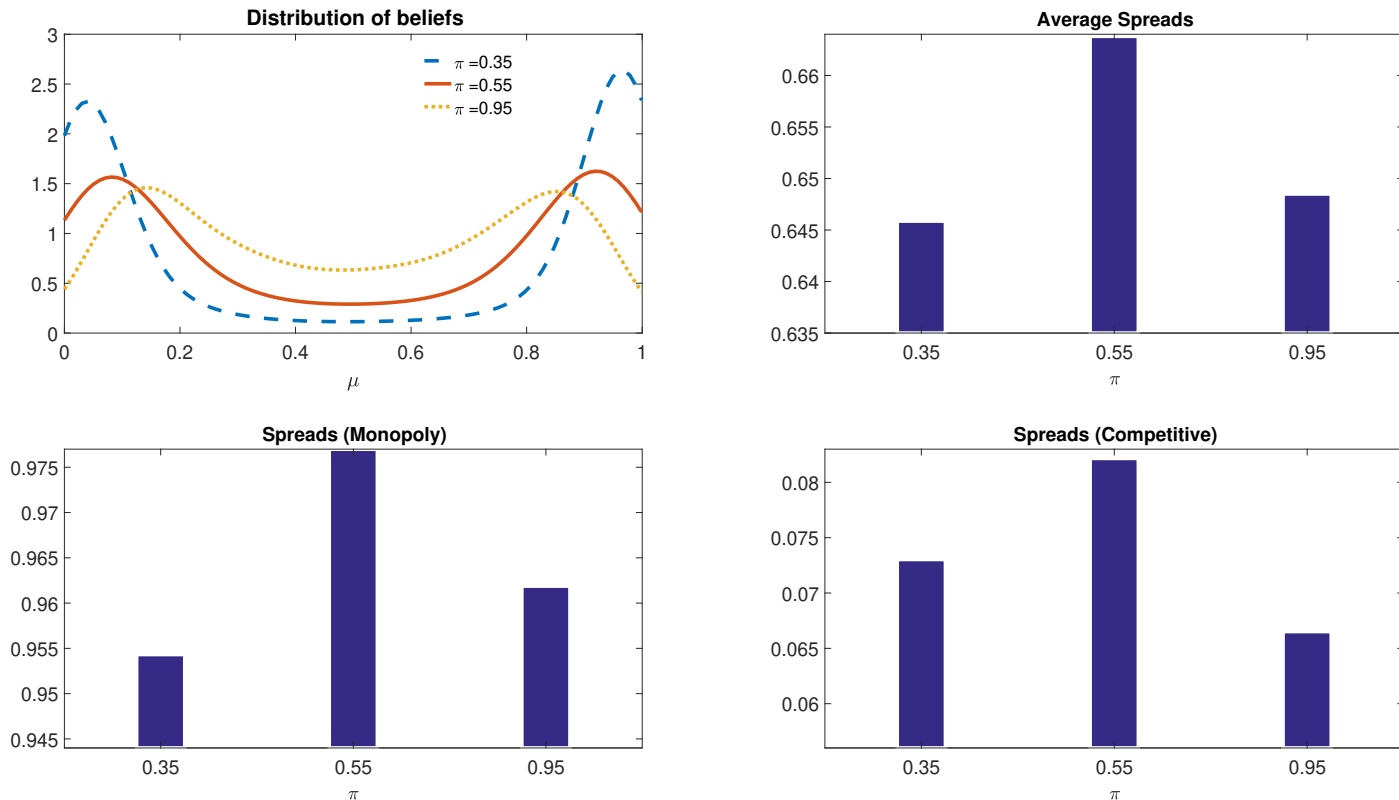


Figure 7: Effect of π , high σ_ω^2 .

C.1 Robustness

In our baseline results, we assumed that aggregate and idiosyncratic liquidity shocks were of equal magnitude. Here, we present results under two alternative assumptions. In the first, we assume that the variance of aggregate liquidity shocks is twice as high as that of the idiosyncratic component. Holding the total variance fixed at its baseline value of 0.32, this implies $\sigma_\omega^2 = 0.21$ and $\sigma_\varepsilon^2 = 0.11$. The results for the stationary version of the model are plotted in Figure 7. It shows the same patterns as in Figure 5. Namely, higher values of π slow down learning, resulting in greater likelihood of intermediate beliefs in the stochastic steady state. Spreads are non-monotonic in π .

Figure 8 repeats the analysis under the assumption that the variance of the idiosyncratic component is twice as large as that of the aggregate one, i.e., $\sigma_w^2 = 0.11$ and $\sigma_\varepsilon^2 = 0.21$. Now, spreads in the competitive meetings are hump-shaped but those under monopoly (as well as the average spread) are *increasing*.

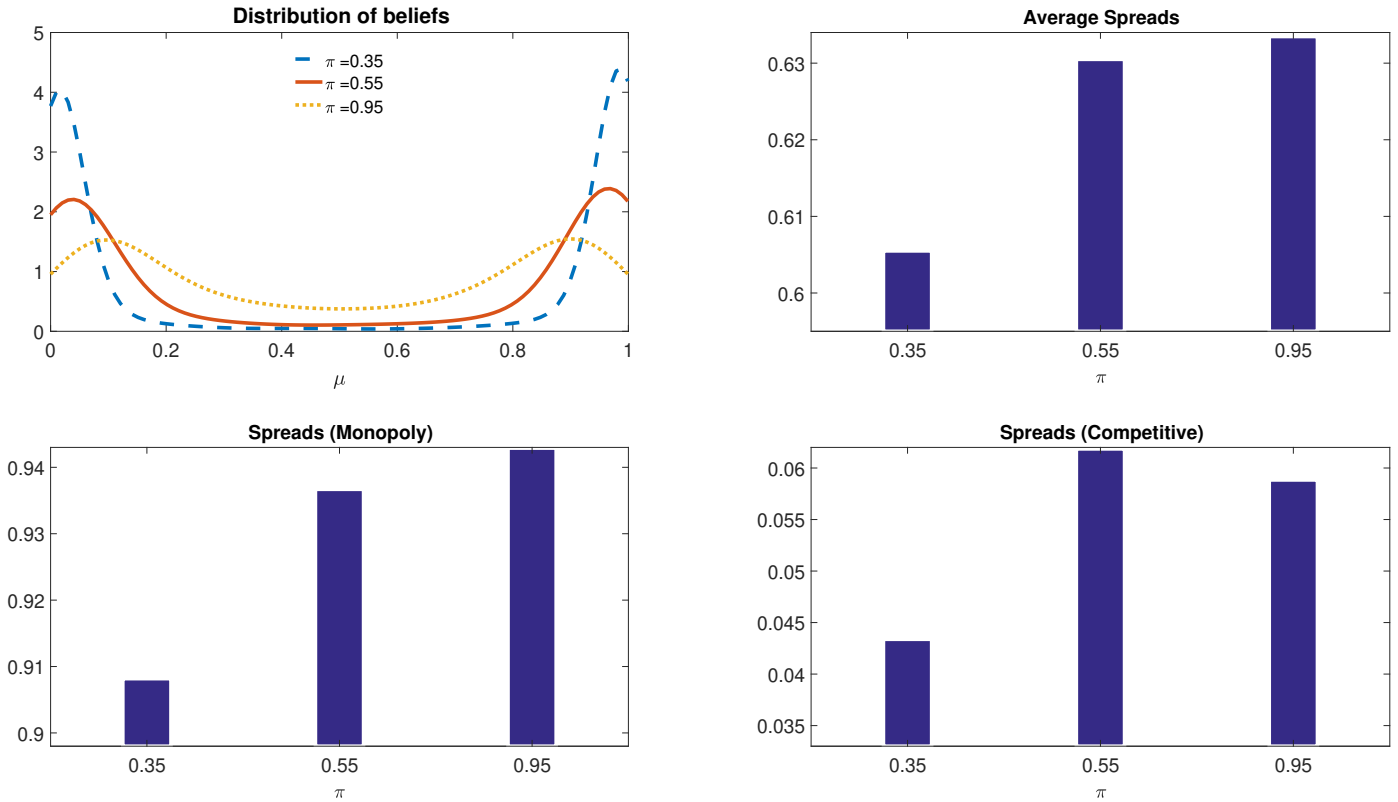


Figure 8: Effect of π , low σ_w^2 .

C.2 The Effect of Additional Information

Here, we demonstrate, by way of numerical example, that reducing asymmetric information could also have counter-intuitive predictions in our environment. In particular, we retain the parameters from our baseline calibration with one exception: we set v_h to a slightly higher number, namely 1.28.³⁴ Figure 9 plots bid-ask spreads, both average (left panel) and in monopolist dealer meetings (right panel).

Importantly, note that bid-ask spreads are convex in μ in some regions of the parameter space. In these regions, providing dealers with additional information—through the form of a noisy public signal about asset quality—would induce a mean-preserving spread in μ . As a result, by Jensen’s inequality, average bid-ask spreads would widen. Hence, in these regions of the parameter space, providing dealers with additional information leads to wider spreads, on average, even though it mitigates adverse selection.

³⁴This corresponds to the price change on a 7-year par bond with a coupon of 2%, when the issuer is downgraded from AAA to AA (recall that the downgrade raises the yield on the bond by 25 bps).





Figure 9: Spreads as a function of beliefs, $v_h = 1.28$.