Liquidity and Information in Limit Order Markets

Ioanid Roşu*

February 26, 2019

Abstract

How does informed trading affect liquidity in limit order markets, where traders can choose between market orders (demanding liquidity) and limit orders (providing liquidity)? In a dynamic model, informed trading overall helps liquidity: a higher share of informed traders (i) improves liquidity as proxied by the bid-ask spread and market resiliency, and (ii) has no effect on the price impact of orders. The model generates other testable implications, and suggests new measures of informed trading.

JEL: C73, D82, G14

KEYWORDS: Limit order book, volatility, trading volume, slippage, informed trading, stochastic game.

*Roşu (corresponding author), rosu@hec.fr, HEC Paris. I thank an anonymous referee, Hendrik Bessembinder (the editor), Peter DeMarzo, Doug Diamond, Thierry Foucault, Johan Hombert, Peter Kondor, Juhani Linnainmaa, Stefano Lovo, Christine Parlour, Uday Rajan, Talis Putnins, Pietro Veronesi, and finance seminar participants at Chicago Booth, Stanford University, University of California at Berkeley, University of Illinois Urbana-Champaign, University of Toronto (Dept. of Economics), Bank of Canada, HEC Lausanne, HEC Paris, University of Toulouse, Ecole Polytechnique, Tilburg University, Erasmus University, Insead, and Cass Business School, for helpful comments and suggestions. I am also grateful to conference participants at the 2010 Western Finance Association meetings, 2010 European Finance Association meetings, NBER microstructure meeting, 4th Central Bank Microstructure Workshop, and the 1st Market Microstructure Many Viewpoints Conference in Paris.

I. Introduction

Market liquidity is a central concept in finance, in particular in relation with asset pricing.[1] According to Bagehot (1971), illiquidity is caused by asymmetric information, via the actions of liquidity providers. The liquidity provider, or market maker, which Bagehot identifies as the "exchange specialist in the case of listed securities and the over-the-counter dealer in the case of unlisted securities," sets prices and spreads so that on average he makes losses from traders who possess superior information, but compensates with gains from uninformed traders, who are motivated by liquidity needs or simply trade on noise. Thus, the stronger the asymmetric information between the informed traders and the market maker, the larger the bid-ask spread needs to be so that the market maker at least breaks even. A large theoretical literature has since made Bagehot's intuition rigorous.[2]

Following Bagehot (1971), most of the theoretical literature assumes that liquidity providers do not to possess any superior fundamental information.[3] More recent evidence, however, has called into question this assumption. One reason is that most financial exchanges around the world have become "limit order markets," meaning that any investor (informed or not) can provide liquidity by posting orders in a limit order book.[4] Moreover, empirical evidence shows that there is an important premium for liquidity provision in limit order markets, and that informed traders do indeed use limit orders extensively.[5] Despite the evidence, the literature has been largely silent on the order choice problem of informed traders, and, importantly, on how this choice affects market liquidity. The goal of the present paper is to fill this gap.

To address these questions, consider a dynamic model of a limit order market. Risk-

[1]See Amihud and Mendelson (1986), Brennan and Subrahmanyam (1996), Easley, Hvidkjaer, and O'Hara (2002), Pástor and Stambaugh (2003), Acharya and Pedersen (2005).

[2]See Kyle (1985), Glosten and Milgrom (1985), or O'Hara (1995) and the references within.

[3]Notable exceptions are Chakravarty and Holden (1995), Kaniel and Liu (2002), Goettler, Parlour, and Rajan (2009), and Brolley and Malinova (2017).

[4]Nowadays, most equity and derivative exchanges are either pure limit order markets (Euronext, Helsinki, Hong Kong, Tokyo, Toronto); or hybrid markets, in which designated market makers must compete with a limit order book (NYSE, Nasdaq, London). See Jain (2005).

[5]See Biais, Hillion, and Spatt (1995), Harris and Hasbrouck (1996), Griffiths, Smith, Turnbull, and White (2000), Sandås (2001), Hollifield, Miller, and Sandås (2004), Anand, Chakravarty, and Martell (2005), Menkhoff, Osler, and Schmeling (2010), Latza and Payne (2011), Hautsch and Huang (2012). Similar findings are reported by Bloomfield, O'Hara, and Saar (2005) in the context of experimental asset markets.

neutral investors arrive randomly to the market and trade in one risky asset. The asset's fundamental value is time varying, and information about it is costly to acquire and process.[6] Informed investors learn the current value of the asset, and decide whether to buy or sell one unit of the asset, and whether to trade with a market order or a limit order. Limit orders can subsequently be modified or cancelled without any cost

The main result is that a larger fraction of informed traders overall improves liquidity. This result is driven by two key features of the model: First, there is competition among informed traders, in the sense that each informed trader must take into account the future arrivals of other informed traders. Second, private information is long-lived, as information about the fundamental value is revealed to the public only via the order flow.[7] Because of these features, a larger share of informed traders produces a dynamic efficiency that can eventually overcome the static increase in adverse selection. To understand in more detail the intuition behind the main result, I briefly describe several key equilibrium results.

The first key result describes the optimal order choice of the informed trader. This is essentially a threshold strategy: the informed trader (referred to in the paper as "she") submits either a market order or a limit order, depending on the magnitude of her privately observed "mispricing," which is the difference between the fundamental value (privately observed) and the "public mean" (the public expectation of the fundamental value). An extreme mispricing causes the informed trader to submit a market order, while a moderate mispricing causes a limit order. This result formalizes an intuition present for instance in Harris (1998), Bloomfield et al. (2005), Hollifield et al. (2004), Large (2009).

The second key result describes the information content of the order flow. Because in equilibrium informed traders can submit both limit orders and market orders, all types of order have "price impact" (defined as the change in public mean caused by the order). Nevertheless, because market orders are associated to more extreme mispricing, the price impact of a buy market order is larger in magnitude (about 4 times larger

[6]Because I am interested in long run liquidity effects, I assume that the asset value is not constant, but follows a random walk. Thus, prices do not eventually reveal all the private information. In Goettler et al. (2009), the fundamental value is also time varying, but follows a Poisson process.

[7]Goettler et al. (2009) obtain different results because in their model the private information is short-lived (the fundamental value is revealed publicly after several periods).

in my model) than the price impact of a buy limit order. In line with this prediction, Hautsch and Huang (2012, p.515) find empirically that market orders have a permanent price impact of about four times larger than limit orders of comparable size.

The third key result describes the equilibrium bid-ask spread, and identifies a new component of this spread: the slippage component. I define "slippage" as the tendency of an informed trader's estimated mispricing to decay over time.[8] Slippage is due to the future arrival of other informed traders who correct the mispricing by submitting their orders. Thus, slippage induces an endogenous waiting cost for the informed trader, called the "slippage cost." In addition, the informed trader suffers from an "adverse selection cost," since at the time of order execution she is potentially less informed than the future informed traders.[9] I define the "decay cost" as the sum of the slippage cost and the adverse selection cost.

The decay cost generates a tradeoff between limit orders and market orders: by trading with a limit order, an informed trader gains half the bid-ask spread, but loses from the decay cost. By trading with a market order instead, the informed trader loses half the bid-ask spread, but pays no decay cost. At the threshold mispricing, the informed trader is indifferent between a market order and a limit order. Hence, the decay cost corresponding to this threshold value is equal to the equilibrium bid-ask spread. From the definition of the decay cost, the bid-ask spread is therefore the sum of a "slippage component" and an "adverse selection component." To my knowledge, the slippage component is new to the literature. Huang and Stoll (1997) decompose the bid-ask spread into order processing costs, adverse selection costs, and inventory holding costs. In my model, I abstract away from inventory issues and order processing costs, but recover the adverse selection component. In addition, however, by allowing informed traders to provide liquidity, the phenomenon of slippage generates a new component of

[8] According to Investopedia, "slippage happens when a trader gets a different [price] than expected between the time he enters the trade and the time the trade is made" (February 22, 2019, available at: `https://www.investopedia.com/terms/s/slippage.asp`). Thus, slippage can also occur if a large, possibly uninformed market order "walks the book," that is, if some units (parts) of the order execute at a worse price. In this paper, slippage applies only to limit orders submitted by informed traders, and it occurs even when limit orders are for just one unit.

[9] This is because the informed trader acquires information only when she enters the market. If instead she continuously observes the fundamental value, the adverse selection component is zero, but the slippage cost is still positive, as competition with future informed traders gradually erodes her initial information advantage.

the bid-ask spread.

The main result describes how liquidity is affected by the fraction, or share of informed traders, henceforth called the "informed share". Surprisingly, a larger informed share overall has a positive effect on liquidity. More precisely, a larger informed share has (i) a negative effect on bid-ask spreads; (ii) no effect on the price impact; and (iii) a strongly positive effect on market resiliency, which is defined in Kyle (1985) as the speed with which prices recover from a random, uninformative shock. Moreover, a larger informed share has a positive effect on market efficiency by reducing the "public volatility." The latter is defined as the publicly inferred volatility of the fundamental value, hence its inverse is a measure of dynamic efficiency: when the public volatility is small, the public has precise information about the fundamental value.

To get intuition for the main result, note that a larger informed share implies that the informed traders exert more pressure on prices to revert to the fundamental value. This explains the strong positive effect of the informed share on market resiliency. Also, it explains the negative effect of the informed share on public volatility: when there are more informed traders, the public eventually learns better about the fundamental value, and the public volatility decreases. But the bid-ask spread is equal to the decay cost corresponding to the threshold mispricing. When the public volatility is smaller, the decay cost is also smaller because the average mispricing tends to be smaller. Hence, a larger informed share generates a smaller bid-ask spread.

To understand the neutral effect of the informed share on market depth, suppose the informed share is small, and a buy market order arrives. There are two opposite effects at play. First, when the informed share is small, it is unlikely that the market order comes from an informed trader. This effect decreases the price impact. But, second, if the buy market order does come from an informed trader, she must have observed a fundamental value far above the public mean; otherwise, knowing there is little competition from other informed traders, she would have submitted a buy limit order. This effect increases the price impact. The two effects exactly offset each other.[10]

The results described thus far are obtained in the "stationary equilibrium," in which the public volatility is constant over time (which in turn makes the bid-ask spread

[10]This is proved rigorously in Proposition 1, and explained in the subsequent discussion.

and price impact also constant). In the stationary equilibrium, the natural increase in uncertainty due to changes in the fundamental value is exactly offset by the new information contained in the order flow. The final set of results arise from the study of "non-stationary equilibria," which can appear for instance after an uncertainty shock (an unobserved shock to the fundamental value) results in a temporary spike in public volatility.

Liquidity is "resilient:" after an uncertainty shock, the bid-ask spread and price impact (as well as the public volatility) decrease over time to their values in the stationary equilibrium. The bid-ask spread and price impact are both increasing in the size the uncertainty shock. The liquidity resiliency is larger when there are more informed traders, as the order flow becomes more informative. Liquidity resiliency is different from market resiliency, as the latter is the tendency of prices to revert to the fundamental value after an uninformative shock.

I introduce a new measure, the "market-to-limit probability ratio," which is the defined as the probability the next order is a market order, divided by the probability that the next order is a limit order. This number is equal to one in the stationary equilibrium, but after an uncertainty shock the market-to-limit probability ratio drops to levels significantly less than one, as the increase in the bid-ask spread temporarily prompts the informed traders to submit more limit orders. The connections among the market-to-limit probability ratio with the liquidity measures and the public volatility, as well as the expected evolution of the equilibrium towards the stationary one, produce new testable implications of the model.

Overall, the theoretical model produces a rich set of implications regarding the connection between the activity of informed traders and the level of liquidity. The main result is that informed traders have on aggregate a positive effect, by making the market more efficient and, at the same time, more liquid. A welfare analysis in Section 3 in the Internet Appendix also shows that a larger number of informed traders (caused for instance by an exogenous decrease in information costs) increases aggregate trader welfare. The model thus provides useful tools to analyze informed trading, and its connection with liquidity, prices, and welfare.

This paper is part of a growing theoretical literature on price formation in limit order markets.[11] Of central interest in this literature is how investors choose between demanding liquidity via market orders and supplying liquidity via limit orders.[12] Several papers, such as Foucault et al. (2005), or Roşu (2009) study order choice by assuming that investors have exogenous waiting costs. One advantage of my model is that waiting costs arise endogenously in the case of an informed investor: these are the decay costs described above.

Goettler et al. (2009) is the first paper that solves a dynamic model of limit order markets with asymmetric information. The focus of their paper is different, however. While I am interested in the effect of informed trading on liquidity, Goettler et al. (2009) analyze the interplay between information acquisition, order choice and volatility. They find that under picking off risks—which are absent in my model—different volatility regimes affect traders' order choice, and make the market act as a volatility multiplier. Moreover, there are two important modeling differences. First, in their model private information is short-lived, because the fundamental value is publicly revealed after several periods. This assumption reduces the effect of dynamic efficiency in their model, as informed traders cannot arrive more quickly to make the market more efficient. By contrast, in my model dynamic efficiency has a strong effect by having private information being incorporated over the long run, and as a result the informed traders have an overall positive effect on liquidity. Second, in their model traders do not continuously monitor the market, which creates stale limit orders and picking off risks. In my model, there are no stale orders since limit orders can be modified instantly.

My main result, that informed trading have a positive overall effect on liquidity, is documented by several empirical papers, starting with Collin-Dufresne and Fos (2013). They find that the bid-ask spread and realized price impact decrease in the presence of informed trading coming from corporate insiders.[13] In my model, I obtain an improvement in the bid-ask spread, but not in the price impact. This latter fact might be due to

[11]See Glosten (1994), Parlour (1998), Foucault (1999), Foucault, Kadan, and Kandel (2005), Goettler, Parlour, and Rajan (2005, 2009), Back and Baruch (2007), Roşu (2009), Biais, Hombert, and Weill (2013), Pagnotta (2010), and the survey by Parlour and Seppi (2008).

[12]For models of order choice without private information, see Cohen, Maier, Schwartz, and Whitcomb (1981), Harris (1998), Foucault (1999), Parlour (1998), Goettler et al. (2005), Roşu (2009).

[13]Their interpretation is based on Admati and Pfleiderer's (1988) intuition that insiders trade more aggressively in periods when they expect noise trading activity to increase. At those times, liquidity is higher, despite the increase in adverse selection coming from informed trading.

the fact that my measure of price impact is instantaneous, while their empirical measure is considered over a longer period, and thus may be affected by market resiliency. Roşu (2019) extends the Glosten and Milgrom (1985) model to allow a moving fundamental value, and finds that the informed share has no effect on the bid-ask spread. In that paper, however, the ask and bid prices are not limit order prices, but rather quote prices, set by a risk-neutral specialist. As a result, the half bid-ask spread is the same as the price impact of a buy order, which, as in the present paper, is not affected by the informed share.

The paper is organized as follows. Section II describes the model. Section III solves for the stationary equilibrium, in which the public volatility (as well as the bid-ask spread and price impact) is constant. Section IV describes the properties of the stationary equilibrium, including the various dimensions of liquidity and information efficiency. Section V explores non-stationary equilibria of the model. Section VI concludes. Proofs of the main results are in the Appendix and the Internet Appendix. The companion Internet Appendix contains additional results and robustness checks.

II. Model

The market consists of a single risky asset. Time is continuous, and traders arrive randomly to the market. After deciding whether to acquire private information regarding the fundamental value of the asset, traders can submit an order to buy or sell one unit of the asset. Traders also choose the price at which they are willing to transact. If an order does not execute, it can be subsequently modified or cancelled. Information can be difficult to process, in a sense made precise below.

A. Trading and Prices

Trading occurs when a buy or sell order is executed against an order of the opposite type. Each order is a limit order, as it specifies a quantity and a price beyond which the trader is no longer willing to transact. The price can be any real number. Limit orders are subject to price priority: buy orders submitted at higher prices and sell orders submitted at lower prices have priority. Limit orders submitted at the same price are subject to time priority: the earlier order is executed first. If several orders arrive at the same time, priority is assigned randomly to them.[14]

The "limit order book" is the collection of all outstanding limit orders (submitted but not yet executed or cancelled). In the book, limit orders form two queues, based on order priority: the "ask queue" on the sell side, and the "bid queue" on the buy side. The lowest price on the ask side is the "ask price," or simply the ask. The highest price on the bid side is the "bid price," or simply the bid. A marketable limit order is a buy limit order with a price above the ask, or a sell limit order below the bid. A marketable limit order is executed immediately and is henceforth called a "market order."

B. Traders and Arrivals

Traders arrive to the market according to a Poisson process with parameter λ. Immediately after arrival, a trader chooses whether to (a) submit a market order, (b) submit a limit order, or (c) submit no order at all. Each order is for one unit of the asset. After submission, a limit order can be either (i) modified, which means the limit price is changed—in which case time priority is lost, or (ii) cancelled. As soon as the order is executed or cancelled, or if no order is submitted, the trader exits the model.

Traders are risk-neutral but their utility also includes a private valuation component and a cost from waiting.[15] Each trader has a type (u, r), which consists of a private valuation u for the asset and a waiting coefficient r. The private valuation u can take three possible values, $\{-\bar{u}, 0, \bar{u}\}$, where $\bar{u} > 0$. A trader is a "natural buyer" if $u = \bar{u}$, a "natural seller" if $u = -\bar{u}$, or "speculator" if $u = 0$. At time t, the instantaneous utility of a trader with private valuation u is:

$$\begin{cases} v_t - P_t + u, & \text{if trader buys at } t, \\ P_t - v_t - u, & \text{if trader sells at } t, \\ 0, & \text{if trader's order does not execute at } t, \end{cases} \tag{1}$$

where v_t is the fundamental value at t, and P_t is the transaction price at t. Traders

[14]With Poisson arrivals, the probability of two or more traders arriving at the same time is zero.

[15]The private valuation can arise from liquidity or hedging needs, or from bias regarding the asset (optimism or pessimism). The waiting cost can arise from trading horizon/deadlines, or from uncertainty regarding future order execution.

incur a waiting cost of the form $r \times \tau$, where τ is the expected waiting time, and r is a constant coefficient. The waiting coefficient r can take two possible values, $\{0, \bar{r}\}$, where $\bar{r} > 0$. A trader is "patient" if $r = 0$, or "impatient" if $r = \bar{r}$.

To simplify presentation, I assume that (i) impatient natural buyers always submit a buy market order, (ii) impatient natural sellers always submit a sell market order, and (iii) impatient speculators do not submit any order. In Section 2 in the Internet Appendix, I show that (i)-(iii) are equilibrium results if $\bar{u}$ and $\bar{r}$ are above certain thresholds.[16] Since traders who submit no order exit the model immediately, I replace (iii) by the assumption that all speculators are patient.

Natural buyers and sellers (traders with valuation $\bar{u}$ or $-\bar{u}$) arrive randomly to the market according to an independent Poisson process with parameter λ^u. They are equally likely to have positive or negative private valuation, and equally likely to be patient or impatient. Patient speculators arrive randomly to the market according to an independent Poisson process with parameter λ^i. The total trading activity is $\lambda = \lambda^u + \lambda^i$. The "informed share" is defined as the ratio:

$$\rho = \frac{\lambda^i}{\lambda^i + \lambda^u}. \tag{2}$$

Thus, ρ is the fraction of traders who are speculators, and $1 - \rho$ is the fraction of traders who are natural buyers or sellers (patient or impatient).

C. Information

At any time t, the asset has a fundamental value v_t, also called common value or full-information price. The asset value follows a diffusion process $\mathrm{d}v_t = \sigma_v \mathrm{d}B_t$, where B_t is a standard Brownian motion, and the "fundamental volatility" parameter σ_v is a positive constant. Because traders arrive to the market according to a Poisson process, inter-arrival times are exponentially distributed with mean $1/\lambda$. For simplicity of notation, throughout the paper I work in event time rather than calendar time: if a trader arrives

[16] In particular, I show that it is not profitable for a sufficiently impatient speculator to acquire information. Ex post, i.e., after seeing the signal, such a speculator might observe an extreme mispricing that could be exploited without waiting, and would therefore justify the information cost, but ex ante such signals are rare and therefore do not justify the cost.

at t, the next trader arrives at $t+1$.[17] The discrete version of the fundamental value process in event time is:

$$(3) \qquad v_t \;=\; v_{t-1} + \sigma_I \varepsilon_t, \quad \text{with} \quad \sigma_I \;=\; \frac{\sigma_v}{\sqrt{\lambda}} \quad \text{and} \quad \varepsilon_t \;\sim\; \mathcal{N}(0,1),$$

where σ_I is the "inter-arrival volatility," and ε_t has the standard normal distribution.

By paying an information acquisition cost, a trader learns the fundamental value at the time of arrival.[18] To simplify presentation, I assume that all patient speculators acquire information, and that no other traders acquire information; this is proved as an equilibrium result in Section 3 in the Internet Appendix. In what follows, I refer to the patient speculators as "informed traders," and to the natural buyers and sellers as "uninformed traders."

All traders observe the history of the game. The history consists of the whole order flow, i.e., submissions, executions, modifications, and cancellations. The evolution of the limit order book and the transaction prices are part of this public information. A trader's type (private valuation and waiting coefficient) is private information for each trader. The fundamental value at the time of arrival is private information for each informed trader.

D. Equilibrium Concept

The model represents a stochastic game, in which Nature moves by drawing randomly new traders at each time $t \in \mathbb{N} = \{0, 1, 2, ...\}$. After traders arrive and decide whether to become informed or not, they engage in a trading game and at each time maximize their expected utility given their information set. Even though the arrivals occur at discrete points in time, traders can later modify their orders at any time in between. The game is therefore set in continuous time, and I use the framework of Bergin and MacLeod (1993) in which traders can react instantly.

The equilibrium concept is Markov Perfect Equilibrium (MPE), as defined for in-

[17] This use of event time has been justified empirically for instance by Hasbrouck (1993). Equivalently, I set the model in discrete time, in which case $t+1$ is replaced by $t+\frac{1}{\lambda}$.

[18] Learning only at arrival is consistent with the assumption below that the informed trader who submits a limit order must use an uninformed trader (broker) to update the limit order until it is executed.

stance in Fudenberg and Tirole (1991). As a refinement of the Perfect Bayesian Equilibrium (PBE) concept, a MPE is defined by a "game assessment," which is the collection of a strategy profile and a belief system such that (i) at every stage of the game, strategies are optimal given the beliefs, and the beliefs are obtained from equilibrium strategies and observed actions using Bayes' rule, and (ii) the game assessment is conditional on a set of state variables which are payoff-relevant. The latter condition implies that in a MPE there are no ad-hoc punishments to support the equilibrium.

E. Information Processing

Solving the model described above is very challenging if traders can do full Bayesian updating. This is because each trader's inference problem involves an infinite number of state variables, which are the moments of the probability density that describes the trader's belief about the fundamental value. As new orders arrive, the belief must be updated based the information contained in each order type. But because informed traders use threshold strategies (see Theorem 1), the update of the density changes its shape in ways which are difficult to quantify precisely.

The modeling approach is to introduce frictions in information processing such that the traders solve a simplified inference problem.[19] These frictions are based on the principle that it is more difficult to process (i) private rather than public information, (ii) conditional rather than unconditional information, and (iii) higher rather than lower moments of a distribution. But rather than explicitly introducing information processing costs, I directly specify what information traders can process.

When updating the belief density, an uninformed trader can compute without cost (i) the first moment of the posterior belief conditional on order type, and (ii) the second moment of the posterior belief conditional on order arrival, but unconditional of order type.[20] Uninformed traders cannot compute higher moments, hence I assume that their

[19] Given the difficulty of the traders' inference problem and the fact that information acquisition is costly in the model, it is plausible to assume that information processing is costly as well.

[20] Formally, condition (ii) means that the uninformed trader correctly computes the average posterior variance conditional on an order being submitted (ignoring, e.g., whether the order submitted is market or limit), and then updates the posterior variance to this same value regardless of the order type. This assumption is necessary because the market and limit orders have different posterior standard deviations (see Footnote 50 below). This difference approaches zero when the informed share ρ is small, and has a maximum possible value of about 13%.

posterior beliefs are always normally distributed.

To avoid different beliefs among uninformed traders, I assume that the initial belief of an uninformed trader is such that after submitting a limit order in the direction of his private valuation, his posterior belief coincides with the posterior belief of the other uninformed traders.[21] Thus, the uninformed traders waiting in the order book have the same normally distributed belief, the "public density." Just before trading at t, I denote the public density by ψ_t, and its mean and standard deviation are, respectively, the "public mean" μ_t and the "public volatility" σ_t.

Private information is much more difficult to process, therefore I assume that an informed trader who chooses to submit a limit order must subsequently use an uninformed trader (who acts as a broker) to update the order.[22]

For tractability, I assume that an informed trader receives a penalty ω if after observing the fundamental value she chooses not to trade.[23] This assumption is equivalent to the informed trader receiving a private benefit ω if she submits an order to the market, which intuitively can arise from "learning by trading." Because ω indicates a commitment to trade by the informed investor, it is called the "commitment parameter." In Section 5.2 in the Internet Appendix, I show that this assumption is necessary only if the number of informed traders is above a threshold.

F. Robustness

The model described thus far can be solved essentially in closed form. It can be used therefore as a benchmark model to study the robustness of the equilibrium results. In Section 5 in the Internet Appendix, I study the effect of relaxing some of the assumptions that are made for tractability. I then verify that the equilibrium is not significantly

[21]This assumption reconciles the divergence in beliefs that private knowledge about own type can create. For instance, an uninformed trader who submits a limit order privately knows that his order is uninformed, but the other uninformed traders do not know and may update their beliefs. See the proof of Lemma A3 in the Appendix for a formal discussion.

[22]This simplifying assumption is justified by two arguments. First, private information processing is indeed difficult: an informed trader must learn not just how the public density evolves, but must also use her signal to form a private belief about the asset value (she only observes the asset value once). Second, even if she could properly update her private belief, while waiting in the book she might not want to deviate from the uninformed strategy, as this would reveal information to the public.

[23]This assumption in needed to avoid no-order regions for the informed trader, which can occur when her perceived mispricing is close to zero.

affected by relaxing these assumptions.

III. Equilibrium

The simplifying assumptions in Section II.E imply that we can consider MPEs in which the only relevant state variables are the public mean and the public volatility, corresponding to the first two moments of the uniformed traders' posterior belief about the asset value.

In this section, I describe a MPE in which the public volatility is constant and equal to the parameter V defined in equation (7) below. Moreover, in Section V below, I examine non-stationary equilibria corresponding to different initial public volatility, and show that all of these equilibria converge to the stationary equilibrium of this section (Result 3). In the rest of the paper, therefore, I refer to the equilibrium in this section as the (unique) stationary equilibrium.

In Section III.A, I describe intuitively the stationary equilibrium, as well as the role played by several key assumptions. In Section III.B, I introduce the notation used throughout the paper. In Section III.C, I describe the optimal strategies of the informed and uninformed traders, their resulting expected utility, and I examine the equilibrium limit order book and its evolution in time.

A. Intuition and Discussion

I first provide a brief description of the stationary equilibrium, and then explain how it relates to the traders' strategies. In equilibrium, the public volatility is a constant V, and the bid-ask spread is also a constant S. (For a definition of these constants, see Section III.B.) Before the arrival of the first trader at time $t = 0$, there is a countably infinite number of traders in the queue at the ask price $\mu_0 + \frac{S}{2}$, and at the bid price $\mu_0 - \frac{S}{2}$.[24] The arrival of a BMO shifts the public mean, the ask price, and the bid price by a constant Δ. The BMO executes the first SLO at the ask, and the ask and bid

[24]As all orders are for one unit, the traders who are not the first in the queue can have their limit orders above the ask or below the bid, as long as the relative positions in the queue do not change. Nevertheless, the equilibrium shape of the limit order book can be fixed if one imposes an infinitesimal cost of modifying limit orders: see the discussion about Figure 2 below.

queues are shifted by Δ. The arrival of a BLO shifts the public mean, the ask price, and the bid price by a constant $\gamma\Delta$, where $\gamma \approx 0.2554$. The BLO is submitted at the new ask price, such that it becomes the first in the bid queue, and the ask and bid queues are shifted by $\gamma\Delta$. The shifts in the ask and bid queues are done such that the traders never switch their relative positions in the queue. The arrival of a SMO and SLO has similar effects, but with a negative sign. Except for these shifts, traders never cancel or modify their limit orders. All limit orders execute with probability one.

At each integer time $t = 0, 1, \ldots$, a new trader arrives, who is either informed with probability ρ, or uninformed with probability $1 - \rho$. All informed traders are patient, while the uninformed traders are, with equal probability, either buyers or sellers, patient or impatient. By assumption, the impatient traders always submit market orders, and thus provide a source of profit for the patient traders (informed or uninformed) who submit limit orders.[25]

The strategy of the patient uninformed traders is simple. Suppose a patient natural buyer (with positive private valuation and zero waiting costs) arrives to the market. He then submits a buy limit order (BLO), after which he waits for his order to be executed, and in the meantime he modifies his bid to account for the information contained in the order flow. This modification is done such that the traders do not switch their relative positions in the queue.[26] The bid and ask queues consist each of a countably infinite number of buyers and sellers, respectively, as all these traders have zero waiting costs.[27] The patient natural buyer chooses a BLO for two reasons: (i) his private valuation is positive, hence he prefers a buy order to sell order, (ii) his waiting costs are zero, hence

[25] In Section 2 in the Internet Appendix, I show that the impatient traders always submit market orders for sufficiently large values of $\bar{u}$ (private valuation) and $\bar{r}$ (waiting cost). Also, I show that impatient informed traders optimally do not participate in the market.

[26] Switching positions in the queue does not matter for uninformed traders, as they have zero waiting costs. The same goes for the informed traders, because by assumption they must hire an uninformed trader (broker) to handle their orders. In principle, however, an informed trader could realize that her average information advantage decreases over time because of the future arrival of competing informed traders. Thus, she could instruct her broker to jump ahead in the queue in order to ensure a faster order execution. To prevent this behavior, I impose the out-of-equilibrium belief that jumping ahead in the queue can come only from an informed trader.

[27] I conjecture that the main results in this paper remain robust to having small positive waiting costs, but the solution of such a model would be much more complicated. Indeed, as seen in models of the limit order book with symmetric information but positive waiting costs, such as Roşu (2009), the numbers of limit orders on each side of the book become additional state variables. In that case, it is plausible that the patient traders start submitting market orders in states when their queue size exceed a particular value, as their expected waiting cost becomes too high.

he prefers a limit order to a market order.[28]

To understand the strategy of informed traders, it is enough to describe their initial order choice, as subsequently their orders are handled by an uninformed trader.[29] This optimal choice problem is difficult to solve. To understand why, consider an informed trader who arrives and observes the asset value, or equivalently the "mispricing," which is the difference between the asset value and the public mean. Then, in order to decide what order to submit, she must be able to estimate for instance the payoff of a BLO. This is a complex problem, because she must take an average over all future order flow sequences that lead to the execution of her BLO. It turns out, however, that this payoff can be described easily if one can compute a certain function of two variables called the "information function" (see Definition 1 below). This function can be estimated only numerically, but otherwise the main formulas in the paper are given in closed form.

I then show that the informed trader's choice is based on a threshold strategy. For instance, if she observes a mispricing above a threshold, she optimally submits a BMO; if the mispricing is below the threshold (but positive), she optimally submits a BLO. With a BMO she loses half the bid-ask spread, but trades immediately. With a BLO she gains half the bid-ask spread, but she expects to lose because her information advantage (the mispricing) decays over time. Indeed, the informed traders who arrive later observe more recent instances of the asset value, and hence reduce the mispricing by their trading.[30] Not surprisingly, a limit order's decay cost is larger when the mispricing is larger. The benefit of a limit order relative to a market order, however, does not depend on the observed mispricing, and is equal to the bid-ask spread. Thus, at the threshold mispricing, the informed trader is indifferent between BMO and BLO. Moreover, the equilibrium bid-ask spread is equal to the expected decay cost incurred by the informed

[28]One also needs to show that, relative to a market order, the benefits of a limit order (the bid-ask spread) are larger than the costs (adverse selection when the BLO is executed by a sell market order, and the initial price impact of the BLO). See the proof of Theorem 1 and equation (A-18).

[29]See Footnote 22 for a justification of this assumption.

[30]This information advantage decay arises from the assumption that informed traders observe the asset value only once, when they arrive. I have not been able to solve a model in which the informed traders continuously learn about the asset value. But even in such a model, it is plausible that traders who wait in the book would imitate the behavior of the uninformed traders (see the second part of Footnote 22), and would thus be adversely selected later by informed market orders. This would not a problem, though, if the limit order traders received enough compensation from uninformed market orders and from a sufficiently large bid-ask spread.

trader at the threshold mispricing. The exact value for the threshold is determined by a "dynamic market clearing" condition: all types of orders must be equally likely ex ante. This condition is true exogenously for the uninformed traders, and hence in a stationary equilibrium the informed traders must follow it as well.[31]

The threshold argument above has two important consequences: first, all orders have information content. In particular, a buy limit order has a positive impact on the public mean: if the market sees a BLO, with positive probability it comes from an informed trader who observed an asset value above the public mean (but below the threshold). A buy market order has an even stronger effect, as the asset value must have been above the threshold.

An important variable in the model is the public volatility, which measures how uncertain the uninformed traders are about the current asset value. There are two opposite forces operating on the public volatility: First, the order flow carries information, and this reduces the public volatility. Second, the asset value diffuses over time, and this increases the public volatility. In a stationary equilibrium, the two effects exactly offset each other, and the public volatility remains constant.[32] A key fact behind many results is the relation between the informed share and the public volatility. A larger informed share implies that the order flow is more informative about the asset value, and thus generates a smaller public volatility (better public knowledge of the asset value).

Note that in the model, the bid-ask spread is determined by the informed traders. Indeed, these traders have zero private valuation, and choose the spread that compensates them for their information decay cost. The patient uninformed traders are not marginal: their private valuation is sufficiently high, and thus they are always willing to trade with limit orders. Note also that the informed traders face adverse selection from future informed traders, hence the bid-ask spread has an adverse selection component.

Finally, I briefly discuss the case in which the informed share ρ is close to one.

[31] In a previous version of this paper (available upon request), I show that the results in the paper are robust if the uninformed limit-to-market order ratio is exogenously chosen different from one.

[32] A key simplifying assumption is that information processing is costly, and as a result uninformed traders always perceive the public density as normal. With perfect Bayesian updating, however, the informed trader's threshold strategy leads to non-normal distributions. In Section 5.1 in the Internet Appendix, I examine perfect Bayesian updating, and find that the departure from normality is small, especially for the average public density. I thus conjecture that the main results remain true on average under perfect Bayesian updating.

Then, there are relatively few uninformed traders (their share is $1-\rho$), and even fewer impatient uninformed traders (their share is $\frac{1-\rho}{2}$). But these traders always submit market orders, and are therefore the source of profit for limit order traders. As a result, the expected profit of an informed trader who submits a limit order is small, although it is still positive, and thus an equilibrium still exists.[33] Note that when ρ is close to one, the adverse selection is not infinite: the order flow coming from informed traders generates public information, in such a way that the public volatility and the bid-ask spread are bounded. In fact, both these variables are decreasing in the informed share, as more informed traders generate better public information.

B. Notation and Parameters

The exogenous parameters in the model are: the fundamental volatility σ_v, the informed trading activity λ^i, the uninformed trading activity λ^u, the private valuation parameter $\bar{u}$, the impatience parameter $\bar{r}$, and the commitment parameter ω. I define other parameters: the trading activity $\lambda = \lambda^u + \lambda^i$, the informed share $\rho = \lambda^i/(\lambda^i + \lambda^u)$, and the inter-arrival volatility $\sigma_I = \sigma_v/\sqrt{\lambda}$.

Let $\phi(\,\cdot\,; m, s)$ be the normal density with mean m and standard deviation s, $\phi(\,\cdot\,) = \phi(\,\cdot\,; 0, 1)$ the standard normal density, and $\Phi(\,\cdot\,)$ its cumulative density. Let $\mathbf{1}_X$ be the indicator function which equals one if X is true and zero if X is false.

I next define four numeric parameters that are used extensively throughout the paper. The first three are:

(4)

$$\alpha = \Phi^{-1}\left(\tfrac{3}{4}\right) \approx 0.6745, \quad \beta = \frac{1}{4\phi(\alpha)} \approx 0.7867, \quad \gamma = \frac{\phi(0)-\phi(\alpha)}{\phi(\alpha)} \approx 0.2554,$$

where $\phi(\,\cdot\,)$ is the standard normal density, and $\Phi(\,\cdot\,)$ is its cumulative density.

In Definition 1, I introduce the fourth numeric parameter, the "information function." Formally, this is a function $I = I(\rho, w, j)$ defined on $(0,1) \times \mathbb{R} \times \mathbb{N}_+$, but I show

[33] This is true unconditionally, before the asset value is observed. Conditionally, however, it is possible that an informed trader who observes an asset value only slightly above the public mean, might prefer not to trade with a BLO. This is where the assumption of a commitment parameter ω comes in: to avoid a penalty for not trading, she now prefers to submit a BLO. In Section 5.2 in the Internet Appendix, I show that the commitment parameter is necessary only if the informed share is above a threshold, approximately equal to 0.156.

below that it has an interpretation in the model. To that end, I refer to the elements of the set $\{\text{BMO}, \text{BLO}, \text{SLO}, \text{SMO}\}$ as "orders," even though this is just an abstract set with four elements.

Definition 1. *Let $\rho \in (0,1)$, $w \in \mathbb{R}$, $j \in \mathbb{N}_+$. For each order $\mathcal{O} \in \{\text{BMO}, \text{BLO}, \text{SLO}, \text{SMO}\}$, define, respectively, $\delta_{\mathcal{O}} \in \left\{\frac{\rho}{\beta}, \gamma\frac{\rho}{\beta}, -\gamma\frac{\rho}{\beta}, -\frac{\rho}{\beta}\right\}$, $i_{\mathcal{O}} \in \{(\alpha,\infty), (0,\alpha), (-\alpha,0), (-\infty,-\alpha)\}$, and $j_{\mathcal{O}} \in \{0, +1, 0, -1\}$. If g is a density over $\mathbb{R}$, and $G_{\mathcal{O}} = \int_{z \in i_{\mathcal{O}}} g(z)\mathrm{d}z$, define the scalar $\pi_{g,\mathcal{O}}$ and the density $f_{g,\mathcal{O}}$ by:*

$$(5)\quad \pi_{g,\mathcal{O}} = \tfrac{1-\rho}{4} + \rho G_{\mathcal{O}}, \quad f_{g,\mathcal{O}}(x) = \frac{\int \left(\frac{1-\rho}{4} + \rho \mathbf{1}_{z \in i_{\mathcal{O}}}\right) g(z) \phi\left(x; z - \delta_{\mathcal{O}}, \rho\sqrt{\frac{1+\gamma^2}{2\beta^2}}\right) \mathrm{d}z}{\pi_{g,\mathcal{O}}}.$$

If $T \in \mathbb{N}_+$, a sequence of orders $Q = (\mathcal{O}_0, \mathcal{O}_1, \dots, \mathcal{O}_T)$ is called a "j-execution sequence" if $j + \sum_{t=1}^{T} j_{\mathcal{O}_t} = 0$, but for any $T' = 0, 1, \dots, T-1$, $j + \sum_{t=1}^{T'} j_{\mathcal{O}_t} \neq 0$. For any j-execution sequence $Q = (\mathcal{O}_0, \mathcal{O}_1, \dots, \mathcal{O}_T)$, and any density g_1 over $\mathbb{R}$, define $P(Q) = \prod_{t=1}^{T} P_t$ and $\nu(Q) = \nu_{T+1} - \frac{\rho}{\beta}$, where one recursively defines $P_t = \pi_{g_t,\mathcal{O}_t}$, $g_{t+1} = f_{g_t,\mathcal{O}_t}$, and $\nu_{t+1} = \mathsf{E}(g_{t+1})$ $(t = 1, \dots, T)$. Let $\mathcal{Q}_j$ be the set of all j-execution sequences of the form $Q = (\text{BLO}, \mathcal{O}_1, \dots, \mathcal{O}_T)$ for some $T \in \mathbb{N}_+$. Then, the information function is:

$$(6)\qquad I(\rho, w, j) = \sum_{Q \in \mathcal{Q}_j} P(Q)\nu(Q), \quad \text{with} \quad g_1 = \mathcal{N}\left(w - \gamma\frac{\rho}{\beta}, \rho^2\frac{1+\gamma^2}{2\beta^2}\right).$$

Moreover, define $J(\rho, w, j)$ as in (6), but with $\nu(Q) = 1$. If $j = 1$, omit the dependence on j, and write $I(\rho, w) = I(\rho, w, 1)$ and $J(\rho, w) = J(\rho, w, 1)$.

Before discussing the information function, I introduce several more parameters:

$$(7)\qquad \begin{aligned} \Delta &= \sqrt{\tfrac{2}{1+\gamma^2}}\frac{\sigma_v}{\sqrt{\lambda}} && = \text{impact parameter}, \\ V &= \beta\rho^{-1}\Delta, && = \text{volatility parameter}, \\ S &= \big(\alpha - I(\rho,\alpha)\big)\, V && = \text{spread parameter}. \end{aligned}$$

I briefly explain how the information function I is interpreted in the model. Consider an informed trader who arrives at $t = 0$ and observes an asset value v_0. Suppose that at $t = 0$, the informed trader submits a BLO such that this order has initial rank $j_0 \in \mathbb{N}_+$ in the bid queue, after which she follows the strategy of an uninformed trader (described in

Corollary 4 below). This implies that she patiently waits in the queue until a sequence of orders $\mathcal{O}_1, \mathcal{O}_2, \ldots, \mathcal{O}_T$ finally executes her BLO. Note that execution occurs at T only if $\mathcal{O}_T = \text{SMO}$. Just before trading at t, the uninformed traders regard the asset value v_t as distributed by the normal density ψ_t (the public density), with mean μ_t (the public mean) and volatility σ_t (the public volatility). In the stationary equilibrium of this section, the public volatility is constant and equal to the parameter V from (7). It is therefore convenient to normalize variables by V. I define the "signal" at t as the normalized mispricing just before trading at t:

$$w_t = \frac{v_t - \mu_t}{V}. \tag{8}$$

Thus, for the uninformed traders, the distribution of w_t in the stationary equilibrium remains the same at all times, namely the standard normal distribution.

The arguments of $I(\rho, w, j)$ are interpreted as follows: First, ρ represents the informed share. Second, w represents the initial signal $w_0 = \frac{v_0 - \mu_0}{V}$, before the informed trader submits the BLO at $t = 0$. Third, j represents the rank $j_0 \in \mathbb{N}_+$ in the bid queue of her BLO. The symbols used in Definition 1 are interpreted as follows: First, g_t is the posterior density of the signal w_t before trading at t, conditional on observing the sequence of orders $\mathcal{O}_0 = \text{BLO}, \mathcal{O}_1, \ldots, \mathcal{O}_{t-1}$. The mean of g_t is $\nu_t = \mathsf{E}(g_t)$, and P_t is the probability of an order $\mathcal{O}_t$ being submitted at t. The rank of the informed trader's BLO in the bid queue after trading at t is j_t. An execution sequence is a sequence of orders $Q = (\mathcal{O}_0 = \text{BLO}, \mathcal{O}_1, \ldots, \mathcal{O}_T = \text{SMO})$ such that the last order (SMO) executes the initial BLO, which translates into the final rank j_T being zero. Next, $P(Q)$ is the ex-ante probability of a particular execution sequence $Q = (\mathcal{O}_0, \mathcal{O}_1, \ldots, \mathcal{O}_T)$, and $\nu(Q) = \nu_{T+1} - \frac{\rho}{\beta}$ is the expected signal w_T after the execution at T.[34]

Thus, the information function $I(\rho, w, j)$ is interpreted as the expected signal w_T immediately after the BLO is executed at T by an SMO, where the expectation is taken over all possible order sequences that eventually execute the BLO. Proposition 2 below shows that this interpretation of I is indeed correct.[35]

[34] The expected signal w_T after T is similar to the expected signal w_{T+1} before $T+1$ (which is ν_{T+1}), except that the public means at T and $T+1$ differ by the price impact of a SMO, which is $-\Delta$. Thus, $\nu(Q) = \nu_{T+1} - \frac{\Delta}{V} = \nu_{T+1} - \frac{\rho}{\beta}$. For more details, see the proof of Lemma A2 in the Appendix.

[35] The function $J(\rho, w, j)$ is interpreted as the probability that the initial BLO (which has initial

Despite the interpretation of the information function within the model, the definition itself is completely independent of this interpretation, and therefore I can be thought as a parameter. Even though it does not have a closed form expression, it can be estimated with good precision by using a numerical Monte Carlo procedure described in detail in Section 4 in the Internet Appendix. The next numerical result describes several properties of the information function which are used in Theorem 1.

Result 1. For all $\rho \in (0, 1)$, the functions $I(\rho, w)$, $w - I(\rho, w)$ and $I(\rho, w) - I(\rho, -w)$ are strictly increasing in w, and satisfy the inequality:

$$\max\left(\frac{\rho(1+\gamma)}{\beta}\,,\ -2I(\rho, 0) - 2\frac{\rho\gamma}{\beta} \right) \ < \ \alpha - I(\rho, \alpha). \tag{9}$$

Moreover, (i) $I(\rho, w, j)$ decreases in j if $w > 0$; and (ii) J is constant and equal to one.

C. Results

To describe a Markov Perfect Equilibrium (MPE), I need to specify the state variables on which the traders' strategies depend. The public state variables are: the public density, determined by its first two moments (the public mean and the public volatility), and the limit order book, determined by the ask and bid queues. The private state variable is the asset value, observed by each informed trader when arriving to the market.

I define the initial state of the system, an instant before $t = 0$. If V is the volatility parameter from (7), the initial public density is $\mathcal{N}(0, V^2)$, with public mean equal to zero, and public volatility equal to V. If S is the spread parameter from (7), the ask price is $S/2$, the bid price is $-S/2$, while the initial order book has countably infinitely many limit orders on each side (see the middle plot in Figure 2).

Theorem 1 shows that there exists a MPE of the model if the conditions stated in Result 1 above are satisfied. These conditions are verified numerically in Section 4 in the Internet Appendix.

Theorem 1. *Suppose the information function I satisfies analytically the conditions from Result 1, and the investor preference parameters satisfy $\bar{u} \geq \frac{S}{2}$ and $\omega \geq \gamma\Delta$. Then,*

rank j in the bid queue) is eventually executed. Numerically, J is identically equal to one, indicating that the BLO is executed almost surely. Thus, there is no need to consider J as another parameter.

there exists a stationary Markov Perfect Equilibrium of the game.

I describe the main properties of the equilibrium in the Corollaries 1–4 below. Corollary 1 describes the evolution of the public mean, bid price, and ask price. Corollary 2 describes the initial order submission strategy of the informed trader. Corollary 3 shows that all types of orders are equally likely. Corollary 4 describes the initial strategy of the uninformed traders, and the subsequent equilibrium behavior of all types of traders in the limit order book.

Corollary 1. *In equilibrium, the public volatility and the bid-ask spread are constant and equal, respectively, to the parameters V and S from* (7)*. If the public mean is μ_t, the ask price is $\mu_t + S/2$, while the bid price is $\mu_t - S/2$.*

The public mean changes only when a new order arrives. Let $\gamma \approx 0.2554$ be as in equation (4)*. If an order arrives at t, the public mean changes from μ_t to* (*i*) *$\mu_t + \Delta$ if the order is* BMO*,* (*ii*) *$\mu_t + \gamma\Delta$ if the order is* BLO*,* (*iii*) *$\mu_t - \gamma\Delta$ if the order is* SLO*, and* (*iv*) *$\mu_t - \Delta$ if the order is* SMO*.*

The first part of Corollary 1, that the public volatility and the bid-ask spread are constant over time, follows from the stationarity of the equilibrium. I postpone this discussion until after Corollary 3.

To get intuition for the second part of Corollary 1, recall that the public mean is the expected asset value given the public information (the information of the uninformed traders). A new order affects the public mean because each type of order contains private information. For instance, according to Corollary 2 below, an informed trader submits a BMO for extreme signals, i.e., w_t larger than $\alpha \approx 0.6745$; and a BLO for positive moderate signals, i.e., w_t lies in $(0, \alpha)$. This implies that BMO increases the public mean by some amount (Δ), while a BLO increases the public mean by a smaller amount ($\gamma\Delta \approx 0.2554\,\Delta$).

Thus, the key to understanding the equilibrium is the strategy of the informed trader, which is described in the next result.

Corollary 2. *Suppose an informed trader arrives at $t \geq 0$, and observes a signal $w_t = \frac{v_t - \mu_t}{V}$. Then, she submits a* (*i*) BMO *if $w_t \in (\alpha, \infty)$,* (*ii*) BLO *if $w_t \in (0, \alpha)$,* (*iii*) SLO *if* $w_t \in (-\alpha, 0)$, *or (iv)* SMO *if* $w_t \in (-\infty, -\alpha)$. *Depending on the order submitted, her expected utility is:*

$$
\begin{aligned}
(10) \qquad \mathcal{U}^I_{\text{BLO}} &= \frac{S}{2} + V\, I(\rho, w_t), & \mathcal{U}^I_{\text{SLO}} &= \frac{S}{2} + V\, I(\rho, -w_t),\\
\mathcal{U}^I_{\text{BMO}} &= -\frac{S}{2} + V\, w_t, & \mathcal{U}^I_{\text{SMO}} &= -\frac{S}{2} - V\, w_t.
\end{aligned}
$$

To understand this result, suppose the informed trader gets a positive signal w_t. Then, her main choice is between submitting a BMO and a BLO. By submitting a BMO, she gains from her signal (w_t), but loses half of the bid-ask spread ($S/2$) because she has to pay the ask price, which is higher than the public mean by $S/2$ (see Corollary 1). By submitting a BLO instead, equation (10) implies that the informed trader gains half of the bid-ask spread, and also benefits from her signal via the information function $I(\rho, w)$.

The information function increases in w at a lower rate than w itself. Formally, this follows from Result 1, according to which $w - I(\rho, w)$ is increasing in w. Intuitively, this is because an informed trader who observes a large signal w_t knows that other informed traders are also likely to receive positive signals in the future, and therefore are more likely to submit buy orders. This bias towards buy orders therefore pushes up the public mean in the future. In other words, the informed trader with a BLO expects to buy at a higher price in the future while she waits in the book. The stronger her signal, the stronger the bias, and therefore the stronger the relative penalty from submitting a BLO compared to a BMO. A more detailed discussion of this phenomenon, which is called "slippage," is left for Section IV.

Because the function $w - I(\rho, w)$ is increasing in w, the payoff difference between BMO and BLO is increasing in w. Therefore, for some threshold α, the informed trader prefers BMO for $w_t > \alpha$, and BLO for $w_t \in (0, \alpha)$. Intuitively, with an extreme signal the informed trader should use a market order, while with a moderate signal the informed trader should use a limit order. At the threshold $w = \alpha$ (which occurs with zero probability), the informed trader is indifferent between BMO and BLO. The threshold $\alpha = \Phi^{-1}(3/4)$ is given by equation (4), and satisfies the property that for a variable w with the standard normal distribution, the probability that $w \in (\alpha, \infty)$ is equal to the probability that $w \in (0, \alpha)$ and is equal to $1/4$. This corresponds to the fact that all order types (BMO, BLO, SLO, SMO) are equally likely, with probability $1/4$ (see Corollary 3 below).

Figure 1: The Order Choice of the Informed Trader
This figure displays (i) the public density, $\psi_t \sim \mathcal{N}(\mu_t, V^2)$, which is the density of the asset value v_t conditional on all public information until t, (ii) the four intervals on the horizontal axis that define the equilibrium order choice of an informed trader after observing v_t. The parameter $\alpha \approx 0.6745$ is as in equation (4).

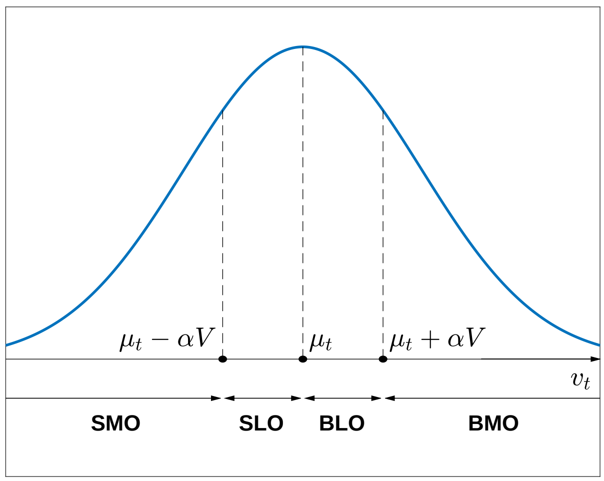

Figure 1 illustrates the equilibrium order choice of the informed trader. The threshold between BMO and BLO is given by $w_t = \alpha$, or equivalently by $v_t = \mu_t + \alpha V$. The normal curve in the figure represents the public density, which is the public belief about the asset value. The four regions under the curve and above the horizontal axis have area equal to $1/4$ each, which reflects the fact that the informed trader submits each of the four order types with the same ex-ante probability. Because the four types of orders are also equally likely for an uninformed trader,[36], and because there are no cancellations in equilibrium, it follows that the four types of orders are equally likely in equilibrium given public information. I state this result in the next corollary.

Corollary 3. *Conditional on public information, all order types* (BMO, BLO, SLO, *and* SMO) *are equally likely in equilibrium, with probability* $1/4$.

[36]Indeed, the four types of uninformed traders arrive with equal probability, and the patient natural buyers submit BLO and the patient natural sellers submit SLO (see Corollary 4), while the impatient natural buyers submit BMO and the impatient natural sellers submit SMO.

I call this equilibrium property "dynamic market clearing." It is equivalent to the following two properties: (i) buy and sell orders are equally likely, and (ii) market and limit orders are equally likely. It is the second property that is key for the intuition regarding dynamic market clearing. Suppose for instance that market orders were more likely than limit orders. Since every market order is executed against a limit order, the limit order book would become thinner over time, and therefore the equilibrium would not be stationary. Thus, dynamic market clearing occurs because the equilibrium in Theorem 1 is stationary. In Section V, I analyze non-stationary equilibria of the model, and find that the dynamic market clearing condition no longer holds.

The next corollary describes the initial order submission decision of the uninformed traders, as well as their subsequent strategy once they submit a limit order. One only needs to understand the patient uninformed traders, since the impatient ones always submit market orders. Also, because the informed traders are essentially uninformed after the initial order choice, the subsequent equilibrium behavior of the informed and uninformed traders coincides.

Corollary 4. *Consider a patient uninformed trader with private valuation u larger in absolute value than $\Delta - S/2$. Then, he submits a* BLO *if he is a natural buyer, and a* SLO *if he is a natural seller. In both cases, his expected utility is:*

$$\mathcal{U}^U = \frac{S}{2} - \Delta + \bar{u}. \tag{11}$$

After the initial limit order is submitted, the uninformed trader modifies his order along with the public mean, as in Corollary 1. If an informed trader chooses to submit a limit order, her subsequent behavior mimicks the behavior of an uninformed trader. Traders in the limit order book modify their orders such that their relative rank in the ask or bid queue is preserved.

The intuition behind Corollary 4 is straightforward. A patient natural buyer who submits a BLO gains half of the bid-ask spread ($S/2$), as well as his private valuation ($\bar{u}$), but loses from the adverse selection of the SMO that eventually executes his order (according to Corollary 1, the price impact of a SMO is $-\Delta$). Hence, as long as his private valuation is large enough to make his expected utility in (11) positive, he optimally submits a BLO. After submitting the initial order, the uninformed trader simply modifies his order according to the evolution of the public mean, because he is risk-neutral and updates his estimate of the asset value according to the public mean.

Figure 2: The Effect of Order Flow on the Limit Order Book
This figure displays the equilibrium shape of the limit order book just before trading at t (middle plot), as well as the shape of the book at $t+1$ after a buy market order BMO (right plot) or a buy limit order BLO (left plot). For simplicity, the public mean is set to $\mu_t = 0$, so that before trading at $t+1$, the public mean becomes $\mu_{t+1} = \Delta$ after BMO, or $\mu_{t+1} = \gamma\Delta$ after BLO. The parameter $\gamma \approx 0.2554$ is as in equation (4).

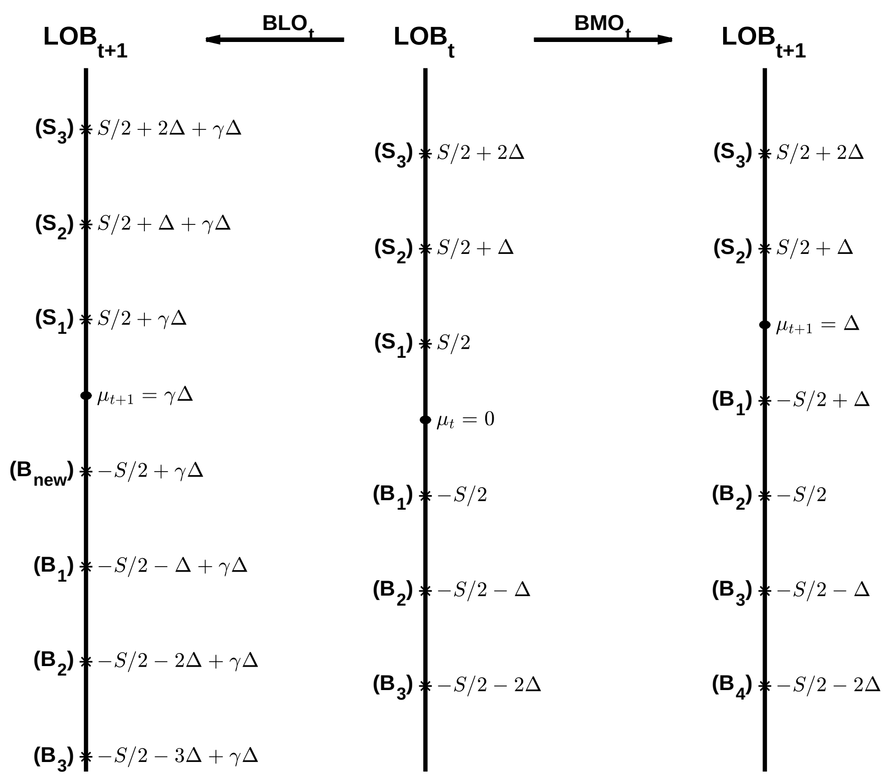

Moreover, limit order traders preserve their relative position in the ask or bid queues. Indeed, uninformed traders have zero waiting costs, and therefore have no incentive to change their position in the queue. By contrast, if they were allowed, informed traders would prefer to jump ahead in the queue, because (if nothing else changed) this would reduce the expected decay in their information advantage (see Section IV.C). Nevertheless, this behavior cannot occur in equilibrium. To see this, suppose a trader were to jump ahead in the bid queue. This out-of-equilibrium behavior would be interpreted immediately as coming from an informed trader with positive information. This new information would then increase the public mean, and reduce the informed trader's information advantage. The reduction in expected payoff would then prevent the trader from deviating in the first place.

Normally, without additional assumptions one should not expect the equilibrium limit order book in the model to have a well defined shape. Indeed, trading is for only one unit, and without any modification cost the exact position of limit orders away from the bid and ask does not matter. However, the equilibrium shape of the limit order book can be fixed if I impose an infinitesimal cost of modifying limit orders. Suppose that when a limit order is executed at the ask, there is an infinitesimal modification cost for all the remaining limit orders on the ask side (and similarly for the bid side).

The resulting equilibrium limit order book is described in Figure 2. The middle plot in Figure 2 describes the typical shape of the limit order book just before trading at t. For simplicity, the public mean is set at $\mu_t = 0$. The right and the left plots, respectively, describe the effect of a BMO or a BLO on the limit order book. To understand the assumption about the infinitesimal modification cost, suppose a BMO arrives at date t, when the limit order book is in the middle plot of Figure 2. Then, the SLO of trader S_1 is executed, and trader S_2 becomes the first in the ask queue. An instant later, S_2 should immediately modify his SLO at $\mu_t + S/2 + \Delta$, and therefore, with an infinitesimal modification cost, S_2 would prefer to have his order at that price already.

IV. Market Quality and Informed Trading

In this section, I consider several measures of market quality and analyze how they are affected by the informed share, which is the fraction of order flow generated by the informed traders. As measures of market quality, I consider the information efficiency, as well as three measures of liquidity: the price impact, the bid-ask spread, and the market resiliency. In the process, I also study the information content of the different types of orders.

A. Information Efficiency

In general, a market is efficient at processing information if pricing errors are small. In the model, the pricing error is the difference between the fundamental value v and the public mean μ, and the standard deviation of the pricing error is the public volatility. According to Corollary 1, in equilibrium the public volatility is constant and equal to the parameter $V = \beta\rho^{-1}\Delta$. I thus propose the following measure of information efficiency:

$$\frac{1}{V^2} \;=\; \frac{\rho^2}{\beta^2\Delta^2}, \quad \text{with} \quad \Delta \;=\; \sqrt{\tfrac{2}{1+\gamma^2}}\,\frac{\sigma_v}{\sqrt{\lambda}}, \tag{12}$$

and β, γ are defined in (4). Note that when the market is informationally efficient, the public volatility is small, and therefore the proposed measure is large.

Because β and Δ are independent of ρ, the information efficiency is increasing in the informed share ρ.[37] It follows that information efficiency is increasing with the informed share. This shows that when there are more informed traders (ρ is large), the order flow is more informative, hence the market is more efficient at processing information. An interesting aspect of the increase in information efficiency is that it arises from the dynamic nature of the equilibrium. In a static equilibrium (see Glosten and Milgrom 1985), the opposite happens: when there are more informed traders the adverse selection is larger, and therefore the market is less informationally efficient. This intuition is discussed in more detail below, after Proposition 1.

The public volatility V can be used to estimate in practice the informed share. The problem is that it depends on other parameters of the model, such as the fundamental volatility σ_v and the total trading activity λ. To remove this dependence, I consider the ratio of the inter-arrival volatility ($\sigma_I = \sigma_v/\sqrt{\lambda}$) to the public volatility (V), which is:

$$\frac{\sigma_I}{V} \;=\; \rho\sqrt{\frac{1+\gamma^2}{2\beta^2}} \;\approx\; 0.9277\,\rho \;<\; 1. \tag{13}$$

The ratio σ_I/V provides a clean estimate of the informed share ρ, in the sense that the ratio does not depend on additional parameters. The inter-arrival volatility σ_I is in principle observable, as the price variance between order arrivals. The public volatility is not observable directly, but it can be proxied by the dispersion of financial analysts' estimates. Since (as I show in Section IV.C), the bid-ask spread S is decreasing in the informed share ρ, a testable implication of equation (13) is that stocks with a lower ratio

[37] The fact that Δ is independent of ρ is obvious from its formula. The economic interpretation of this fact, however, is not obvious, and I discuss it in Proposition 1 and the paragraphs that follow it.

of inter-arrival volatility to public volatility have larger bid-ask spreads.

B. Price Impact

I define the price impact of an order as the effect of one additional unit of trading on the transaction price. Since all trades in the model are for one unit, the marginal price impact measure is the same as the effect of one unit on the public mean.[38] Because there are four types of orders, each order type $\mathcal{O} \in \{\mathrm{BMO}, \mathrm{BLO}, \mathrm{SLO}, \mathrm{SMO}\}$ has a different price impact, which I denote by $\Delta_{\mathcal{O}}$. Corollary 1 implies the following result.

Proposition 1. *The price impact $\Delta_{\mathcal{O}}$ of any order $\mathcal{O} \in \{\mathrm{BMO}, \mathrm{BLO}, \mathrm{SLO}, \mathrm{SMO}\}$ is:*

$$\Delta_{\mathrm{BMO}} = \Delta, \quad \Delta_{\mathrm{SMO}} = -\Delta, \quad \Delta_{\mathrm{BLO}} = \gamma\Delta, \quad \Delta_{\mathrm{SLO}} = -\gamma\Delta, \tag{14}$$

where $\gamma \approx 0.2554$ is as in equation (4), and $\Delta = \sqrt{\frac{2}{1+\gamma^2}}\,\frac{\sigma_v}{\sqrt{\lambda}} \approx 1.3702\,\frac{\sigma_v}{\sqrt{\lambda}}$ is as in equation (7). In particular, $\Delta_{\mathcal{O}}$ does not depend on the informed share ρ. Moreover, the variance of the price impact is equal to the inter-arrival variance $\sigma_I^2 = \frac{\sigma_v^2}{\lambda}$, that is,

$$\mathsf{Var}\big(\Delta_{\mathcal{O}}\big) = \frac{1+\gamma^2}{2}\,\Delta^2 = \frac{\sigma_v^2}{\lambda}. \tag{15}$$

The reason why all order types have price impact is given by the usual adverse selection argument. Indeed, when setting the public mean, the uninformed traders take into account the information contained in the order flow. For instance, if a BMO is submitted at t, then with positive probability it comes from an informed trader with a large signal $w_t = \frac{v_t - \mu_t}{V} \in (\alpha, \infty)$. Then, the public mean should increase (by Δ). Similarly, if a BLO is submitted at t, then with positive probability it comes from an informed trader with a moderate signal, $w_t \in (0, \alpha)$. Then, the public mean should increase as well, although by a smaller amount (by $\gamma\Delta$).

A surprising implication of Proposition 1 is that the informed share ρ has no effect on Δ. To give intuition for this result, note that there are two opposite effects of the

[38] Alternatively, given the equilibrium shape of the limit order book (see Figure 2), one can also define the "instantaneous" price impact of a multi-unit market order, even though such orders are not part of the model. Then, as the size of the market order increases, each additional unit trades at a price changed by Δ. This shows that the two definitions are consistent.

informed share on Δ. Suppose the informed share is small, and a buy market order arrives. The first effect is the usual "adverse selection effect" (see for instance Glosten and Milgrom (1985)): because ρ is small, it is unlikely that the market order comes from an informed trader. This reduces the adverse selection coming from informed traders, and therefore decreases the price impact. But there is a second effect, the "dynamic efficiency effect:" if the buy market order does come from an informed trader, she must have observed an asset value far above the public mean; otherwise, knowing there is little competition from the other informed traders, she would have submitted a buy limit order.[39] This effect increases the price impact.

Intuitively, the fact that the two effects exactly cancel each other follows from the equilibrium being stationary. Indeed, in Section 7 in the Internet Appendix, I show more generally that in a stationary equilibrium the change in asset value and the change in public mean must have the same variance. In the present context, this translates to $\mathsf{Var}(v_{t+1} - v_t) = \mathsf{Var}(\mu_{t+1} - \mu_t)$. But the variance of the asset value change is the inter-arrival variance σ_I^2, which does not depend on the informed share, while the variance of the public mean change is $\mathsf{Var}(\Delta_{\mathcal{O}})$, which according to Proposition 1 is a constant multiple of Δ^2. Therefore, the price impact Δ is independent of the informed share ρ.

Proposition 1 yields a testable implication of the model, namely that the ratio of the price impact of a buy market order to the price impact of a buy limit order is:

$$\frac{\Delta_{\text{BMO}}}{\Delta_{\text{BLO}}} = \frac{1}{\gamma} \approx 3.9152, \tag{16}$$

which is close to 4. Interestingly, Hautsch and Huang (2012, p.515) find empirically that market orders have a permanent price impact of about four times larger than limit orders of comparable size.

C. Bid-Ask Spread

Another measure of liquidity is the bid-ask spread, which is by definition the difference between the ask price and the bid price. Corollary 1 implies the equilibrium bid-ask

[39] Formally, when ρ is small, the informed trader's threshold for the choice between BMO and BLO is large. Indeed, Corollary 2 implies that the threshold signal is $w_t = \alpha$, or equivalently $v_t = \mu_t + \alpha V$. But, as discussed in Section IV.A, the public volatility V is decreasing in ρ.

spread is constant and is equal to the parameter S from equation (7).

Corollary 5. *The equilibrium bid-ask spread is constant over time, and is equal to:*

$$S \;=\; \big(\alpha - I(\rho,\alpha)\big)\, V. \tag{17}$$

To get more intuition about the equilibrium bid-ask spread, I explain how the information function I is interpreted in the model. Consider an informed trader who arrives at $t = 0$, observes a signal $w = \frac{v_0-\mu_0}{V}$ and submits a BLO (which is not necessarily optimal). Assuming that subsequently all investors follow their equilibrium strategies, the informed trader then forms an expectation about the average asset value, based on all possible future order flow that executes her BLO at a later random time $T > 0$. The fact that the BLO is executed at T means that (i) the BLO is the first order in the bid queue before trading at T, and (ii) a SMO is submitted at T.

To state the next result, I introduce some notation. Let E_t be the informed trader's expectation conditional on her information set before trading at t, $\mathcal{J}_t = \{w, \mathcal{O}_1, \ldots, \mathcal{O}_{t-1}\}$, and let E^{e} be the informed trader's expectation at $t = 0$ over all future "execution sequences," i.e., over order sequences $\mathcal{O}_1, \ldots, \mathcal{O}_T$ that execute the BLO at some $T > 0$.[40]

Proposition 2. *Consider an informed trader who observes at $t = 0$ a signal w, and submits a* BLO*, which is executed at a random time $T > 0$. Let $\rho \in (0,1)$ be the informed share. Then, the information function I satisfies:*

$$I(\rho, w) \;=\; \mathsf{E}^{\mathrm{e}}\, \mathsf{E}_{T+1}(w_T). \tag{18}$$

According to Proposition 2, I is the informed trader's initial expectation of the signal at execution (w_T) conditional on the execution sequence, including the final SMO (hence the subscript "$T+1$" for the expectation in equation (18)). Then, the difference $w - I(\rho, w)$ can be interpreted as the signal decay between the initial submission of the BLO until after its execution. It is therefore a cost that the informed trader faces when submitting a BLO (relative to submitting a BMO). I call $\big(w - I(\rho,w)\big)\, V$ the

[40] The expectation operator E^{e} is biased, because it is taken only on a subset of the future order flow sequences (the executable ones). As a result, the law of iterated expectations does not hold. As shown below, this bias is caused by the phenomenon of "slippage."

"information decay cost," or simply the "decay cost." Corollary 5 implies that the bid-ask spread S is precisely equal to the decay cost at the threshold signal $w = \alpha$.

Corollary 6. *Let* Decay Cost$_w = \big(w - I(\rho, w)\big)\, V$ *be the information decay cost faced by an informed trader. Then, the equilibrium bid-ask spread S satisfies:*

$$S \;=\; \text{Decay Cost}_\alpha. \tag{19}$$

The intuition for this result is as follows. If the informed trader submits a BMO, she immediately captures her whole signal (w), but loses half of the bid-ask spread ($S/2$). If she submits a BLO instead, she expects the future informed traders to increase the public mean by also submitting buy orders, resulting in a decrease of her future signal. In other words, she expects that by the execution time T the signal w_T will decrease significantly. But this is exactly what the information function I measures. Thus, if the informed trader submits a BLO, she gains half of the bid-ask spread ($S/2$), but captures only part of the signal ($I(\rho, w)$). Hence, the relative payoff difference between BMO and BLO is Decay Cost$_w - S$. Since at the threshold ($w = \alpha$) the informed trader is indifferent between BMO and BLO, it follows that the equilibrium bid-ask spread is equal to the information decay cost at the threshold.

The next numerical result analyzes the connection between the bid-ask spread and the informed share (see also Figure 3).

Result 2 (Part 1)**.** The bid-ask spread S is decreasing in the informed share ρ.

The formula $S = (\alpha - I(\rho, \alpha))V$ indicates that an increase in the informed share has two opposite effects on the bid-ask spread. First, the bid-ask spread is proportional to the public volatility $V = \beta\rho^{-1}\Delta$, and therefore more informed traders cause a tighter public density and a negative effect on the bid-ask spread. Second, the bid-ask spread is proportional to the term $\alpha - I(\rho, \alpha)$, which turns out to be increasing in the informed share. Intuitively, more informed traders cause a faster rate of information decay, as the mispricing is corrected more quickly over time. But the bid-ask represents a compensation for the decay cost (see Corollary 6), hence more informed traders cause a quicker information decay and a positive effect on the bid-ask spread. According to Result 2

(Part 1), the net effect of the informed share on the bid-ask spread is negative, which is not surprising, since the public volatility V is strongly decreasing in the informed share.

To get further intuition about the bid-ask spread, I decompose it into two components. The first component, called the "slippage component," corresponds to the informed trader's information decay from the initial submission of the BLO until just before its execution by the final SMO. The second component, called the "adverse selection component," corresponds to the informed trader's information decay due to the final SMO. To define these components, I introduce two functions similar to the information function I.

Figure 3: Components of the Bid-Ask Spread
This figure plots the bid-ask spread (S), as well as the slippage component (S^s) and the adverse selection component (S^a). On the horizontal axis is the informed share $\rho = 0.05, 0.10, \ldots, 0.95$. The bid-ask spread and its components are written in units of the impact parameter Δ.

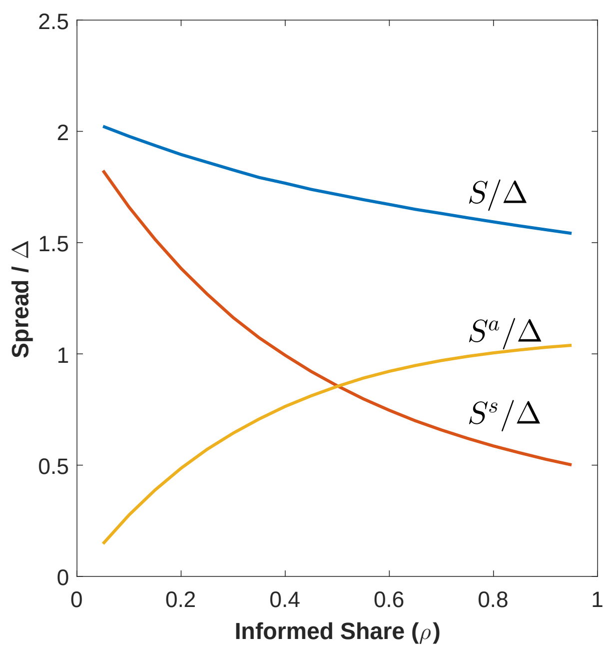

Definition 2. *For $\rho \in (0,1)$ and $w \in \mathbb{R}$, define the "slippage function" $I^s(\rho, w)$ in the same way as the information function $I(\rho, w)$ from Definition 1, except that the expression $\nu(Q) = \nu_{T+1} - \frac{\rho}{\beta}$ is replaced with $\nu(Q) = \nu_T$. Define the "adverse selection function" as the difference $I^a = I - I^s$. Define the "slippage component" S^s and the "adverse selection component" S^a as follows:*

$$S^s \;=\; \big(\alpha - I^s(\rho,\alpha)\big)\,V, \qquad S^a \;=\; S - S^s \;=\; \big(I^s(\rho,\alpha) - I(\rho,\alpha)\big)\,V. \tag{20}$$

Proposition 3 provides an interpretation of the functions I^s and I^a.

Proposition 3. *In the context of Proposition 2, the slippage function I^s and the adverse selection function I^a satisfy:*

$$(21) \qquad I^s(\rho, w) \;=\; \mathsf{E}^{\mathrm{e}}\,\mathsf{E}_T\big(w_T\big), \qquad I^a(\rho, w) \;=\; \mathsf{E}^{\mathrm{e}}\big(\mathsf{E}_{T+1}(w_T) - \mathsf{E}_T(w_T)\big).$$

Recall that the information function I can be interpreted as the informed trader's initial expectation of the signal at execution (w_T) conditional on the execution sequence including the final SMO. By Proposition 3, the slippage function I^s is the same expectation, but conditional on the execution sequence without the final SMO. The difference is the adverse selection that the informed trader faces at T from the final SMO.

Similar to Corollary 6, the next result shows that both the components of the bid-ask spread are equal to certain information decay costs. The slippage component is equal to the information decay cost until the arrival of the final SMO, while the adverse selection is equal to the information decay cost due to the final SMO.

Corollary 7. *Define the following cost functions:* Slippage $\text{Cost}_w = \big(w - I^s(\rho, w)\big)\, V$, *and* Adverse Selection $\text{Cost}_w = -I^a(\rho, w)\, V$. *Then, the two components of the bid-ask spread satisfy:*

$$(22) \qquad S^s \;=\; \text{Slippage Cost}_\alpha, \qquad S^a \;=\; \text{Adverse Selection Cost}_\alpha.$$

The next numerical result shows how the components of the bid-ask spread depend on the informed share ρ.

Result 2 (Part 2)**.** Both component S^s and S^a are positive. As functions of the informed share ρ, the slippage component S^s is decreasing in ρ, while S^a is increasing in ρ.

Figure 3 displays the bid-ask spread and its components against the informed share ρ. The bid-ask spread and its components are expressed in Δ-units, meaning that I consider the ratios S/Δ, S^s/Δ, and S^a/Δ. Using (7) and Definition 2, I compute:

$$(23) \quad \frac{S}{\Delta} \;=\; \big(\alpha - I(\rho, \alpha)\big)\beta\rho^{-1}, \quad \frac{S^s}{\Delta} \;=\; \big(\alpha - I^s(\rho, \alpha)\big)\beta\rho^{-1}, \quad \frac{S^a}{\Delta} \;=\; -I^a(\rho, \alpha)\beta\rho^{-1}.$$

The normalization by Δ does not affect the inferences, because Δ is independent of ρ (see Proposition 1). Note that all three terms in (23) contain the factor $\beta\rho^{-1} = \frac{V}{\Delta}$, which is strongly decreasing in ρ. This is because, as discussed in Section IV.A, the market is more informationally efficient when there are more informed traders, which translates into the public volatility V being smaller. One may thus expect that all three terms in (23) are decreasing in ρ. Result 2 (Part 2) shows that this is indeed true for the bid-ask spread S and the slippage component S^s, but not for the adverse selection component S^a.

The adverse selection component of the bid-ask spread, S^a, reflects the fact that the initial BLO is eventually executed by a SMO coming potentially from a future informed trader, with superior information. Thus, as expected, adverse selection increases when there are more future informed traders, that is, S^a is increasing in the informed share ρ. Moreover, S^a is close to zero when the ρ approaches zero. This is intuitive, since when there are few informed traders, there is little adverse selection.[41]

The slippage component of the bid-ask spread, S^s, reflects the phenomenon of "slippage," which is the signal decay caused by competition with the other informed traders. When there are more informed traders (the informed share ρ is higher), more informed traders are likely to arrive in the future, and therefore the *rate* at which slippage occurs is higher. The total amount of slippage, however, is multiplied by the public volatility V (recall that signals are normalized by the public volatility). Since the public volatility is strongly decreasing in the informed share, the slippage component is actually decreasing in the informed share, as can be seen in Figure 3.

The bid-ask spread S is the sum of the adverse selection component and the slippage component. In Figure 3, the spread S is indeed decreasing in the informed share ρ, although the overall effect is not as strong as the effect of ρ on public volatility. When the informed share increases from $\rho = 0.05$ to $\rho = 0.95$, the bid-ask spread decreases by about 25%. When the informed share ρ is small, the adverse selection component is close to zero, and therefore most of the bid-ask spread is determined by the slippage

[41] In the model, informed traders only observe the asset value once, when they arrive to the market. Alternatively, informed traders could be allowed to continuously observe the fundamental value. Then, the adverse selection cost would be zero, as all informed traders would have the same information. I conjecture, however, that the slippage cost would remain positive, as competition among the informed traders would still impose a cost on the submission of limit orders.

component.

Note that the slippage of limit orders can be interpreted as an endogenous waiting cost for the informed trader who decides to submit a limit order. Indeed, even though the actual waiting cost of a patient investor is zero, the informed investors' expected payoff decreases gradually over time because of slippage.[42]

The bid-ask spread can be used to construct a clean empirical proxy for the informed share ρ (which does not depend on other parameters such as the fundamental volatility σ_v or trading activity λ). Using equation (7), I compute the ratio of inter-arrival volatility $\sigma_I = \sigma_v/\sqrt{\lambda}$ to the bid-ask spread S as follows:

$$\frac{\sigma_I}{S} = \frac{\rho\sqrt{\frac{1+\gamma^2}{2\beta^2}}}{\alpha - I(\rho,\alpha)}. \tag{24}$$

By taking this ratio, the dependence of both σ_I and S on the other parameters of the model is removed. The inter-arrival volatility σ_I does not depend on ρ, while Result 2 (Part 1) implies that S is decreasing in ρ. Therefore, the ratio σ_I/S is increasing in ρ.

I summarize the effect of the informed share ρ on the first two liquidity measures: (i) the equilibrium price impact Δ is independent on ρ, while (ii) the equilibrium bid-ask spread S is decreasing in ρ. The reason for this difference is that the price impact is determined by the uninformed traders, while the bid-ask spread is determined by the informed traders. Indeed, the price impact is determined by the evolution of the public mean over time, which in turn is determined by the uninformed traders. And, because the traders with limit orders in the book behave identically whether they are informed or not (see Corollary 4), the exact breakdown between informed and uninformed traders becomes irrelevant. By contrast, the bid-ask spread is determined by the optimal order choice of the informed traders, and their order submission strategy can be shown to be elastic in the bid-ask spread. Thus, the share of informed traders affects the equilibrium bid-ask spread.

[42] A larger informed share implies higher endogenous waiting costs for an informed trader, holding the mispricing volatility constant. However, her mispricing volatility increases over time, as the informed trader gradually becomes less informed. Therefore, the exact behavior of the average waiting costs is ambiguous, and I leave this analysis for future research.

D. Resiliency

The third dimension of liquidity considered is market resiliency. Kyle (1985) defines resiliency as "the speed with which prices recover from a random, uninformative shock." Because, as shown below, in the model the speed of price correction is nonlinear in the size of the shock, I define resiliency as the rate at which a small uninformative shock is corrected, in the limit when the size of the shock approaches zero.

More formally, I define resiliency from the point of view of an econometrician who observes a small uninformative shock to the public mean. Before the shock, the econometrician has the same belief about the asset value as the uninformed traders (the public density). Suppose at date t the public mean shifts down by a small positive amount x, while the public volatility remains the same (V). The cause of this price shift is not made explicit, but one can imagine it as the reaction to the arrival of some trades that the econometrician knows to be uninformed. Therefore the econometrician knows that the shock x is uninformative, and expects the mispricing to be corrected. I then define resiliency as the rate at which the mispricing is corrected, in the limit when the shift x approaches zero.[43]

Definition 3. *Suppose before trading at t, the econometrician believes that the asset value has a normal density $v_t \sim \mathcal{N}\big(\mu_t + x, V^2\big)$, or equivalently perceives a mispricing $v_t - \mu_t \sim \mathcal{N}\big(x, V^2\big)$, with $x \in \mathbb{R}$. Denote by $f(x)$ the average mispricing $v_{t+1} - \mu_{t+1}$ after observing the order at t. The "market resiliency coefficient" K is defined by:*

$$K \;=\; 1 - f'(0). \tag{25}$$

Intuitively, if the order at t comes from an informed trader, she is more likely to observe a positive mispricing, just as the econometrician does. Hence, she is more likely to submit a buy order, which pushes up the public mean and reduces the mispricing. Thus, the econometrician expects the forecast error to become smaller on average, which

[43] Note that market resiliency is defined in the context of a stationary equilibrium, in which the public volatility (the uncertainty about v_t) is constant and equal to V. In this sense, there is an average mispricing which is never fully corrected: informed traders reduce the mispricing over time, but diffusion in v restores the mispricing. Nevertheless, the definition of resiliency in this section is based on the correction of an *additional* mispricing x by the informed traders.

for a positive shock x translates into $0 < f(x) < x$. If f is linearized near $x = 0$, one gets $f(x) \approx f'(0)x = (1-K)x$. Hence, K is the rate at which the mispricing x gets corrected when x is small, which is indeed the definition of resiliency. Note that the mechanism behind resiliency is essentially the same as the mechanism behind slippage: the existence of informed traders corrects mispricings over time (resiliency), which generates a cost for an informed trader who submits a limit order (slippage).

Proposition 4. *The market resiliency coefficient equals:*

$$K \; = \; \frac{2\big(\gamma\phi(0) + (1-\gamma)\phi(\alpha)\big)}{\beta}\,\rho^2 \;\approx\; 0.8606\,\rho^2. \tag{26}$$

Proposition 4 implies that the market is more resilient when the informed share is larger. This confirms the intuition that a larger share of informed traders results in a faster correction of pricing errors. However, even if the informed share ρ is very close to one, there is an upper bound on how quickly the mispricing is corrected. This is because each informed trader has a threshold strategy, and therefore her information cannot be fully revealed.[44]

Market resiliency is related to information efficiency. Indeed, the market resiliency coefficient K from equation (26) is proportional to the information precision measure, $1/V^2$. Thus, in the model, a larger share of informed traders ρ causes the market to be both more resilient and more informationally efficient.

V. Non-Stationary Equilibria

In the stationary equilibrium of Section III, the public volatility is constant and equal to V. In this section, the public volatility can take a different initial value than V. This could happen, for instance, if an uncertainty shock (an unobserved shock to the fundamental value) suddenly pushes the public volatility above V. Then, the equilibrium is fully determined by the initial value of the public density σ_0, or equivalently by the

[44]For example, after a BMO, the uninformed traders cannot infer the informed trader's signal $w_t = \frac{v_t - \mu_t}{V}$, they can only infer that her signal belongs to the interval (α, ∞).

initial value of the “normalized public volatility,” which is:

$$\theta_0 \;=\; \frac{\sigma_0}{V}. \tag{27}$$

I thus define a “non-stationary equilibrium” as an equilibrium for which the initial normalized public volatility θ_0 is different from one.

A. Properties of Non-Stationary Equilibria

In Definition A1 in the Appendix, I introduce several new parameters: $\tilde{\alpha}$, $\tilde{\beta}$, $\tilde{\gamma}$, $\tilde{I}$, $\tilde{V}$, $\tilde{\Delta}$, $\tilde{S}$, which are all functions of two variables, ρ and θ. In addition, $\tilde{I}$ is also a function of a third variable, w. Intuitively, one thinks of ρ as the informed share; of θ as the normalized public volatility, and of w as the signal, or normalized asset value, $(v-\mu)/V$. However, just like the corresponding parameters from Section III.B (without the tilde above them), the new parameters are defined completely formally.

Proposition 5 shows that for any initial normalized public volatility θ_0 there exists a non-stationary Markov Perfect Equilibrium (MPE) of the model, as long as the conditions in Result IA.4 in the Internet Appendix (Section 5) are satisfied. I verify these conditions numerically in Section 6 in the Internet Appendix, for θ sufficiently close to one. The next result also describes several properties of the equilibrium.

Proposition 5. *Let $\theta_0 > 0$. If the conditions in Result IA.4 are satisfied, there exists a Markov Perfect Equilibrium of the game in which the initial normalized public volatility is θ_0. In equilibrium, the normalized public volatility $\theta_t = \sigma_t/V$ evolves according to:*

$$\theta_{t+1}^2 \;=\; \rho^2\,\frac{1+\gamma^2}{2\beta^2} + \theta_t^2 - 2\rho^2\theta_t^2\left(\frac{\left(\phi\left(\frac{\tilde{\alpha}_t}{\theta_t}\right)\right)^2}{\frac{1-\rho}{4}+\rho\left(1-\Phi\left(\frac{\tilde{\alpha}}{\theta_t}\right)\right)} + \frac{\left(\phi\left(\frac{0}{\theta_t}\right)-\phi\left(\frac{\tilde{\alpha}_t}{\theta_t}\right)\right)^2}{\frac{1-\rho}{4}+\rho\left(\Phi\left(\frac{\tilde{\alpha}}{\theta_t}\right)-\Phi(0)\right)}\right), \tag{28}$$

where $\tilde{\alpha}_t = \tilde{\alpha}(\rho,\theta_t)$. Let $\tilde{\gamma}_t = \tilde{\gamma}(\rho,\theta_t)$, $\tilde{\Delta}_t = \tilde{\Delta}(\rho,\theta_t)$, and $\tilde{S}_t = \tilde{S}(\rho,\theta_t)$. Then, an order arriving at $t \geq 0$ changes the public mean from μ_t to (i) $\mu_t + \tilde{\Delta}_t$ if the order is BMO, *(ii) $\mu_t + \tilde{\gamma}_t\tilde{\Delta}_t$ if the order is* BLO, *(iii) $\mu_t - \tilde{\gamma}_t\tilde{\Delta}_t$ if the order is* SLO, *and (iv) $\mu_t - \tilde{\Delta}_t$ if the order is* SMO. *At date t, the bid-ask spread is $\tilde{S}_t$, the ask price is $\mu_t + \tilde{S}_t/2$, and the bid price is $\mu_t - \tilde{S}_t/2$.*

In equilibrium, the bid-ask spread and the price impact of an order are no longer constant. The next result, however, provides a linear combination that remains constant over time. The result involves the parameters S and Δ from equation (7).

Corollary 8. *The equilibrium bid-ask spread $\tilde{S}_t$ and price impact coefficient $\tilde{\Delta}_t$ satisfy:*

$$\frac{\tilde{S}_t}{2} - \tilde{\Delta}_t = \frac{S}{2} - \Delta. \tag{29}$$

Equation (29) is the indifference condition for the uninformed traders. Consider an uninformed trader who is the first in the bid queue, and suppose his BLO is executed at date t by a SMO. Then, net of his private valuation, his expected payoff is $\tilde{S}_t/2 - \tilde{\Delta}_t$, where $\tilde{S}_t/2$ represents the difference between the public mean and the bid price, and $\tilde{\Delta}_t$ represents the adverse selection loss from a potentially informed SMO. If his expected payoff were not the same at $t+1$, then the uninformed traders would have an incentive to modify their position in the bid queue. The discussion thus far explains why the expected payoff $\tilde{S}_t/2 - \tilde{\Delta}_t$ is constant. That the constant is equal to $S/2 - \Delta$ is due to the fact that the equilibrium converges to the stationary equilibrium of Section III. I state this as a numerical result.

Result 3. As t becomes large, the public volatility σ_t approaches the parameter V, the bid-ask spread $\tilde{S}_t$ approaches S, and the price impact coefficient $\tilde{\Delta}_t$ approaches Δ.

B. Market Quality in Non-Stationary Equilibria

I now describe non-stationary equilibria in more detail, in particular regarding the market quality measures introduced in Section IV: information efficiency, price impact, bid-ask spread, and market resiliency. To obtain other testable predictions, I also analyze observable measures such as the limit-to-market impact ratio (the price impact ratio of a limit order to a market order) and the market-to-limit probability ratio (the probability ratio of the next order being a market order or a limit order).

First, I analyze a measure that is specific to non-stationary equilibria: the speed of convergence to the stationary equilibrium. Intuitively, this speed is related to the resiliency of certain market quality measures, such as public volatility, bid-ask spread, or price impact. Indeed, after an uncertainty shock that raises the initial normalized public volatility θ above one, Result 3 above shows that θ (as well as the bid-ask spread and the price impact) reverts to its stationary value at a certain rate. It is then perhaps not surprising that this speed of convergence is closely connected to the previous measure of market (or price) resiliency, which is the rate at which a mispricing is corrected.

Figure 4: Dynamic Information Efficiency
This figure plots the time evolution of two market quality measures in non-stationary equilibria. On the horizontal axis is time. Each curve in each graph corresponds to a value of the informed share (ρ) ranging from 0.05 to 0.95 (the right plot considers the subset where ρ ranges from 0.40 to 0.55). The left graph plots the normalized public volatility (θ_t), which is the public volatility (σ_t) divided by its stationary value (V); time on the horizontal axis is in logarithmic scale and is shifted by one, such that time 1 corresponds to $t = 0$ in the model; the initial normalized public volatility in all cases is $\theta_0 = 2$. The right graph plots the bid-ask spread ($\tilde{S}_t$) in units of the impact parameter Δ; the initial (non-normalized) public volatility in all cases is $\sigma_0 = 2\Delta$. In the right plot, the time scale is linear and starts from $t = 0$, as in the model.

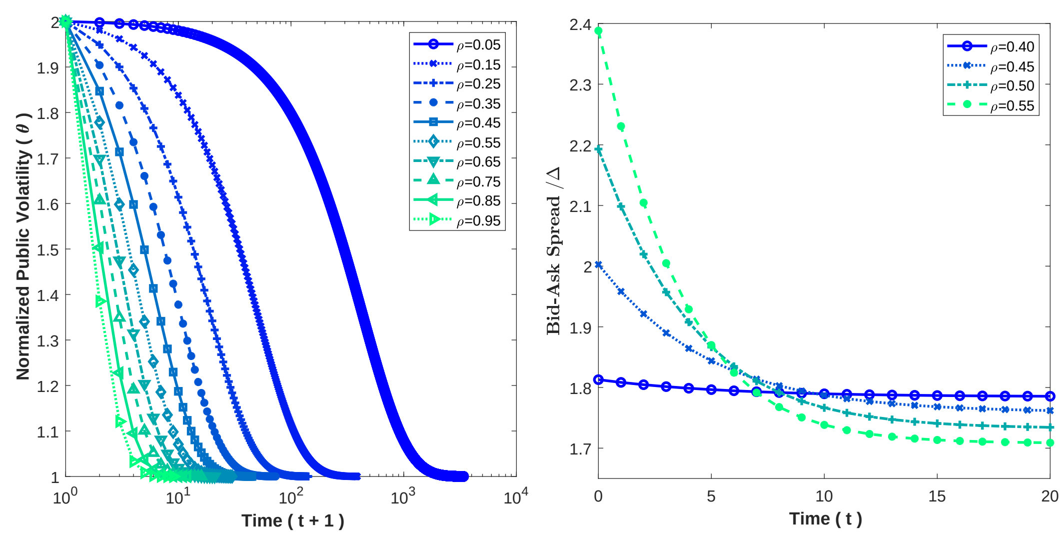

Figure 4 (left graph) displays the evolution over time of the normalized public volatility θ_t according to equation (28), starting from an initial value $\theta_0 = 2$. Each curve in the plot corresponds to an informed share ρ ranging from 0.05 to 0.95. One observes that in all cases the normalized public volatility indeed converges to one, and furthermore, that the speed of convergence is inversely related to the informed share.

More formally, I define the "recovery time" as the number of trading rounds it takes for the normalized public volatility to revert within a neighborhood of one after a positive or negative shock. In Figure 4, I choose as neighborhood of one the interval $(1-\varepsilon, 1+\varepsilon)$, with $\varepsilon = 10^{-4}$. Numerically, the recovery time appears linear in the inverse informed share, $1/\rho^2$, regardless of the choice of neighborhood or shock size. I report this fact as a numerical result.

Result 4. The recovery time after a shock to the normalized public volatility ($\theta_t = \sigma_t/V$) is linear in the inverse squared informed share ($1/\rho^2$).

This result confirms the previous intuition that informed traders make the market more dynamically efficient. Indeed, when there are more informed traders (the informed share is higher), a shock that moves the public volatility away from its stationary value ($\theta = 1$) is followed by a quicker reversal to the stationary value, and hence it requires a shorter recovery time. The quicker convergence occurs because orders carry more information when the informed share is higher, since the probability of each order being submitted by an informed trader is higher.

The inverse recovery time is thus a measure of information efficiency, and Result 4 shows that this measure is linear in the squared informed share (ρ^2). In Section IV.A, another measure of information efficiency is the inverse stationary public variance ($1/V^2$), which is also linear in the squared informed share.[45] The two measures share the same dynamic efficiency intuition, but the inverse recovery time measure is more explicit in how dynamic efficiency is achieved.

I now discuss the bid-ask spread $\tilde{S} = \tilde{S}(\rho, \theta)$ and the price impact coefficient $\tilde{\Delta} = \tilde{\Delta}(\rho, \theta)$. From the results of the previous section, the two measures are connected by $\tilde{S}/2 - \tilde{\Delta} = S/2 - \Delta$.[46] Thus, they have a similar evolution over time. Moreover, as shown in Result 5 below, both are increasing functions of θ.

Figure 4 (right graph) displays the evolution over time of the bid-ask spread $\tilde{S}_t$ if one starts with the same (non-normalized) public volatility $\sigma_0 = 2\Delta$.[47] Intuitively, this graph shows the effect of an absolute uncertainty shock at $t = 0$ on the bid-ask spread $\tilde{S}$, and how that effect depends on the information share ρ. Initially, a higher informed share makes the bid-ask spread $\tilde{S}$ jump to a higher value. This is because there is more

[45] By equation (12), the inverse stationary public variance is proportional to ρ^2.

[46] More formally, Corollary 8 translates into $\tilde{S}(\rho, \theta_t)/2 - \tilde{\Delta}(\rho, \theta_t) = S(\rho)/2 - \Delta(\rho)$. By setting $\theta_0 = \theta$, it follows that $\tilde{S}/2 - \tilde{\Delta} = S/2 - \Delta$ for any $\theta > 0$ and $\rho \in (0, 1)$.

[47] I only consider the values $\rho = 0.40$ to $\rho = 0.55$ because I want $\theta_0 = \sigma_0/V \in (1, 1.5]$. I require $\theta_0 > 1$ because I want a positive shock to the normalized public volatility, and I require $\theta_0 \leq 1.5$ because the numerical algorithm has only been made to work for $0.5 \leq \theta_0 \leq 1.5$.

static adverse selection when there are more informed traders. Over time, however, a higher informed share pushes the bid-ask spread to a lower value, as the market is dynamically more efficient. Indeed, the bid-ask spread $\tilde{S}$ converges over time to the stationary bid-ask spread S (see Result 3), which is decreasing in the informed share ρ (see Result 2).

Figure 5: Market Quality in Non-Stationary Equilibria
This figure plots three market quality measures in non-stationary equilibria. On the horizontal axis is the normalized public volatility (θ), which is the public volatility (σ) divided by its stationary value (V). Each curve in each graph corresponds to a value of the informed share (ρ) ranging from 0.05 to 0.95. The left graph plots the relative price impact coefficient, which is the price impact coefficient $\tilde{\Delta} = \tilde{\Delta}(\rho, \theta)$ divided by its stationary value Δ. The middle graph plots the limit-to-market impact ratio $\tilde{\gamma}$, which is equal to the price impact of a limit order divided by the price impact of a market order ($\tilde{\gamma}\tilde{\Delta}/\tilde{\Delta}$); the horizontal line corresponds to the equilibrium value $\gamma \approx 0.2554$ in the stationary equilibrium. The right graph plots the market-to-limit probability ratio $\frac{P_{\text{MO}}}{P_{\text{LO}}}(\rho, \theta)$, which is the probability the next order is a market order, divided by the probability that the next order is a limit order. If the numerical procedure does not yield a unique value, the corresponding point in the plot is omitted.

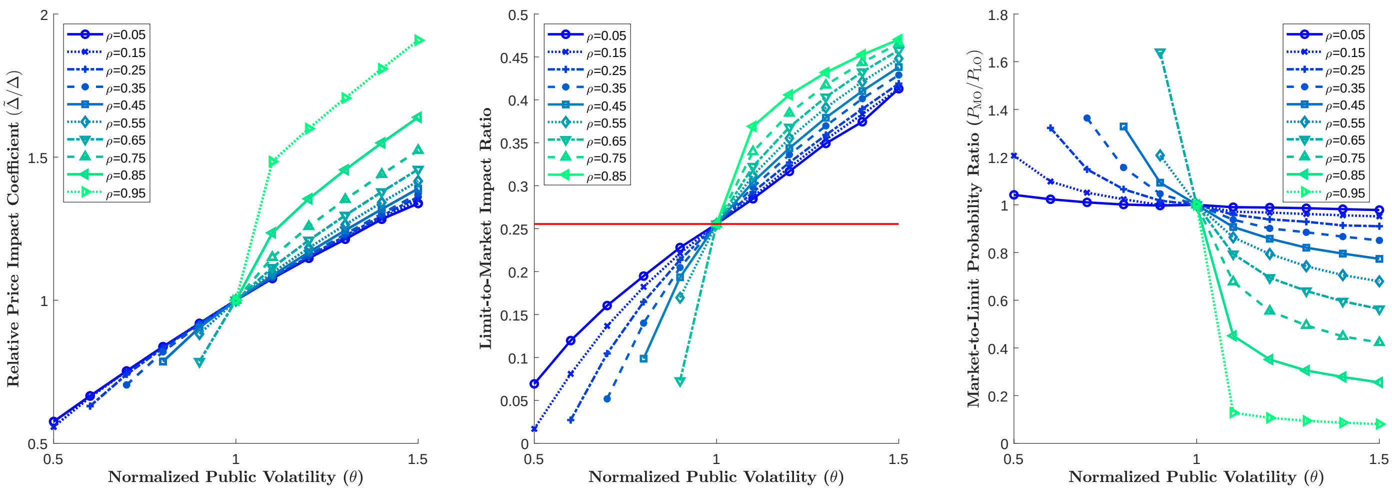

Because the price impact and bid-ask spread depend on the normalized public volatility θ in the same way, I can focus on either of these measures. I thus consider the price impact $\tilde{\Delta}$, or equivalently the "relative price impact coefficient," which is $\tilde{\Delta}$ divided by its stationary value, the parameter Δ. Equation (A-22) in the Appendix then implies that the relative price impact coefficient is:

$$\frac{\tilde{\Delta}}{\Delta} = \frac{\tilde{\beta}(\rho, \theta)}{\beta}\,\theta. \tag{30}$$

Figure 5 (left graph) plots the dependence of the relative price impact coefficient on both ρ and θ.[48] Each curve in the plot corresponds to a value of the informed share ρ ranging from 0.05 to 0.95. Note that in all cases the relative price impact is increasing in θ. I report this fact as a numerical result. Since Δ does not depend on either ρ or θ, the next result is equally true for the price impact coefficient $\tilde{\Delta}$. Also, equation (8) shows that the same is true for the bid-ask spread $\tilde{S}$.

Result 5. The price impact coefficient ($\tilde{\Delta}$) and the bid-ask spread ($\tilde{S}$) are increasing in the normalized public volatility (θ).

Intuitively, when the normalized public volatility θ is larger, the uninformed traders have imprecise knowledge about the fundamental value, and therefore the adverse selection is stronger. That implies that the price impact of a buy market order, $\tilde{\Delta}$, is larger, as confirmed by Result 5. The bid-ask spread, $\tilde{S}$, is also larger, to compensate the uninformed traders for the increase in adverse selection. Formally, the bid-ask spread and price impact vary with θ in the same way, since equation (29) implies that the difference $\tilde{S}/2 - \tilde{\Delta}$ is equal to $S/2 - \Delta$, which does not depend on θ.

Putting together the previous results, it follows that after a positive shock in the public volatility (or equivalently in θ), the bid-ask spread $\tilde{S}$ initially increases, to adjust for the higher value of θ, after which it decreases gradually to its stationary value S. The same effect occurs for the price impact $\tilde{\Delta}$. Thus, in the model, the bid-ask spread and the price impact coefficient both display resiliency, in the sense that they eventually recover to their stationary values after a shock in the public volatility. I call this phenomenon "liquidity resiliency."

Liquidity resiliency is different from market resiliency. As discussed in Section IV.D, market resiliency is defined as the recovery of prices after an uninformative shock. In the context of non-stationary equilibria, I obtain the following result similar to Proposition 4.

[48]The figure displays results computed with the function I instead of $\tilde{I}$; very similar results are obtained by using instead the estimated function $\tilde{I}$. The numerical procedure used to solve for the equilibrium is explained in Section 6 in the Internet Appendix. I impose the strict condition that the solution to the first equation in (A-21) must be unique. When the threshold $\tilde{\alpha}$ is close to zero, this condition is not satisfied because of estimation errors. This explains why there are missing points in Figure 5 when ρ is large and θ is small. Intuitively, this occurs because the increase in adverse selection makes the indifference condition (29) for the uninformed traders harder to satisfy.

Proposition 6. *The equilibrium market resiliency coefficient $\tilde{K} = \tilde{K}(\rho, \theta)$ satisfies:*

$$\tilde{K} \;=\; \frac{2\rho^2}{\theta}\,\frac{\tilde{\gamma}\phi(0) + (1-\tilde{\gamma})\phi\big(\frac{\tilde{\alpha}}{\theta}\big)}{\tilde{\beta}}. \tag{31}$$

Numerically, the market resiliency coefficient $\tilde{K}$ is increasing in the informed share ρ, and decreasing in the normalized public volatility θ. The intuition for why market resiliency is increasing in the informed share is the same as in the stationary equilibrium. The new result is that market resiliency is decreasing in the normalized public volatility. Intuitively, when the public volatility is large, the informed traders become less aggressive and are more likely to submit limit orders (as explained below). Therefore, it takes longer for the price to converge to the fundamental value.

I introduce a new market of market quality, the "market-to-limit probability ratio," which is defined as the probability the next order is a market order, divided by the probability that the next order is a limit order. In the stationary equilibrium of Section III, this ratio is equal to one since all types of orders are equally likely (see Corollary 3). In non-stationary equilibria, the market-to-limit probability ratio varies with both the informed share and the normalized public volatility.

Proposition 7. *The market-to-limit probability ratio is equal to:*

$$\frac{P_{\mathrm{MO}}}{P_{\mathrm{LO}}} \;=\; \frac{\frac{1-\rho}{4} + \rho\big(1 - \Phi\big(\frac{\tilde{\alpha}}{\theta}\big)\big)}{\frac{1-\rho}{4} + \rho\big(\Phi\big(\frac{\tilde{\alpha}}{\theta}\big) - \Phi(0)\big)}. \tag{32}$$

Figure 5 (right graph) plots the dependence of the market-to-limit probability ratio on both ρ and θ. Each curve corresponds to a value of the informed share ρ ranging from 0.05 to 0.95. In all cases, the market-to-limit probability ratio is decreasing in θ. Intuitively, as explained before, when the normalized public volatility θ is larger, there is an increase in adverse selection for the uninformed traders. This causes the bid-ask spread, as well as the price impact, to be larger. But the increase in the bid-ask spread changes the informed traders' tradeoff between market orders and limit orders, and makes limit orders more attractive. Thus, the market-to-limit probability ratio is smaller when the public volatility is larger. This result, along with the previous results regarding the resiliency of the bid-ask spread and the price impact after a public volatility shock provide new testable empirical implications.

Figure 5 (middle graph) plots a related measure, the "limit-to-market impact ratio" $\tilde{\gamma}$, which is the ratio of the price impact of a buy limit order ($\tilde{\gamma}\tilde{\Delta}$) to the price impact of a buy market order ($\tilde{\Delta}$). In all cases, the limit-to-market impact ratio is increasing in θ. The intuition is based on the fact that limit orders are relatively more likely when θ is larger, which implies that their price impact is also larger. This result is however dependent on the public density being normal, and thus might be considered less robust.

VI. Conclusion

I have presented a model of a limit order market with asymmetric information, in which investors can choose between demanding liquidity (with a market order) and providing liquidity (with a limit order). Despite the difficulty of the problem, the model is tractable, and, except for an information function that must be computed numerically, the results are obtained in closed form.

The main result is that informed trading, as proxied in the model by the informed share, has an overall positive effect on liquidity, under its three dimensions: tightness (bid-ask spread), depth (price impact), and resiliency (the speed at which pricing errors are corrected). In particular, a larger informed share (i) leads to a smaller bid-ask spread, (ii) generates a stronger market resiliency, yet (iii) does not affect the price impact of one additional unit of trading. From the perspective of the informed trader, limit orders have a slippage cost, which measures the erosion in information advantage due to the competition from future informed traders. Slippage costs represent an endogenous waiting cost for informed traders, and generate a new component of the bid-ask spread.

I also estimate the information content of order flow. In particular, because in equilibrium informed traders also use limit orders—whereas in much of the theoretical literature informed traders only use market orders,—in the model limit orders also have a nonzero price impact. Quantitatively, the price impact of a limit order is roughly one fourth of the price impact of a market order.

The results described thus far are true in the context of the stationary equilibrium, in which the public volatility is constant. If an uncertainty shock suddenly increases the public volatility, the results predict that the public volatility (as well as the bid-ask spread and price impact) decrease over time toward the stationary equilibrium value, at a speed that is increasing in the informed share. I introduce a new measure, the market-to-limit ratio, which measures the probability of a trader to submit a market order relative to a limit order. After an uncertainty shock, the market-to-limit ratio drops significantly below one, as the increase in the bid-ask spread temporarily convinces the informed traders to submit more limit orders. The connections among the market-to-limit ratio with the liquidity measures and the public volatility, as well as the expected evolution of the equilibrium towards the stationary one, produce new testable implications of the model.

The results show that informed trading has an important effect on liquidity, especially under its resiliency aspect. But estimating market resiliency directly is difficult, since that would involve having access to information that is not public. Instead, the results regarding non-stationary equilibria suggest that one can use an estimate of liquidity resiliency, which is observable as long as the uncertainty shocks can be identified.

Yet another approach is to use rigidities such as stale prices as evidence of low market resiliency, and study the connection with informed trading. I argue that market resiliency is inversely related to the price delay measure of Hou and Moskowitz (2005, in short HM05). HM05 find empirically that firms in which the price responds with a delay to information command a large return premium.[49] Interestingly, HM05 find that the delay premium has little relation with the PIN measure of Easley et al. (2002), which is another measure of informed trading. This suggests that the informed share in my model may in fact be measuring a different aspect of informed trading than PIN. Since PIN is based on large imbalances between buyers and sellers, I postulate that PIN is related to informed trading done by large traders, possibly corporate insiders. By contrast, the informed share in my model may be more related to trading done by small informed traders that are not necessarily insiders, and are just better informed than the public.

[49]Indeed, it is plausible that firms in which prices respond with a delay to information are also firms for which prices move more slowly toward the fundamental value. It is true that HM05 consider delay at weekly (or in some robustness checks at daily) frequency, while in my model it is more natural to think of events as occurring at higher, intra-day frequencies. Then, my identification is correct if delay at lower frequencies is correlated with delay at higher frequencies.

Overall, my theoretical model produces a rich set of implications regarding the connection between the activity of informed traders and the level of liquidity. Informed traders have on aggregate a positive effect, by making the market more efficient and, at the same time, more liquid. A welfare analysis also shows that a larger number of informed traders (caused for instance by an exogenous decrease in information costs) increases aggregate trader welfare. The model thus provides useful tools to analyze informed trading, and its connection with liquidity, prices, and welfare.

Appendix. Proofs

Before proving Theorem 1, I explain how investors' beliefs are updated after observing the order flow. For an order $\mathcal{O} = \{\text{BMO}, \text{BLO}, \text{SLO}, \text{SMO}\}$, define, respectively, $\delta_{\mathcal{O}} \in \left\{\frac{\rho}{\beta}, \gamma\frac{\rho}{\beta}, -\gamma\frac{\rho}{\beta}, -\frac{\rho}{\beta}\right\}$, and $i_{\mathcal{O}} \in \{(\alpha, \infty), (0, \alpha), (-\alpha, 0), (-\infty, -\alpha)\}$. Let $\phi(\cdot; m, s)$ be the normal density with mean m and standard deviation s, $\phi(\cdot)$ the standard normal density ($m = 0,\ s = 1$), and $\Phi(\cdot)$ the cumulative normal density. Denote the normalized inter-arrival volatility by:

$$\hat{\sigma_I} \;=\; \frac{\sigma_I}{V} \;=\; \rho\sqrt{\frac{1+\gamma^2}{2\beta^2}}. \tag{A-1}$$

Lemma A1. *In the context of Theorem 1, consider a trader who, before trading at t, believes that the signal $w_t = \frac{v_t - \mu_t}{V} = z$ has probability density function $g_t(z)$. Then, the following are true:*

(a) The probability of observing $\mathcal{O}$ at t is:

$$P_{\mathcal{O}} \;=\; \frac{1-\rho}{4} + \rho\int_{z\in i_{\mathcal{O}}} g_t(z)\mathrm{d}z. \tag{A-2}$$

After seing the order $\mathcal{O}$ at t, the posterior density of $w_{t+1} = \frac{v_{t+1}-\mu_{t+1}}{V}$ is:

$$g_{t+1,\mathcal{O}}(x) \;=\; \frac{\int\left(\frac{1-\rho}{4} + \rho\mathbf{1}_{z\in i_{\mathcal{O}}}\right)g_t(z)\phi\big(x; z-\delta_{\mathcal{O}}, \hat{\sigma_I}\big)\mathrm{d}z}{P_{\mathcal{O}}}. \tag{A-3}$$

(b) Suppose g_t is not necessarily normal, and has mean ν_t and standard deviation τ_t.

Define the "normalized price impact" $\delta_{t+1,\mathcal{O}}$ as the change in the expectation of w_t after observing $\mathcal{O}$ at t. It satisfies:

$$\text{(A-4)} \qquad \delta_{t+1,\mathcal{O}} \;=\; \mathsf{E}(w_t \mid g_t, \mathcal{O}) - \mathsf{E}(w_t \mid g_t) \;=\; \frac{\rho \int_{i_{\mathcal{O}}} g_t(z)(z-\nu_t)\mathrm{d}z}{P_{\mathcal{O}}}.$$

Denote by $\nu_{t+1,\mathcal{O}}$ and $\tau_{t+1,\mathcal{O}}$ the mean and standard deviation, respectively, of the posterior density $g_{t+1,\mathcal{O}}(x)$. Let $V_{t+1,\mathcal{O}} = \frac{1}{P_{\mathcal{O}}} \int g_t(z)\big(\frac{1-\rho}{4}+\rho \mathbf{1}_{z\in i_{\mathcal{O}}}\big)\Big(\big(\frac{z-\nu_t}{\tau_t}\big)^2 - 1\Big)\,\mathrm{d}z$. The following formulas hold:

$$\text{(A-5)} \quad \nu_{t+1,\mathcal{O}} \;=\; \nu_t - \delta_{\mathcal{O}} + \delta_{t+1,\mathcal{O}}, \qquad \tau^2_{t+1,\mathcal{O}} \;=\; \tau_t^2(1+V_{t+1,\mathcal{O}}) + \hat{\sigma_I}^2 - \delta^2_{t+1,\mathcal{O}},$$

$$\text{(A-6)} \qquad \mathsf{E}(w_t \mid g_t, \mathcal{O}) \;=\; \nu_{t+1,\mathcal{O}} + \delta_{\mathcal{O}} \;=\; \nu_t + \delta_{t+1,\mathcal{O}}.$$

Let $\bar{\nu}_{t+1} = \mathsf{E}_{\mathcal{O}}\big(\nu_{t+1,\mathcal{O}}\big)$ and $\bar{\tau}^2_{t+1} = \mathsf{E}_{\mathcal{O}}\big(\tau^2_{t+1,\mathcal{O}}\big)$, where $\mathsf{E}_{\mathcal{O}}$ represents the average over $\mathcal{O} \in \{\text{BMO}, \text{BLO}, \text{SLO}, \text{SMO}\}$, with weights $P_{\mathcal{O}}$. Then, $\mathsf{E}_{\mathcal{O}}(\delta_{t+1,\mathcal{O}}) = \mathsf{E}_{\mathcal{O}}(V_{t+1,\mathcal{O}}) = 0$, and:

$$\text{(A-7)} \qquad \bar{\nu}_{t+1} \;=\; \nu_t - \mathsf{E}_{\mathcal{O}}\,\delta_{\mathcal{O}}, \qquad \bar{\tau}^2_{t+1} \;=\; \tau_t^2 + \hat{\sigma_I}^2 - \mathsf{E}_{\mathcal{O}}\,\delta^2_{t+1,\mathcal{O}}.$$

(c) If $g_t = \mathcal{N}\big(\nu_t, \tau_t^2\big)$ is normal, let $i_{\mathcal{O}} = (L_{\mathcal{O}}, H_{\mathcal{O}})$, $\ell_{\mathcal{O}} = \frac{L_{\mathcal{O}}-\nu_t}{\tau_t}$, $h_{\mathcal{O}} = \frac{H_{\mathcal{O}}-\nu_t}{\tau_t}$. Then,

$$\begin{aligned} P_{\mathcal{O}} \;&=\; \frac{1-\rho}{4} + \rho\big(\Phi(h_{\mathcal{O}}) - \Phi(\ell_{\mathcal{O}})\big), \\ \text{(A-8)} \qquad \nu_{t+1,\mathcal{O}} \;&=\; \nu_t - \delta_{\mathcal{O}} + \delta_{t+1,\mathcal{O}}, \quad \tau^2_{t+1,\mathcal{O}} \;=\; \tau_t^2(1+V_{t+1,\mathcal{O}}) + \hat{\sigma_I}^2 - \delta^2_{t+1,\mathcal{O}}, \\ \delta_{t+1,\mathcal{O}} \;&=\; \frac{\rho\tau_t\big(\phi(\ell_{\mathcal{O}}) - \phi(h_{\mathcal{O}})\big)}{P_{\mathcal{O}}}, \quad V_{t+1,\mathcal{O}} \;=\; \frac{\rho\big(\ell_{\mathcal{O}}\phi(\ell_{\mathcal{O}}) - h_{\mathcal{O}}\phi(h_{\mathcal{O}})\big)}{P_{\mathcal{O}}}. \end{aligned}$$

If one writes $\bar{\nu}_{t+1} = f(\nu_t)$, then the derivative of f is:

$$\text{(A-9)} \qquad f'(\nu_t) \;=\; 1 - \frac{\rho^2}{\beta\tau_t}\left(2\gamma\phi\Big(\frac{\nu_t}{\tau_t}\Big) + (1-\gamma)\Big(\phi\Big(\frac{\alpha+\nu_t}{\tau_t}\Big) + \phi\Big(\frac{\alpha-\nu_t}{\tau_t}\Big)\Big)\right).$$

(d) If $g_t = \mathcal{N}(0,1)$ is the standard normal density, with $\nu_t = 0$ and $\tau_t = 1$, then for

all orders $\mathcal{O}$ at t:

$$\text{(A-10)} \qquad P_{\mathcal{O}} = \frac{1}{4}, \qquad \delta_{t+1,\mathcal{O}} = \delta_{\mathcal{O}}, \qquad \nu_{t+1,\mathcal{O}} = 0, \qquad \bar{\tau}_{t+1} = 1.$$

Hence, the normalized density g_t has constant volatility.[50]

Proof. Conditional on observing $w_t = \frac{v_t - \mu_t}{V} = z$, the probability of an order $\mathcal{O}$ at t is $\mathsf{P}\big(\mathcal{O}_t = \mathcal{O} \mid w_t = z\big) = (1-\rho)\frac{1}{4} + \rho \mathbf{1}_{z \in i_{\mathcal{O}}}$. Indeed, if the trader at t is uninformed (with probability $1-\rho$), he submits an order $\mathcal{O}$ with equal probability $\frac{1}{4}$; if the trader at t is informed (with probability ρ), she submits an order $\mathcal{O}$ if and only if $z \in i_{\mathcal{O}}$. Integrating over z, one obtains $P_{\mathcal{O}} = \frac{1-\rho}{4} + \rho \int_{z \in i_{\mathcal{O}}} g_t(z) \mathrm{d}z$, which proves (A-2).

I now compute the density of the normalized asset value at $t+1$ after observing an order $\mathcal{O}$ at t. Immediately after t the public mean moves to $\mu_{t+1} = \mu_t + \Delta_{\mathcal{O}}$, where $\Delta_{\mathcal{O}} \in \{\Delta, \gamma\Delta, -\gamma\Delta, -\Delta\}$. Since $\frac{\Delta}{V} = \frac{\rho}{\beta}$, note that $\delta_{\mathcal{O}} = \frac{\Delta_{\mathcal{O}}}{V} \in \Big\{\frac{\rho}{\beta}, \gamma\frac{\rho}{\beta}, -\gamma\frac{\rho}{\beta}, -\frac{\rho}{\beta}\Big\}$. If $z = w_t$ and $\delta_v = \frac{v_{t+1} - v_t}{V}$, write $x = w_{t+1} = \frac{v_{t+1} - (\mu_t + \Delta_{\mathcal{O}})}{V} = \delta_v + z - \delta_{\mathcal{O}}$. But δ_v has a normal distribution given by $\mathcal{N}\big(0, \frac{\sigma_I^2}{V^2}\big) = \mathcal{N}(0, \hat{\sigma_I}^2)$, hence $\mathsf{P}\big(w_{t+1} = x \mid \mathcal{O}_t = \mathcal{O}, w_t = z\big) = \mathsf{P}\big(\delta_v = x - z + \delta_{\mathcal{O}}\big) = \phi(x - z + \delta_{\mathcal{O}}; 0, \hat{\sigma_I}) = \phi(x; z - \delta_{\mathcal{O}}, \hat{\sigma_I})$. Compute also $\mathsf{P}\big(w_{t+1} = x, \mathcal{O}_t = \mathcal{O} \mid w_t = z\big) = \mathsf{P}\big(w_{t+1} = x \mid \mathcal{O}_t = \mathcal{O}, w_t = z\big)\, \mathsf{P}\big(\mathcal{O}_t = \mathcal{O} \mid w_t = z\big) = \phi(x; z - \delta_{\mathcal{O}}, \hat{\sigma_I})\big(\frac{1-\rho}{4} + \rho\mathbf{1}_{z \in i_{\mathcal{O}}}\big)$. Thus, the posterior density is $g_{t+1,\mathcal{O}}(x) = \mathsf{P}(w_{t+1} = x \mid w_t \sim g_t(z), \mathcal{O}_t = \mathcal{O}) = \frac{\int \mathsf{P}(w_{t+1} = x, \mathcal{O}_t = \mathcal{O} \mid w_t = z) g_t(z) \mathrm{d}z}{\int \mathsf{P}(\mathcal{O}_t = \mathcal{O} \mid w_t = z) g_t(z) \mathrm{d}z} = \frac{\int (\frac{1-\rho}{4} + \rho \mathbf{1}_{z \in i_{\mathcal{O}}}) \phi(x; z - \delta_{\mathcal{O}}, \hat{\sigma_I}) g_t(z) \mathrm{d}z}{P_{\mathcal{O}}}$. This proves (A-3).

To prove part (b), start by computing as above $\mathsf{P}\big(w_t = z \mid \mathcal{O}_t = \mathcal{O}\big) = \frac{\frac{1-\rho}{4} + \rho\mathbf{1}_{z \in i_{\mathcal{O}}}}{P_{\mathcal{O}}}$. Multiplying by z and integrating, one gets $\mathsf{E}\big(w_t \mid g_t, \mathcal{O}\big) = \frac{\int z (\frac{1-\rho}{4} + \rho\mathbf{1}_{z \in i_{\mathcal{O}}}) g_t(z) \mathrm{d}z}{P_{\mathcal{O}}}$, and by subtracting $\nu_t = \mathsf{E}\big(w_t \mid g_t\big)$ one gets $\delta_{t+1,\mathcal{O}} = \frac{\int (\frac{1-\rho}{4} + \rho\mathbf{1}_{z \in i_{\mathcal{O}}})(z - \nu_t) g_t(z) \mathrm{d}z}{P_{\mathcal{O}}}$. But $\int (z - \nu_t) g_t(z) \mathrm{d}z = 0$, hence $\delta_{t+1,\mathcal{O}} = \frac{\int \rho \mathbf{1}_{z \in i_{\mathcal{O}}} (z - \nu_t) g_t(z) \mathrm{d}z}{P_{\mathcal{O}}}$, which proves (A-4).

To compute the mean of $g_{t+1,\mathcal{O}}(x)$, integrate the formula (A-3) over x, and obtain $\nu_{t+1,\mathcal{O}} = \frac{\int (\frac{1-\rho}{4} + \rho\mathbf{1}_{z \in i_{\mathcal{O}}})(z - \delta_{\mathcal{O}}) g_t(z) \mathrm{d}z}{P_{\mathcal{O}}}$. This is similar to the formula for $\delta_{t+1,\mathcal{O}}$, except that ν_t is replaced by $\delta_{\mathcal{O}}$. One gets $\nu_{t+1,\mathcal{O}} = \delta_{t+1,\mathcal{O}} + \nu_t - \delta_{\mathcal{O}}$, which proves the first part

[50] One computes $\tau_{t+1,\mathrm{BMO}} = \tau_{t+1,\mathrm{SMO}} = 1 + \rho\frac{\alpha}{\beta} - \rho^2 \frac{1-\gamma^2}{2\beta^2}$, and $\tau_{t+1,\mathrm{BLO}} = \tau_{t+1,\mathrm{SLO}} = 1 - \rho\frac{\alpha}{\beta} + \rho^2\frac{1-\gamma^2}{2\beta^2}$. The average of $\tau_{t+1,\mathcal{O}}$ is indeed $\bar{\tau}_{t+1} = 1$. Also, let $D(\rho) = \rho\frac{\alpha}{\beta} - \rho^2\frac{1-\gamma^2}{2\beta^2}$ be the absolute deviation of $\tau_{t+1,\mathcal{O}}$ from one. Then, $D(0) = 0$, $D(1) = 0.1022$, and $D(\rho)$ attains a maximum value of 0.2433 at $\rho = \frac{2\alpha\beta}{1-\gamma^2} = 0.5677$. The posterior standard deviation $\tau_{t+1,\mathcal{O}}^{1/2} = \big(1 \pm D(\rho)\big)^{1/2}$ has a maximum deviation from one equal to $1 - (1 - 0.2433)^{1/2} = 0.1301$. Note that $D(\rho)$ is small when ρ is close to zero.

of (A-5).

For the second part of (A-5), note that for any (not necessarily normal) distribution g with mean ν and variance σ_I^2, $\int (x+a)^2 g(x)\mathrm{d}x = \sigma_I^2 + (\nu+a)^2$. Then,

$$\text{(A-11)}\quad \int (x-\nu_t+\delta_{\mathcal{O}})^2 g_{t+1,\mathcal{O}}(x)\mathrm{d}x \;=\; \tau_{t+1,\mathcal{O}}^2 + (\nu_{t+1,\mathcal{O}}-\nu_t+\delta_{\mathcal{O}})^2 \;=\; \tau_{t+1,\mathcal{O}}^2 + \delta_{t+1,\mathcal{O}}^2.$$

One integrates directly $\int (x-\nu_t+\delta_{\mathcal{O}})^2 g_{t+1,\mathcal{O}}(x)\mathrm{d}x$ by replacing $g_{t+1,\mathcal{O}}(x)$ as in (A-3). Using the formula $\int (x-\nu_t+\delta_{\mathcal{O}})^2 \phi(x; z-\delta_{\mathcal{O}}, \hat{\sigma_I})\mathrm{d}x = (z-\nu_t)^2 + \hat{\sigma_I}^2$, one obtains:

$$\text{(A-12)}\qquad \int (x-\nu_t+\delta_{\mathcal{O}})^2 g_{t+1,\mathcal{O}}(x)\mathrm{d}x \;=\; \hat{\sigma_I}^2 + \frac{\int g_t(z)\big(\frac{1-\rho}{4} + \rho \mathbf{1}_{z\in i_{\mathcal{O}}}\big)(z-\nu_t)^2 \mathrm{d}z}{P_{\mathcal{O}}}.$$

Putting together (A-11) and (A-12), one gets the desired formula for $\tau_{t+1,\mathcal{O}}^2$. Equation (A-6) follows directly from (A-4) and (A-5). Finally, proving $\mathsf{E}_{\mathcal{O}}(\delta_{t+1,\mathcal{O}}) = 0$ and $\mathsf{E}_{\mathcal{O}}(V_{t+1,\mathcal{O}}) = 0$ is straightforward, which also implies equation (A-7).

To prove part (c), first use (A-2) to compute $P_{\mathcal{O}} = \frac{1-\rho}{4} + \rho\big(\Phi(h_{\mathcal{O}}) - \Phi(\ell_{\mathcal{O}})\big)$. To prove the formula for $\delta_{t+1,\mathcal{O}}$, make the change of variable $z' = \frac{z-\nu_t}{\tau_t}$ and denote by $i'_{\mathcal{O}} = (\ell_{\mathcal{O}}, h_{\mathcal{O}})$. Then, $\delta_{t+1,\mathcal{O}} = \frac{\rho\tau_t \int_{i'_{\mathcal{O}}} \phi(z')z'\mathrm{d}z}{P_{\mathcal{O}}} = \frac{\rho\tau_t\big(\phi(\ell_{\mathcal{O}})-\phi(h_{\mathcal{O}})\big)}{P_{\mathcal{O}}}$. A similar computation for $V_{t+1,\mathcal{O}}$ finishes the proof of (A-8). Finally, $\bar{\nu}_{t+1} = f(\nu_t) = \nu_t - \sum_{\mathcal{O}} P_{\mathcal{O}}\delta_{\mathcal{O}} = \nu_t - \rho\sum_{\mathcal{O}}\big(\Phi(h_{\mathcal{O}}) - \Phi(\ell_{\mathcal{O}})\big)\delta_{\mathcal{O}}$. If one differentiates the endpoints of $i'_{\mathcal{O}}$ with respect to ν_t, one gets $-\frac{1}{\tau_t}$ in all cases, hence $f'(\nu_t) = 1 - \rho\sum_{\mathcal{O}}\big(\phi(h_{\mathcal{O}}) - \phi(\ell_{\mathcal{O}})\big)\big(-\frac{1}{\tau_t}\big)\delta_{\mathcal{O}}$. Using $\delta_{\mathcal{O}} \in \Big\{\frac{\rho}{\beta}, \gamma\frac{\rho}{\beta}, -\gamma\frac{\rho}{\beta}, -\frac{\rho}{\beta}\Big\}$, a straightforward calculation proves (A-9).

To prove part (d), substitute $\nu_t = 0$ and $\tau_t = 1$ in the formulas above. I only prove the results for $\mathcal{O} = \text{BMO}$ and BLO, the proof for the other order types being symmetric. The probability of a BMO is $P_{\text{BMO}} = \frac{1-\rho}{4} + \rho\int_\alpha^\infty \phi(z)\mathrm{d}z = \frac{1-\rho}{4} + \rho\frac{1}{4} = \frac{1}{4}$. The probability of a BLO is $P_{\text{BLO}} = \frac{1-\rho}{4} + \rho\int_0^\alpha \phi(z)\mathrm{d}z = \frac{1-\rho}{4} + \rho\frac{1}{4} = \frac{1}{4}$.

The normalized price impact of a BMO is $\delta_{t+1,\text{BMO}} = \frac{\rho\int_\alpha^\infty \phi(z)z\mathrm{d}z}{P_{\text{BMO}}} = \frac{\rho\phi(\alpha)}{1/4} = \frac{\rho}{\beta} = \delta_{\text{BMO}}$. The normalized price impact of a BLO is $\delta_{t+1,\text{BLO}} = \frac{\rho\int_0^\alpha \phi(z)z\mathrm{d}z}{P_{\text{BLO}}} = \frac{\rho(\phi(0)-\phi(\alpha))}{1/4} = \frac{\phi(0)-\phi(\alpha)}{\phi(\alpha)}\frac{\rho\phi(\alpha)}{1/4} = \gamma\frac{\rho}{\beta} = \delta_{\text{BLO}}$. By symmetry, it follows that $\delta_{t+1,\mathcal{O}} = \delta_{\mathcal{O}}$ for all orders $\mathcal{O} \in \{\text{BMO}, \text{BLO}, \text{SLO}, \text{SMO}\}$.

I now compute $\nu_{t+1,\mathcal{O}} = \nu_t - \delta_{\mathcal{O}} + \delta_{t+1,\mathcal{O}} = \nu_t = 0$. Also, $\bar{\tau}_{t+1}^2 = \mathsf{E}_{\mathcal{O}}\big(\tau_{t+1,\mathcal{O}}^2\big) = \mathsf{E}_{\mathcal{O}}\big(\tau_t^2 + \hat{\sigma_I}^2 - \delta_{\mathcal{O}}^2\big)$. But $\mathsf{E}_{\mathcal{O}}(\delta_{\mathcal{O}}^2) = \frac{1}{4}\Big((\frac{\rho}{\beta})^2 + (\gamma\frac{\rho}{\beta})^2 + (-\gamma\frac{\rho}{\beta})^2 + (-\frac{\rho}{\beta})^2\Big) = \hat{\sigma_I}^2$, hence

$\bar{\tau}_{t+1}^2 = \tau_t^2 + \hat{\sigma_I}^2 - \hat{\sigma_I}^2 = \tau_t^2$, from which $\bar{\tau}_{t+1} = \tau_t = 1$. Thus, the posterior mean is equal to zero irrespective of the order $\mathcal{O}$ at t, while the posterior variance is equal to one on average. This means that the normalized density $\mathcal{N}(0,1)$ corresponds to a stationary equilibrium. □

In the next two lemmas, I describe the continuation payoff from submitting a BLO for either a patient speculator (Lemma A2), or for an uninformed patient natural buyer (Lemma A3), assuming that all investors follow their equilibrium strategies.

To state the next result, let the "execution probability function" $J(\rho, w)$ be as in Definition 1. Numerically, I verify that J is constant and equal to one (see Result 1), but for the next lemma, no particular expression for J is necessary.

Lemma A2. *In the context of Theorem 1, consider an informed trader who submits a* BLO *at t after observing the signal $w_t = \frac{v_t - \mu_t}{V}$. Then, if subsequently all traders follow the equilibrium strategies, the continuation payoff of the informed trader is:*

$$\text{(A-13)} \qquad \mathcal{U}^I_{\text{BLO}} \;=\; \frac{S}{2}\, J(\rho, w_t) \;+\; V\, I(\rho, w_t).$$

Proof. I simplify notation and assume that the initial BLO is submitted at $t = 0$. Denote by $\mathcal{Q}$ the set of all execution sequences $Q = \big(\mathcal{O}_0 = \text{BLO}, \mathcal{O}_1, \ldots, \mathcal{O}_{T-1}, \mathcal{O}_T = \text{SMO}\big)$ for the initial BLO. Let $\mathcal{J}_t$ be the information set of the informed trader just before trading at t, which consists of the signal w_0 observed at $t = 0$, and the orders $\mathcal{O}_0, \ldots, \mathcal{O}_{t-1}$. Let E_t be the expectation operator conditional on $\mathcal{J}_t$. At the execution time T, the bid price is $\mu_T - S/2$, therefore:

$$\text{(A-14)} \qquad \begin{aligned} \mathcal{U}^I_{\text{BLO}} \;&=\; \sum_{Q \in \mathcal{Q}} \mathsf{E}_0\Big(v_T - (\mu_T - \tfrac{s}{2}) \mid Q\Big)\, P_0(Q) \\ &=\; \frac{S}{2} \sum_{Q \in \mathcal{Q}} P_0(Q) + \sum_{Q \in \mathcal{Q}} e(Q)\, P_0(Q), \end{aligned}$$

where $e(Q) = \mathsf{E}_0\big(v_T - \mu_T \mid Q\big)$.

For $t = 1, \ldots, T+1$, let P_t be the probability of observing the order $\mathcal{O}_t$ at t conditional on $\mathcal{J}_t$, g_t the density of w_t before trading at t, and $\nu_t = \mathsf{E}_t(w_t)$ the mean of g_t. I show that the sequence of probabilities $(P_1, \ldots, P_T)$, densities $(g_1, \ldots, g_{T+1})$, and means $(\nu_1, \ldots, \nu_{T+1})$ is indeed associated to the execution sequence Q, in the sense of Definition 1. From equations (A-2) and (A-3) in Lemma A1, it follows that $P_t = \pi_{g_t, \mathcal{O}_t}$ and $g_{t+1} = f_{g_t, \mathcal{O}_t}$, where π and f are given by equation (5) in Definition 1.

Next, I show that $P_0(Q)$ coincides with $P(Q) = \prod_{t=1}^{T} P_t$ from Definition 1. Indeed, $P_0(Q) = P\big(\mathcal{O}_1, \ldots, \mathcal{O}_T \mid \mathcal{O}_0\big) = \prod_{t=1}^{T} P\big(\mathcal{O}_t \mid \mathcal{O}_0, \ldots, \mathcal{O}_{t-1}\big) = \prod_{t=1}^{T} P_t = P(Q)$. Since by definition $J(\rho, w_0) = \sum_{Q \in \mathcal{Q}} P(Q)$, using (A-14) one obtains the first half of (A-13).

It remains to prove that $\sum_{Q \in \mathcal{Q}} e(Q) P(Q) = V I(\rho, w_0)$. First, I show that $g_1 = \mathcal{N}\big(w_0 - \gamma \frac{\rho}{\beta}, \frac{1+\gamma^2}{2\beta^2}\big)$, as specified in the definition of I. To see this, note that $\delta_{\mathcal{O}_0} = \delta_{\mathrm{BLO}} = \gamma \frac{\rho}{\beta}$. Then, $w_1 = w_0 + \frac{v_1 - v_0}{V} - \frac{\mu_1 - \mu_0}{V} = w_0 + \frac{v_1 - v_0}{V} - \gamma \frac{\rho}{\beta}$. Because $v_1 - v_0 \sim \mathcal{N}(0, \sigma_I^2)$, one has $\mathsf{Var}\big(\frac{v_1 - v_0}{V}\big) = \frac{\sigma_I^2}{V^2} = \rho^2 \frac{1+\gamma^2}{2\beta^2}$, where the last equality follows from (13). Since g_1 is the density of w_1, one obtains indeed $g_1 = \mathcal{N}\big(w_0 - \gamma \frac{\rho}{\beta}, \rho^2 \frac{1+\gamma^2}{2\beta^2}\big)$.

Finally, I show that $e(Q) = V\, \nu(Q)$, where $\nu(Q) = \nu_{T+1} - \frac{\rho}{\beta}$. The executing order at T is a SMO, therefore (A-6) implies that $\mathsf{E}_T\big(w_T \mid \mathrm{SMO}_T\big) = \nu_{T+1} + \delta_{\mathrm{SMO}} = \nu_{T+1} - \frac{\rho}{\beta}$. Thus, $e(Q) = V\ \mathsf{E}_0\, \mathsf{E}_T\big(w_T \mid \mathrm{SMO}_T\big) = V\ \big(\nu_{T+1} - \frac{\rho}{\beta}\big)$.

For future reference, note that according to (A-6) one has the following decomposition:

$$\text{(A-15)} \qquad \nu_{T+1} - \frac{\rho}{\beta} \;=\; \nu_{T+1} + \delta_{\mathrm{SMO}_T} \;=\; \nu_T + \delta_{T+1, \mathrm{SMO}_T},$$

where, as shown in (A-4), $\delta_{T+1, \mathrm{SMO}_T}$ is the normalized adverse selection from SMO_T. □

Lemma A3. *In the context of Theorem 1, consider a patient uninformed trader with private valuation $\bar{u}$, who submits a* BLO *at t. Then, if subsequently all traders follow the equilibrium strategies, the continuation payoff of the uninformed trader is:*

$$\text{(A-16)} \qquad \mathcal{U}^{U}_{\mathrm{BLO}} \;=\; \bar{u} + \frac{S}{2} - \Delta.$$

***Proof*.** See Section 1 in the Internet Appendix. □

Verification of Result 1. See Section 4 in the Internet Appendix. □

Proof of Theorem 1. The proof depends on the conditions in Result 1 being analytically true. Thus, for all $\rho \in (0,1)$, it is assumed that:

$$
\begin{aligned}
&I(\rho,w),\ w - I(\rho,w),\ \text{ and } I(\rho,w) - I(\rho,-w) \text{ are strictly increasing in } w,\\
&\max\Big(\frac{\rho(1+\gamma)}{\beta},\ -2I(\rho,0) - 2\frac{\rho\gamma}{\beta}\Big) \ <\ \alpha - I(\rho,\alpha),\\
&I(\rho,w,j) \text{ decreases in } j \text{ for all } w > 0, \text{ and}\\
&J(\rho,w,j) = 1 \text{ for all } w \text{ and } j \geq 1,
\end{aligned}
\tag{A-17}
$$

where I and J are as in Definition 1.

I next define a Markov Perfect Equilibrium (MPE) of the game, by specifying a strategy profile and a belief system that are compatible with each other. In addition, I specify a set of state variables that summarizes the payoff-relevant information contained in each history of the game. As public state variables, I choose: the public mean (μ_t) and the public volatility (σ_t), the ask price (a_t) and the bid price (b_t), as well as the bid and ask queues.[51] As private state variable, I choose the fundamental value (v_t), which is observed by the informed trader at the time of her arrival (t).

Because I want the game in stationary equilibrium, I choose $\mathcal{N}(0, V^2)$ as the initial public density (before trading at $t = 0$). Moreover, the ask price is $S/2$, the bid price is $-S/2$, with S as in (7), while the initial limit order book has countably many limit orders on each side (see the middle plot in Figure 2).

To define the strategy profile $\mathcal{S}$, I first describe the action of a new trader who arrives at t. Then, I describe the reaction of the other traders remaining in the limit order book to the new arrival at t. Finally, in Section 1 in the Internet Appendix, I describe the reaction of the existing traders to any out-of-equilibrium deviation that might occur from either the new trader or an existing trader. Recall that impatient traders are assumed to automatically submit market orders. I therefore describe only the strategies of patient traders, who can be informed (with private valuation 0), uninformed buyers (with private valuation $\bar{u}$), or uninformed sellers (with private valuation $-\bar{u}$). The strategy profile $\mathcal{S}$ is then given by the following set of rules:

(a) The uninformed buyer arriving at t submits a BLO at the price $(\mu_t + \gamma\Delta) - S/2$.

[51] Because in the model traders can submit orders only for one unit, the limit prices for orders other than the first ones in the bid and ask queues are not relevant.

(b) The uninformed buyer arriving at t submits a SLO at the price $(\mu_t - \gamma\Delta) + S/2$.

(c) The informed trader who observes an asset value v_t when she arrives at t submits an order $\mathcal{O} \in \{\text{BMO}, \text{BLO}, \text{SLO}, \text{SMO}\}$ whenever her signal $\frac{v_t - \mu_t}{V}$ lies, respectively, in the interval $i_{\mathcal{O}} \in \{(\alpha, \infty), (0, \alpha), (-\alpha, 0), (-\infty, -\alpha)\}$.

(d) After the initial order submission, all traders behave as described in (e) and (f).

(e) If a BMO is submitted at t, then an instant later the public mean is updated to $\mu_t + \Delta$, the ask price to $\mu_t + \Delta + S/2$, and the bid price to $\mu_t + \Delta - S/2$, and all other limit traders shift their orders by Δ such that the relative ranks in the ask and bid queues are preserved. After that, no other trader moves until $t + 1$.

The reaction to a SMO at t is symmetric to the reaction to a BMO.

(f) If a BLO is submitted at t, then an instant later the public mean is updated to $\mu_t + \gamma\Delta$, the ask price to $\mu_t + \gamma\Delta + S/2$, and the bid price to $\mu_t + \gamma\Delta - S/2$, and all other limit traders shift their orders by $\gamma\Delta$ such that the relative ranks in the ask and bid queues are preserved. After that, no other trader moves until $t + 1$.

The reaction to a SLO at t is symmetric to the reaction to a BLO.

For brevity, I leave the description of out-of-equilibrium moves $\mathcal{S}$(g) and $\mathcal{S}$(h) to Section 1 in the Internet Appendix.

The belief system is described by the following rules: At $t = 0$, the uninformed investors perceive the asset value distributed according to $\mathcal{N}(0, V^2)$. Subsequently, the uninformed investors' belief about the asset value (the public density) is updated using the approximate Bayes' rule described in Section II. At the time of arrival to the market, the informed investor observes the asset value and can compute the average payoff of a limit order based on updating her belief according to the exact Bayes' rule. After the arrival, however, the informed trader cannot update her belief, and becomes essentially uninformed. In the limit order book at $t = 0$, all traders are uninformed with probability one. At $t \geq 0$, each new trader is believed to be informed with probability ρ by the other traders. Subsequently, traders' beliefs about the other investors' types are updated according to the Bayes' rule.

Because the strategy profile $\mathcal{S}$ defined above depends only on the current value of the state variables, the strategies are indeed Markov. I now show that the strategy of each type of investor is a best response to the other investors' strategies.

Uninformed Traders

I consider a patient natural buyer (with private valuation $\bar{u}$ and zero waiting costs). I need to show that the strategy specified by $\mathcal{S}$ is optimal for this trader. Because the proof is straightforward but tedious, I present only the intuition behind the results, and leave the complete proof of this statement to Section 1 in the Internet Appendix.

Intuitively, it is clear that the patient natural buyer chooses a buy order, since with a sell order he would lose the private valuation $\bar{u}$. Hence, the main choice is between a BMO, a BLO, and NO (no order). Recall that a simplifying assumption in Section II is that the uninformed trader starts with a prior belief at t such that after submitting his order his posterior belief coincides with the public density. In the proof of Lemma A3 I compute that his prior belief is $v_t \sim \mathcal{N}\big(\mu_t + \gamma\Delta, V^2 - \sigma_I^2\big)$ (see equation (IA.2) in Section 1 in the Internet Appendix).

Then, Lemma A3 shows that the trader's continuation payoff from submitting a BLO is $\mathcal{U}^U_{\text{BLO}} = S/2 - \Delta + \bar{u}$. If instead he submits a BMO, he gets $\mathcal{U}^U_{\text{BMO}} = \mathsf{E}_t(v_t) - a_t + \bar{u} = (\mu_t + \gamma\Delta) - (\mu_t + S/2) + \bar{u} = \bar{u} + \gamma\Delta - S/2$. Finally, if he submits no order, he gets by convention zero.

First, I note that BMO is preferred to NO, since $\bar{u} \geq S/2$. To compare BMO with BLO, note that condition (A-17) implies $\alpha - I(\rho, \alpha) > \frac{\rho}{\beta}(1+\gamma)$, which if one multiplies by $V = \beta\rho^{-1}\Delta$ implies:

$$S \;>\; \Delta(1+\gamma), \tag{A-18}$$

meaning that the relative benefit of a limit order (the bid-ask spread S) is larger than the relative cost of a market order (the adverse selection Δ coming from the execution with a market order, plus the price impact of a limit order $\gamma\Delta$). One obtains $\mathcal{U}^U_{\text{BLO}} > \mathcal{U}^U_{\text{BMO}}$, therefore the patient natural buyer optimally submits a BLO. The rest of the proof is in Section 1 in the Internet Appendix.

Informed Traders

I prove that the strategy of an informed trader is as specified in $\mathcal{S}$(c), $\mathcal{S}$(e), $\mathcal{S}$(f). Consider a (patient) informed trader who arrives at t and observes the asset value v_t, or equivalently the signal $w_t = \frac{v_t - \mu_t}{V}$. The informed trader has the option to submit either (i) BMO, (ii) SMO, (iii) NO (no order), (iv) BLO at $b^* = (\mu_t + \gamma\Delta) - S/2$, and later follow $\mathcal{S}$, (v) SLO at $a^* = (\mu - \gamma\Delta) + S/2$, and later follow $\mathcal{S}$. I show that the informed trader submits $\mathcal{O} \in \{\text{SMO}, \text{SLO}, \text{BLO}, \text{BMO}\}$ whenever w_t lies, respectively, in the interval $\{(-\infty, -\alpha), (-\alpha, 0), (0, \alpha), (\alpha, \infty)\}$; for this, I show that option (iii) is eliminated by a penalty for not trading that satisfies $\omega \geq \gamma\Delta$. Then, if $w_t > 0$, I show that option (iv) is less profitable if the BLO is submitted at a different price than b^*; symmetrically, if $w_t < 0$, I show that option (v) is less profitable if the SLO is submitted at a different price than a^*. After submitting (iv) or (v), the informed trader behaves in the same way as the uninformed buyer.

Let $\mathcal{U}^I_{\mathcal{O}}$ be the continuation payoff from submitting $\mathcal{O}$ and later following $\mathcal{S}$. As in the case of the uninformed buyer, I assume that the current limit order book has the ask price $a_t = \mu_t + S/2$, and the bid price $b_t = \mu_t - S/2$. Let $\hat{\mathcal{U}}^I_{\mathcal{O}} = \frac{\mathcal{U}^I_{\mathcal{O}}}{V}$ be the normalized payoff from $\mathcal{O}$; $\hat{S} = \frac{S}{V}$ the normalized spread parameter; and $\hat{\omega} = \frac{\omega}{V}$ the normalized commitment parameter, which is a penalty for non-trading. From Lemma A2, $\hat{\mathcal{U}}^I_{\text{BLO}} = \frac{\hat{S}}{2} J(\rho, w_t) + I(\rho, w_t)$. But, by condition (A-17), $J(\rho, w_t) = 1$, hence $\hat{\mathcal{U}}^I_{\text{BLO}} = \frac{\hat{S}}{2} + I(\rho, w_t)$. Putting together all the formulas, one obtains:

$$\text{(A-19)} \qquad \begin{aligned} \hat{\mathcal{U}}^I_{\text{BMO}} &= w_t - \frac{\hat{S}}{2}, \qquad \hat{\mathcal{U}}^I_{\text{SMO}} = -\frac{\hat{S}}{2} - w_t, \qquad \hat{\mathcal{U}}^I_{\text{NO}} = -\hat{\omega}, \\ \hat{\mathcal{U}}^I_{\text{BLO}} &= \frac{\hat{S}}{2} + I(\rho, w_t), \qquad \hat{\mathcal{U}}^I_{\text{SLO}} = \frac{\hat{S}}{2} + I(\rho, -w_t). \end{aligned}$$

Denote by $A(w) = w - I(\rho, w)$, $B(w) = w - I(\rho, -w)$, $D(w) = I(\rho, w) - I(\rho, -w)$; and note that $B(w) = A(w) + D(w)$. With these notations, one gets $\hat{\mathcal{U}}^I_{\text{BMO}} - \hat{\mathcal{U}}^I_{\text{BLO}} = A(w_t) - \hat{S}$, $\hat{\mathcal{U}}^I_{\text{BMO}} - \hat{\mathcal{U}}^I_{\text{SLO}} = B(w_t) - \hat{S}$, $\hat{\mathcal{U}}^I_{\text{BLO}} - \hat{\mathcal{U}}^I_{\text{SLO}} = D(w_t)$, $\hat{\mathcal{U}}^I_{\text{BLO}} - \hat{\mathcal{U}}^I_{\text{SMO}} = \hat{S} - B(-w_t)$, $\hat{\mathcal{U}}^I_{\text{SLO}} - \hat{\mathcal{U}}^I_{\text{SMO}} = \hat{S} - A(-w_t)$. From (A-17), it follows that that A, D, and $B = A + D$ are strictly increasing in w, therefore all the payoff differences above are strictly increasing in w_t. Note that by the definition of S, one has $A(\alpha) = \alpha - I(\rho, \alpha) = \hat{S}$, therefore

BMO is preferred to BLO if and only if $w_t > \alpha$. Similarly, SMO is preferred to SLO if and only if $w_t < -\alpha$. Also, $D(0) = 0$, therefore BLO is preferred to SLO if and only if $w_t > 0$. Because all the payoff differences are strictly increasing in w_t, a straightforward analysis shows that indeed the informed trader prefers $\mathcal{O} \in \{\text{BMO}, \text{BLO}, \text{SLO}, \text{SMO}\}$ whenever w_t lies, respectively, in the interval $i_{\mathcal{O}} \in \{(\alpha, \infty), (0, \alpha), (-\alpha, 0), (-\infty, -\alpha)\}$.

Next, I make sure that NO ("No Order") is never optimal. For that, I use equation (A-19) to compute the minimum payoff for each type of order. According to condition (A-17), I is strictly increasing in w, therefore $\hat{\mathcal{U}}^I_{\mathcal{O}}$ is increasing in w_t for BMO and BLO, and decreasing in w_t for SMO and SLO. Thus, it is sufficient to verify that $\hat{\mathcal{U}}^I_{\text{BLO}} \geq -\hat{\omega}$ when $w_t = 0$. Since by assumption $\omega \geq \gamma\Delta$, the formula $V = \beta\rho^{-1}\Delta$ implies $\hat{\omega} \geq \frac{\rho\gamma}{\beta}$. Hence, it is sufficient to verify that $\frac{\hat{S}}{2} + I(\rho, 0) \geq -\frac{\rho\gamma}{\beta}$. But this follows from condition (A-17), which implies that $\alpha - I(\rho, \alpha) > -2I(\rho, 0) - 2\frac{\rho\gamma}{\beta}$.

I now show that the continuation payoff for the informed trader from submitting his BLO at the equilibrium price b^* is higher than the payoff obtained by choosing BLO at either $b > b^*$ or $b < b^*$. I first rule out BLO at $b > b^*$. Based on the out-of-equilibrium reaction $\mathcal{S}(\text{h})$ described in Section 1 in the Internet Appendix, overshooting a bid is interpreted as coming with probability one from an informed trader with a positive signal. This leads to a positive shift in the public mean μ_t and therefore to a negative shift in the informed trader's signal $w_t = \frac{v_t - \mu_t}{V}$. Condition (A-17) then implies that $I(\rho, w_t)$ strictly decreases, and along with it the informed trader's expected payoff.

I also rule out BLO at $b < b^*$. Based on the out-of-equilibrium reaction $\mathcal{S}(\text{h})$ described in Section 1 in the Internet Appendix, this deviation does not bring any new information about the transgressor's type, but prompts another trader in the bid queue to immediately modify his BLO at b^*. The informed trader thus loses his first rank in the bid queue, which according to Lemma A2 generates a normalized continuation payoff of $\frac{\hat{S}}{2} J(\rho, w_t, j) + I(\rho, w_t, j)$, where $j > 1$ is the informed trader's new rank in the bid queue.[52] By condition (A-17), $J(\rho, w_t, j) = 1$ and $I(\rho, w_t, j)$ is decreasing in j, which implies that the informed traders gets a smaller payoff than $\frac{\hat{S}}{2} + I(\rho, w_t, j = 1) = \hat{\mathcal{U}}^I_{\text{BLO}}$. Hence, the informed trader reduces his payoff by deviating from $b = b^*$.

Finally, after the initial order choice the strategy of the informed trader is the same

[52]By condition (A-17), $J = 1$.

as for the uninformed trader, since they now have the same information set. $\square$

Proof of Corollary 1. The corollary follows directly from the description of the equilibrium strategy profile $\mathcal{S}$, and in particular from $\mathcal{S}$(e) and $\mathcal{S}$(f). $\square$

Proof of Corollary 2. This corollary follows directly from the description of the equilibrium strategy profile $\mathcal{S}$, and in particular from $\mathcal{S}$(c). The formula for the expected utility of the informed trader follows from equation (A-19). $\square$

Proof of Corollary 3. As proved in Theorem 1, the public density at t is $\mathcal{N}(\mu_t, V^2)$, which implies that the normalized public density (the density of the signal $w_t = \frac{v_t - \mu_t}{V}$) is standard normal. Then, part (d) of Lemma A1) shows that all orders have probability equal to 1/4. $\square$

Proof of Corollary 4. This corollary follows directly from the description of the equilibrium strategy profile $\mathcal{S}$, and in particular from $\mathcal{S}$(a) and $\mathcal{S}$(b). The formula for the expected utility of the uninformed trader follows from Lemma A3. $\square$

Proof of Proposition 1. From Corollary 1, any order $\mathcal{O} \in \{\mathrm{BMO}, \mathrm{BLO}, \mathrm{SLO}, \mathrm{SMO}\}$ moves the public mean μ_t by $\Delta_{\mathcal{O}} \in \{\Delta, \gamma\Delta, -\gamma\Delta, -\Delta\}$, respectively. Because each type of order occurs with probability $\frac{1}{4}$, and the public mean moves by an element of $\{\Delta, \gamma\Delta, -\gamma\Delta, -\Delta\}$, it is simple to show that the variance of $\mu_{t+1} - \mu_t$ is indeed equal to $\frac{1+\gamma^2}{2}\Delta^2$. $\square$

Proof of Corollary 5. By Corollary 1, if the public mean is μ, at any time the ask price is $\mu + S/2$, and the bid price is $\mu - S/2$. This implies that the bid-ask spread is equal to the parameter $S = \big(\alpha - I(\rho, \alpha)\big)V$ from (7), and is therefore constant. $\square$

Proof of Proposition 2. The proposition follows from the proof of Lemma A2 in the Appendix. $\square$

Proof of Corollary 6. By equation (7), $S = \big(\alpha - I(\rho, \alpha)\big)V = \text{Decay Cost}_{\alpha}$. $\square$

Proof of Proposition 3. Recall that the slippage function I^s follows Definition 1, except that $\nu(Q) = \nu_T$ instead of $\nu(Q) = \nu_{T+1} - \frac{\rho}{\beta}$ for the information function I. This proves the formula $I^s(\rho, w) = \mathsf{E}^{\mathrm{e}}\,\mathsf{E}_T\big(w_T\big)$. The adverse selection function $I - I^s = I^a$

therefore follows Definition 1, except that $\nu(Q) = (\nu_{T+1} - \frac{\rho}{\beta}) - \nu_T$. But equation (A-15) implies $\nu_{T+1} - \frac{\rho}{\beta} = \nu_T + \delta_{T+1,\mathrm{SMO}_T}$. Hence, I^a is defined using $\nu(Q) = \delta_{T+1,\mathrm{SMO}_T}$, which is the price impact of the SMO at execution time T. But this is equal to $\mathsf{E}_{T+1}(w_T) - \mathsf{E}_T(w_T)$, which proves that indeed $I^a(\rho, w) = \mathsf{E}^{\mathrm{e}}\big(\mathsf{E}_{T+1}(w_T) - \mathsf{E}_T(w_T)\big)$. $\square$

Proof of Corollary 7. By equation (20), the slippage component satisfies $S^s = \big(\alpha - I^s(\rho, \alpha)\big)\, V = \text{Slippage Cost}_\alpha$. Also, the adverse selection component satisfies $S^a = \big(I^s(\rho, \alpha) - I(\rho, \alpha)\big)\, V = -I^a(\rho, w)\, V = \text{Adverse Selection Cost}_\alpha$. $\square$

Verification of Result 2. See Section 4 in the Internet Appendix. $\square$

Proof of Proposition 4. With the notation of Lemma A1 in this Appendix, consider a trader that perceives the signal $w_t = \frac{v_t - \mu_t}{V}$ distributed according to the normalized density prior, $\mathcal{N}(\nu_t, \tau_t^2) = \mathcal{N}(0, 1)$. Denote by $\nu_{t+1} = f(\nu_t)$ the average posterior mean. Then, by setting $\nu_t = 0$ and $\tau_t = 1$ in equation (A-9), one obtains the desired formula for the resiliency coefficient K. $\square$

I now define formally the parameters $\tilde{\alpha}$, $\tilde{\beta}$, $\tilde{\gamma}$, $\tilde{I}$, $\tilde{V}$, $\tilde{\Delta}$ and $\tilde{S}$ that are used in Section V. Recall that $\phi(\,\cdot\,)$ is the standard normal density, and $\Phi(\,\cdot\,)$ is its cumulative density.

Definition A1. *Let $\rho \in (0, 1)$, $\theta \in (0, \infty)$, and $w \in \mathbb{R}$. Define the functions $\tilde{I}(\rho, w, \theta)$ and $\tilde{J}(\rho, w, \theta)$ as in Definition 1, except that in the recursive step at t, instead of the numeric parameters α, β, and γ, one uses the functions $\tilde{\alpha}(\rho, \theta_t)$, $\tilde{\beta}(\rho, \theta_t)$, and $\tilde{\gamma}(\rho, \theta_t)$ defined below, with $\theta_0 = \theta$, and:*

(A-20)

$$\theta_{t+1}^2 \;=\; \rho^2\, \frac{1+\gamma^2}{2\beta^2} + \theta_t^2 - 2\rho^2\theta_t^2 \left(\frac{\Big(\phi\big(\frac{\tilde{\alpha}_t}{\theta_t}\big)\Big)^2}{\frac{1-\rho}{4} + \rho\big(1 - \Phi\big(\frac{\tilde{\alpha}}{\theta_t}\big)\big)} + \frac{\Big(\phi\big(\frac{0}{\theta_t}\big) - \phi\big(\frac{\tilde{\alpha}_t}{\theta_t}\big)\Big)^2}{\frac{1-\rho}{4} + \rho\big(\Phi\big(\frac{\tilde{\alpha}}{\theta_t}\big) - \Phi(0)\big)} \right),$$

where $\tilde{\alpha}_t = \tilde{\alpha}(\rho, \theta_t)$.[53] *The functions* $\tilde{\alpha}$, $\tilde{\beta}$, *and* $\tilde{\gamma}$ *are defined by the implicit equations:*

$$\text{(A-21)} \quad \begin{aligned} &\tilde{\alpha} - \tilde{I}(\rho, \tilde{\alpha}, \theta) - 2\, \frac{\rho\, \theta\, \phi\big(\frac{\tilde{\alpha}}{\theta}\big)}{\frac{1-\rho}{4} + \rho\big(1 - \Phi\big(\frac{\tilde{\alpha}}{\theta}\big)\big)} \;=\; \alpha - I(\rho, \alpha) - 2\, \frac{\rho}{\beta}, \\ &\tilde{\beta} \;=\; \frac{\frac{1-\rho}{4} + \rho\big(1 - \Phi\big(\frac{\tilde{\alpha}}{\theta}\big)\big)}{\phi\big(\frac{\tilde{\alpha}}{\theta}\big)}, \qquad \tilde{\gamma} \;=\; \frac{\phi(0) - \phi\big(\frac{\tilde{\alpha}}{\theta}\big)}{\phi\big(\frac{\tilde{\alpha}}{\theta}\big)}\, \frac{\frac{1-\rho}{4} + \rho\big(1 - \Phi\big(\frac{\tilde{\alpha}}{\theta}\big)\big)}{\frac{1-\rho}{4} + \rho\big(\Phi\big(\frac{\tilde{\alpha}}{\theta}\big) - \Phi(0)\big)}, \end{aligned}$$

Also, if $V = V(\rho)$ *as in equation* (7), *define* $\tilde{V} = \tilde{V}(\rho, \theta)$, $\tilde{\Delta} = \tilde{\Delta}(\rho, \theta)$, *and* $\tilde{S} = \tilde{S}(\rho, \theta)$ *by:*

$$\text{(A-22)} \qquad \tilde{V} \;=\; \theta\, V, \qquad \tilde{\Delta} \;=\; \frac{\rho}{\tilde{\beta}}\, \theta\, V, \qquad \tilde{S} \;=\; \big(\tilde{\alpha} - \tilde{I}(\rho, \tilde{\alpha}, \theta)\big)\, V.$$

Finally, the proofs of Propositions 5–7 and Corollary 8, as well as the verification of Results 4–5 are left to the Internet Appendix.

References

Admati, A., and P. Pfleiderer (1988): "A Theory of Intraday Patterns: Volume and Price Variability," *Review of Financial Studies*, 1, 3–40.

Amihud, Y., and H. Mendelson (1986): "Asset Pricing and the Bid-Ask Spread," *Journal of Financial Economics*, 17, 223–249.

Anand, A., S. Chakravarty, and T. Martell (2005): "Empirical Evidence on the Evolution of Liquidity: Choice of Market versus Limit Orders by Informed and Uninformed Traders," *Journal of Financial Markets*, 8, 289–309.

Back, K., and S. Baruch (2007): "Working Orders in Limit Order Markets and Floor Exchanges," *Journal of Finance*, 62, 1589–1621.

Bagehot, W. (1971): "The Only Game in Town," *Financial Analysts Journal*, 27, 12–14 and 22.

Bergin, J., and B. MacLeod (1993): "Continuous Time Repeated Games," *International Economic Review*, 34, 21–37.

Biais, B., P. Hillion, and C. Spatt (1995): "An Empirical Analysis of the Limit Order Book and the Order Flow in the Paris Bourse," *Journal of Finance*, 50, 1655–1689.

Biais, B., J. Hombert, and P.-O. Weill (2013): "Equilibrium Pricing and Trading Volume under Preference Uncertainty," *Review of Economic Studies*, forthcoming.

Bloomfield, R., M. O'Hara, and G. Saar (2005): "The 'Make or Take' Decision in an Electronic Market: Evidence on the Evolution of Liquidity," *Journal of Financial Economics*, 75, 165–199.

Brennan, M., and A. Subrahmanyam (1996): "Market Microstructure and Asset Pricing: On the Compensation for Illiquidity in Stock Returns," *Journal of Financial Economics*, 41, 441–464.

[53]The only case when a tilde is not added over a parameter is when the term $\rho\sqrt{\frac{1+\gamma^2}{2\beta^2}}$ occurs. This term is equal to $\frac{\sigma_I}{V}$, the normalized inter-arrival volatility, where σ_I is the volatility of the change in fundamental value between order arrivals.

Brolley, M., and K. Malinova (2017): "Informed Trading in a Low-Latency Limit Order Market," Working Paper.

Chakravarty, S., and C. Holden (1995): "An Integrated Model of Market and Limit Orders," *Journal of Financial Intermediation*, 4, 213–241.

Cohen, K., S. Maier, R. Schwartz, and D. Whitcomb (1981): "Transaction Costs, Order Placement Strategy, and Existence of the Bid-Ask Spread," *Journal of Political Economy*, 89, 287–305.

Collin-Dufresne, P., and V. Fos (2013), "Do Prices Reveal the Presence of Informed Trading," Working Paper.

Easley, D., S. Hvidkjaer, and M. O'Hara (2002): "Is Information Risk a Determinant of Asset Returns?" *Journal of Finance*, 57, 2185–2221.

Foucault, T. (1999): "Order Flow Composition and Trading Costs in a Dynamic Limit Order Market," *Journal of Financial Markets*, 2, 99–134.

Foucault, T., O. Kadan, and E. Kandel (2005): "Limit Order Book as a Market for Liquidity," *Review of Financial Studies*, 18, 1171–1217.

Fudenberg, D., and J. Tirole (1991): *Game Theory*, MIT Press.

Glosten, L. (1994): "Is the Electronic Open Limit Order Book Inevitable?," *Journal of Finance*, 49, 1127–1161.

Glosten, L., and P. Milgrom (1985): "Bid, Ask and Transaction Prices in a Specialist Market with Heterogeneously Informed Traders," *Journal of Financial Economics*, 14, 71–100.

Goettler, R., C. Parlour, and U. Rajan (2005): "Equilibrium in a Dynamic Limit Order Market," *Journal of Finance*, 60, 2149–2192.

Goettler, R., C. Parlour, and U. Rajan (2009): "Informed Traders and Limit Order Markets," *Journal of Financial Economics*, 93, 67–87.

Griffiths, M., B. Smith, A. Turnbull, and R. White (2000): "The Costs and Determinants of Order Aggressiveness," *Journal of Financial Economics*, 56, 65–88.

Harris, L. (1998): "Optimal Dynamic Order Submission Strategies in Some Stylized Trading Problems," *Financial Markets, Institutions and Instruments*, Vol. 7, No. 2.

Harris, L., and J. Hasbrouck (1996): "Market versus Limit Orders: the Superdot Evidence on Order Submission Strategy," *Journal of and Financial and Quantitative Analysis*, 31, 213–231.

Hasbrouck (1993): "Assessing the Quality of a Security Market: A New Approach to Transaction-Cost Measurement," *Review of Financial Studies*, 6, 191–212.

Hautsch, N., and R. Huang (2012): "The Market Impact of a Limit Order," *Journal of Economic Dynamics and Control*, 36, 501-522.

Hollifield, B., R. Miller, and P. Sandås(2004): "Empirical Analysis of Limit Order Markets," *Review of Economic Studies*, 71, 1027–1063.

Hou, K., and T. Moskowitz (2005): "Market Frictions, Price Delay, and the Cross-Section of Expected Returns," *Review of Financial Studies*, 18, 981–1020.

Huang, R., and H. Stoll (1997): "The Components of the Bid-Ask Spread: A General Approach," *Review of Financial Studies*, 10, 995–1034.

Jain, P. (2005): "Financial Market Design and the Equity Premium: Electronic versus Floor Trading," *Journal of Finance*, 60, 2955–2985.

Kaniel, R., and H. Liu (2006): "So What Orders Do Informed Traders Use?" *Journal of Business*, 79, 1867–1913.

Kyle, Albert (1985): "Continuous Auctions and Insider Trading," *Econometrica*, 53, 1315–1335.

Large, J. (2009): "A Market-Clearing Role for Inefficiency on a Limit Order Book," *Journal of Financial Economics*, 56, 65–88.

Latza, T., and R. Payne (2011): "Forecasting Returns and Trading Equities Intra-Day Using Limit Order and Market Order Flows," Working Paper.

Menkhoff, L., C. Osler, and M. Schmeling (2010): "Limit-Order Submissions Strategies under Asymmetric Information," *Journal of Banking and Finance*, 34, 2665–2677.

O'Hara, M. (1995): *Market Microstructure Theory*, Blackwell.

Pagnotta, E. (2010): "Information and Liquidity Trading at Optimal Frequencies," Working Paper.

Parlour, C. (1998): "Price Dynamics in Limit Order Markets," *Review of Financial Studies*, 11, 789–816.

Parlour, C., and D. Seppi (2008): "Limit Order Markets: A Survey," Chapter 3 in *Handbook of Financial Intermediation and Banking*, edited by A.W.A. Boot and A.V. Thakor.

Pástor, Ľ., and R. Stambaugh (2003): "Liquidity Risk and Expected Stock Returns?" *Journal of Political Economy*, 111, 642–685.

Roşu, I. (2009): "A Dynamic Model of the Limit Order Book," *Review of Financial Studies,* 22, 4601–4641.

Roşu, I. (2019): "Dynamic Adverse Selection and Liquidity," Working Paper, February.

Sandås, P. (2001): "Adverse Selection and Competitive Market Making: Empirical Evidence from a Limit Order Market," *Review of Financial Studies*, 14, 705–734.