Information droughts on the limit order book

Corey Garriott [1]

Department of Economics
University of California, Los Angeles

January 5, 2011

Abstract

This paper studies a financial exchange in which liquidity may "dry up" endogenously. The drought is generated on a model of a limit order book under conditions of asymmetric information and an evolving fundamental. The book exhibits an equilibrium spread that is in steady state when it balances the rate that traders acquire information with the rate at which their stock of information decays. If the decay rate is too high, beliefs mean-revert almost everywhere to an uninformative prior, raising asymmetric information risk and motivating the departure of liquidity. A drought arises endogenously if traders lose confidence in their information stock. The drought has an element of stability because it reduces trading volume. Traders cannot learn enough from the reduced volume to motivate reentry. The result is interpreted as a limitation on information aggregation that is institutional.

1 Introduction

A financial exchange is charged with at least two functions: the provision of liquidity and the aggregation of information. The object of this paper is to link them by studying a learning problem faced by traders in a financial exchange. The paper gives conditions under which traders cannot sustain beliefs about an asset's value, and it shows these conditions may arise endogenously. It interprets the result as a limitation on information aggregation that is institutional. The trading institution, here the limit order book, can elicit information about the asset only by providing liquidity, *i.e.* the opportunity to trade, and vice versa. When it fails to do either the exchange enters a liquidity drought that is also an information drought.

Liquidity is defined here as the ability to trade without delay and without paying transactions fees. A liquidity ideal is the Arrow-Debreu full-information price, a single price at which any quantity may be traded immediately and for no fee. This perfectly liquid benchmark price would require, among other things, the coincidence of agents who have a contemporaneous desire to trade. Since traders arrive separately and at various times, exchanges have created institutions to match them intertemporally. The primary matching institution on financial exchanges is called the limit order book.[2] Order books necessarily deviate from the liquidity ideal because traders who supply limit orders must be compensated for supplying immediacy and for risking trading with the better informed (Copeland and Galai 1983; Foucault, Kadan and Kandel 2005). This paper is interested in modeling the order book institution because one form of its liquidity failure, the spread, has dynamic consequences for information aggregation.

In addition to providing liquidity, exchange traders are supposed to perform a statistical function by eliciting information held privately and expressing it as the price. Intuitively, traders push the market price toward such a statistic through a sequential *tâtonnement* (Hellwig 1982). A trader who believes an asset price is wrong can profit by demanding or supplying the asset, driving the price toward the trader's beliefs. A sequence of well-capitalized informed traders will draw the asset price closer and closer to what they know.

To link learning with liquidity, this paper studies a theoretical dynamic limit order book with free entry under asymmetric information. Its contribution is to relate information aggregation to the limit order book institution. The model is simplied enough that the paper can prove some results about the learning regime. It finds the level of liquidity supply interacts dynamically with the ability of traders to learn. Moreover, illiquidity and a lack of information can feed back on one another. Exchange failures of various degrees of severity can arise endogenously and persist, and they can be a generic feature of financial trade.

---

[2]On an order book, traders may announce offers to buy or to sell some fixed, prespecified quantity at a fixed, prespecified strike price. Such a quote is called a *limit order* because it executes a buy at no more than some limiting price or a sell at no less than some limiting price. So long as the limit offer remains in force, other traders may fill it by posting a *market order*, which is an order to buy or to sell some quantity at the best prices available. The book can match traders because limit orders remain in force for some time. Order books are a primary financial institution. The NYSE Arca book handles one-sixth of the trading volume in the United States.

The model concentrates on a pricing game played by liquidity suppliers competing for book position. There is one asset whose net present value follows a stochastic process in continuous time. A large number of risk-neutral agents called "book agents" may enter the limit order book by quoting limit orders, and while active they may continuously change quoted prices, and they may exit by canceling orders. Liquidity suppliers are not interested in holding but in making short-term profits through liquidity supply (see Vives 1995). Market orders[3] arrive at a Poisson rate for exogenous reasons that are not modeled. When a market order arrives, its size (the amount to buy or to sell) is informative about the current true value of the asset. For example, if the prevailing limit order prices are lower than the true value, the market orders that arrive will tend to be buy orders.

Book agents know the sizes of past market orders are related to the true value, so they learn about the value of the asset by updating from the order history. In addition to learning from realized orders, book agents also gather information from periods of no observed trade, because periods of market order silence send a signal there has been no major change in the asset value (Easley O'Hara 1992). Last, the book can also lose information. Due to the changing asset value old signals diminish in informativeness, which acts as a mean-reverting drag on the posterior.

The paper contains two sets of results. First it studies what happens as the variance of the asset value's stochastic process increases. If the variance is too high the book cannot sustain any beliefs about the asset because the book loses information faster than it can gain it, as in Hellwig (1982). Liquidity then "dries up." The second result is an endogenous liquidity drought, given in section 4. There the model acquires a second dimension of uncertainty in addition to the asset's unknown value. The asset value's autocovariance will also follow a stochastic process, complicating signal inference and creating the possibility liquidity can dry up randomly.

The key driver of the results is the way two model entities feed on one another: the spread and the fill rate. The spread is the distance between the best buying and best selling prices prevailing. When it is relatively wide, traders looking to trade immediately find the best buying price too high and the best selling price too low to motivate a market order, so they do not transact. A wide spread therefore decreases the rate of arrival of market orders, which is called the fill rate. A low fill rate is bad for learning because the financial exchange learns from incoming orders. Anything that stems the order flow also stems the flow of new information. The feedback effect is that the exchange only provides liquidity when it is informed. Hence if the book cannot learn much about the asset, its spread dilates, drying up the flow of information, which keeps the spread wide.

---

[3]A market order is a direction to buy or sell some quantity of the asset at the best available prices. See the first footnote.

1.1 Comparison with the literature

The model falls in the tradition of pure limit order books with sequential trade, inspired by Glosten and Milgrom (1985). The presence of informed traders induces a spread to compensate liquidity suppliers for risking trading with the better informed. The closest papers to this work in spirit are Easley and O'Hara (1987, 1992). In their models the size and sequence of the order flow matter to the determination of the spread, block trades dilate the spread, and market silence in between orders contracts the spread. This model encapsulates these effects and shows market silence may dilate the spread as well as contract it.

Limit order books with a full price schedule appear in Glosten (1989) and Glosten (1994), which give monopolistic and perfectly competitive price schedules. Biais, Martimort and Rochet (2000) extend the result to imperfect competition. These books are static; dynamic books are studied in Parlour (1998), Foucault (1999) and Goettler, Parlour, and Rajan (2005), all of which wrestle with the tractability problem of modeling strategic agents competing in a dynamic setting. This paper is indebted to the continuous-time game setup given by Roşu (2009) and Roşu (2006wp), which can isolate rigid equilibria. The dynamic limit order books above allow agents to choose between quoting a limit order or posting a market order. This paper reverts to the older convention of treating agents who post limit orders as a separate class. The methodological "step back" makes room to prove things that would be difficult in a more sophisticated setting.

The model can reproduce some of the behaviors of empirical order books. As in Biais, Hillion and Spatt (1995), the fill rate is high when the spread is tight; order flow volatility dilates the spread; the spread mean-reverts; large contrarian trades indicate informed trading. As in Goldstein and Kavajecz (2004), extreme market movements cause limit order exit. As in Greene and Smart (1999), liquidity improves in the measure of noninformed trading.

Studies of information aggregation in financial settings often assume the traded asset has a static value (or one that settles). The assumption is fruitful for thinking about motivating traders to disclose private information independently. It is used to study heterogenous private information in Foster and Viswanathan (1996) and decentralized private information in Golosov, Lorenzoni and Tsyvinski (2008). Ostrovsky (2010wp) gives general results for "separable" assets. This paper uses a stochastic value because it inquires not into how learning interacts with incentives but into how learning interacts with an institution.

Many studies of information in finance focus on traders with market power who trade strategically. This paper mostly sidesteps dynamic strategic concerns. Market orders take their value by maximizing a one-shot utility function; the convention receives theoretical defense in Back and Baruch (2004). Book agents are more strategic as they play a dynamic price-setting game, but they choose Markov-perfect strategies in equilibrium due to continuous-time undercutting.

Dramatic liquidity events such as droughts have begun to warrant attention. Brunnermeier and Pedersen (2008) give an account of liquidity droughts in the spirit of limits to arbitrage.

Morris and Shin (2004) finds a similar feedback effect setting limits on the amount of money agents can tolerate losing. Persistent liquidity or illiquidity can driven by dynamic complementarities in search externalities, as in Vayanos and Weill (2008). These papers share an emphasis on feedback through real factors such as available capital, regulations, loss limits and search costs. This paper adds to these real explanations one in which beliefs can have feedback effects, such as with the self-fulfilling prophecies of Farmer and Guo (1994).

The model assumes book agents may enter freely, and since the game is set in continuous time they enter instantly. The assumption amounts to what Kyle (1985) calls resiliency. If limit orders entered at a staggered rate, after a market order the spread would dilate beyond the competitive distance until more book agents arrived. DeGryse et al. (2003) and Coppejans, Domowitz, and Madhavan (2003) have studied this phenomenon. Continuous-time entry simplifies the model because there are no intermediate stages of liquidity recovery. Relaxing the assumption of resiliency would strengthen the result because spreads would stay wider longer. Resiliency is not unreasonable due to the first-mover advantage in the race to supply liquidity (Harris 1994).

The model drought bears some resemblance to information cascades (Bikhchandani et al. 1992) because market orders do not execute (and hence reveal information) with positive probability. However, on a book in which there is any trade Avery and Zemsky (1998) emphasize there can be no information cascade. If there is a positive probability of trading, the choice not to trade is informative, hence the posterior distribution over the history of observed actions cannot be independent of the state. Cascades occur only if the adverse selection problem is so great agents cannot agree to disagree (Aumann 1976). Last, the drought is not a herd behavior because the market order flow is modeled in reduced form; there are no agents to herd.

1.2 Order of the paper

Section 2 delineates the model objects, defines the mechanism of trade and defines an equilibrium. Section 3 studies a benchmark result and gives informational conditions for a liquid market: the book must be able to learn faster than it loses information. Section 4 demonstrates the possibility of an endogenous liquidity drought. Section 5 draws conclusions.

2 The model

For ease of notation I suppress all time subscripts where it is unambiguous that a variable is endogenous. All exogenous variables are constants and need no time subscript.

2.1 State variables

An equity is the object of trade in this model. Its fundamental value is the net present discounted flow of its future dividends ν. The NPV variable evolves as a continuous-time Markov process.

For tractability ν takes one of two numbers

$$\nu \in \{\nu_g, \nu_b\}, \quad \nu_g > \nu_b \tag{1}$$

where ν_g is the "good" value and ν_b is the "bad" value. The asset is traded in an economic state described by a second state variable x,

$$x \in \{x_s, x_v\} \tag{2}$$

The economic state is said to be "stable" when $x = x_s$ or "volatile" when $x = x_v$. The economic state is assumed to be *public knowledge* here and through section 3. The assumption is relaxed in section 4.

The value switches between ν_g and ν_b at a state-contingent Poisson rate λ_x,

Table 1: Switching rates for ν

		asset switches at rate
while $\{$	$x = x_s$	λ_s
	$x = x_v$	λ_v

This model is interested in parameterizations in which $\lambda_v >> \lambda_s$.

2.2 Agents and preferences

2.2.1 Liquidity supply

A *limit order* is defined as a strike price p at which a trader commits to buy or sell one unit[4] of the book's asset so long as the order remains in force.

A countably infinite stock of identical book agents supplies limit orders in continuous time. They have risk-neutral preferences, start with common priors, and monitor trade on the book continually. Book agents may enter the book by quoting limit order strike prices, while active they may change their prices, and they may exit at any time by canceling their orders. While active, book agents suffer a per-order flow opportunity cost of the next-best use of their capital.

Because agents are identical and have identical beliefs, for the sake of exposition I refer to the set of active agents as "the book" and use the convention of one limit order per agent. The value space is of size two, so the book's belief distribution on the asset's value is just a number,

$$b = \Pr(\nu = \nu_g) \tag{3}$$

---

[4]Empirical limit orders usually offer to trade more than one unit. In this model such an order would be represented as multiple single-unit orders with the same strike price.

2.2.2 Liquidity demand

A *market order* is defined as a quantity $q \in \mathbb{Z}$ a trader orders for immediate transaction at the best prices available.

Market orders are submitted by a single representative "market agent." The market agent receives a Poisson opportunity to move and, upon receiving the opportunity, submits an order size q that maximizes a stochastic utility function. The market's utility takes the asset value and the economic state as arguments, so q is informative about the asset value. By convention $q \in \mathbb{Z}$, where positive q signifies a market buy, negative q signifies a market sell, and $q = 0$ signifies the market agent chooses not to trade. Importantly, the book cannot observe a "market order of size zero."

The market agent's utility is assumed quasiconcave as in Glosten (1994). Therefore the market agent does not explicitly take into account the affect its current behavior has on future prices. A fully strategic market trader would consider manipulating the price or at least staggering its trade. Back and Baruck (2004) give theoretical defense for the convention of an agent with a Poisson move rate who maximizes a quasiconcave utility function. In models conforming to the Glosten-Milgrom (1985) tradition, Back and Baruck find informed agents would capture more of their informational surplus by staggering their market orders at a stochastic rate and limiting the order size. The market agent is this model represents the reduced-form outcome of their structural order-submission strategy.

2.3 Mechanism of trade

The *limit order book* is a four-tuple:

$$\mathcal{B} = \left\{ A, B, \{p_q\}_{q=1}^{A}, \{p_q\}_{q=-1}^{-B} \right\} \tag{4}$$

$$\text{using } p_q \leq p_{q+1} \ \forall q$$

where $A \in \mathbb{N}_0$ is the current number of ask-side limit orders and $B \in \mathbb{N}_0$ is the current number of bid-side limit orders.[5] The $p_q \in \mathbb{R}_+$ are the strike prices of the outstanding limit orders indexed by q. The index of the strike price is called the "position" of the limit order.[6] A positive q signifies an order on the ask side of the book; a negative q signifies an order on the bid side of the book. The same notation q is used for position and for the order size because a market agent who buys q of the asset is matched to q limit orders on the ask side; a market agent who

---

[5]Sell limit orders are said to lie on the "ask side" of the book because they are asking prices; buy limit orders are said to lie on the "bid side" of the book because they bid up the best buying price. A market buy order is matched to an ask-side limit order because a buyer is matched with a seller. A market sell order is matched to a bid-side limit order because a seller is matched with a buyer. To avoid confusion I will refer to limit orders by the side of the book they are on.

[6]Book agents will compete for position, and the equilibrium will have the envy-free characteristic of position auctions (Edelman, Ostrovsky and Schwarz 2007). There may be a revenue equivalence property between the limit order book and position auctions.

sells q of the asset is matched to q limit orders on the bid side.

The book is the only organizer and executor of trade for its asset. It recognizes only limit orders and market orders. Trade takes place in continuous time. It is costless to submit market orders. When the book receives a market order, it fills the order using the best available prices, *i.e.* by buying (selling) at the lowest (highest) limit order prices outstanding. If the book were to receive more than one market order simultaneously, it would execute one of the market orders randomly and ignore the others. If filling a market order were to require more than the available quantity of limit orders, the book fills as much of the market order as possible (exhausting one side of the book) and cancels the remainder.

It is costless[7] to quote limit orders. Limit orders may be quoted, changed or withdrawn at any time. If the book were to receive a limit buy order with a strike above the ask, or a limit sell order with a strike below the bid, it would regard the limit order as spurious and cancel it. If many limit orders have the same price, the book fills the orders by their priority in time or, if they arrived at the same time, randomly.

2.4 Equilibrium concept

The equilibrium concept is subgame-perfect Nash equilibrium. An *equilibrium on the limit order book* is a book state $\mathcal{B}$ such that

1. There are $A, B \in \mathbb{N}_0$ active orders on the ask and bid side.

2. No active book agent can gain by quoting a different price.

3. No inactive book agent can gain by entering, and no active agent wishes to exit.

4. When the market agent moves it posts q maximizing its utility.

2.5 Existence and uniqueness

Under general conditions proving the existence of equilibrium would require solving an intertemporal fixed-point problem. Two assumptions will simplify the task: quasilinear utility and continuous-time undercutting.

Let the market agent maximize the parameterized function,

$$U_\theta(q, p) = \theta q - \alpha(x)q^2 - p(q) \tag{5}$$
$$\text{where } \theta = \nu + l$$

which is a version of quadratic mean-variance utility. The market chooses an integer $q \in \mathbb{Z}$ because the book contains a discrete number of limit orders. The current NPV ν and the state

---

[7]Costless is the sense that the book does not charge for the privilege. Liquidity suppliers experience the per-order flow opportunity cost for maintaining a limit order.

x enter directly into the function, cementing the way information is conveyed by order size. The assumption of quasilinearity eliminates wealth effects.

Call the market agent's demand parameter θ its "type," as in mechanism design. The type is a linear combination of the asset's current value ν and a normally distributed liquidity shock,

$$l \sim \mathcal{N}(0, \sigma_l) \tag{6}$$

The liquidity shock l is a random variable that is uninformative (in the sense of being independent of ν and x).

The second term in (5) is the market's risk aversion coefficient, which is a function of the state of the economy. The coefficient increases in the volatile state, $\alpha(x_v) > \alpha(x_s)$, due to the added risk. The third term $p(q)$ is the total money transfer as a function of the order size

$$p(q) = \begin{cases} \sum_{i=1}^{q} p_i & \text{if } q \in \mathbb{N}^+ \\ -\sum_{i=-1}^{q} p_i & \text{if } q \in \mathbb{N}^- \\ 0 & \text{else} \end{cases} \tag{7}$$

where p_i are the strike prices of the i^{th} positioned limit orders.

The fourth equilibrium condition requires the market agent to post $\arg\max_{q \in \mathbb{Z}} U_\theta(q, p(q))$ when it moves. Due to the quasilinear form it is possible and convenient to write an indicator function equaling one when the market buys at least q at current prices,

$$\mathbb{I}_{buy}(q, p_q) = \begin{cases} 1 & \text{if } \theta > 2\alpha(x)q - \alpha(x) + p_q, \ q > 0 \\ 1 & \text{if } \theta < 2\alpha(x)q + \alpha(x) + p_q, \ q < 0 \\ 0 & \text{else} \end{cases} \tag{8}$$

The convenience is that the indicator is not functionally related to the other prices on the book due to the lack of wealth effects. The book agents' risk-neutral profit function can be written

$$\pi_q(p) = \int_\Theta \left[p - \mathbb{E}(\nu(t) \mid \theta, b) \right] f_\theta(\theta; s, b) \, \mathbb{I}_{buy}(q, p) \, d\theta \tag{9}$$

Such a profit function is well-behaved for the purposes of equilibrium:

Proposition 1. *If $\sigma_l > \frac{\nu_g + \nu_b}{2}$, the profit function is single-peaked and ordered[8] in the positive orthant.*

The profit function is single-peaked and ordered because liquidity suppliers face a tradeoff between revenue on the one hand and the fill rate and picking-off risk (Foucault 1999) on the other. For intuition, consider an ask-side order. Its profit function is single-peaked because higher prices increase revenue but decrease the probability the market will find the price agree-

---

[8]On the ask side, if $\pi_q(p) > 0$ then $\pi_q(p) > \pi_{q+1}(p)$. On the bid side, if $\pi_q(p) > 0$ then $\pi_q(p) > \pi_{q-1}(p)$.

able enough to fill. In addition the higher the price the more likely the market agent would fill the order because the asset's value is high, increasing picking-off risk. The profit functions are ordered due to the same tradeoff—the same price in a worse position wins fewer expected profits because it is less likely to fill and more likely to bear picking-off risk.

A quadratic or mean-variance parameterization is typical in microstructure because it is often the outcome of a more complicated dynamic investment problem; for a discussion of the dynamic CARA-normal setting see Vives (2008) chapter 8.

The second simplifying device is continuous-time undercutting. Technical details are given in section 2.7. The continuous-time game structure in this model allows agents to undercut one another at speeds arbitrarily close to instantly. Instantaneous undercutting pushes gains to deviation to zero. A deviator would like to raise prices but could gain nothing by it because before a market order arrived another agent would Bertrand undercut the deviator. The deviator could undercut again, initiating a cycle of Bertrand undercutting, but due to continuous time the cycle finishes immediately.

An *equilibrium on the limit order book* can now be redefined as the set of prices satisfying

---

Figure 1: Profit functions and equilibrium prices on the ask side for $L = 3$ orders.
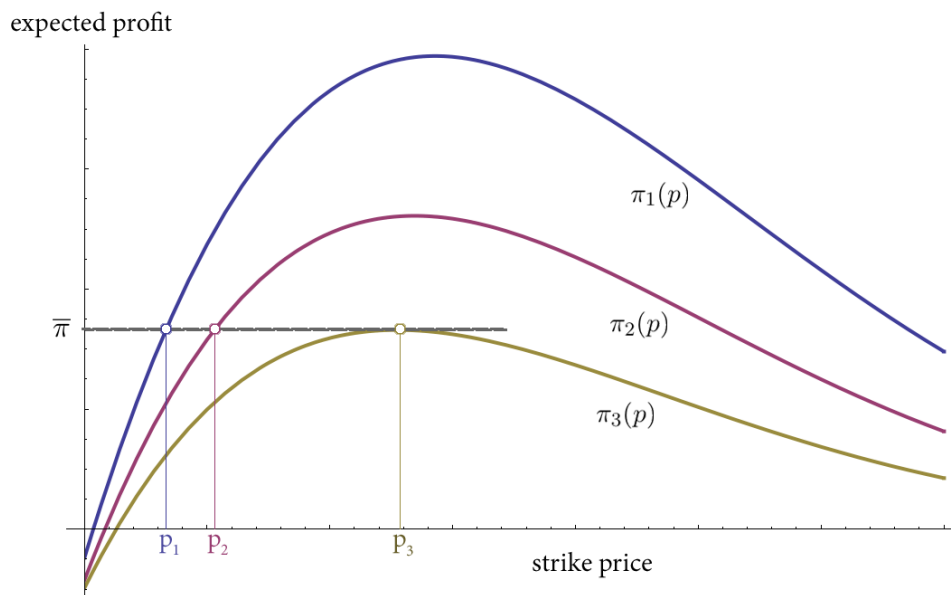
These are the profit functions conditional on position for three unit limit orders. It is easy to see graphically why there is a unique set of prices at which $\pi_q(p_q) = \max \pi_3(p) = \overline{\pi}$ in order.

Nash-subgame perfection and the free entry condition ($\pi_q \geq c$):

$$A = \arg\max_{A \in \mathbb{N}_0} \left\{ A \mid \max_p \pi_A(p) > c \right\} \tag{10a}$$

$$B = \arg\max_{B \in \mathbb{N}_0} \left\{ B \mid \max_p \pi_{-B}(p) > c \right\} \tag{10b}$$

$$\pi_1(p_1) = \pi_2(p_2) = ... = \pi_A(p_A) = \max_p \pi_A(p) \equiv \overline{\pi}_A \tag{10c}$$

$$\pi_{-1}(p_{-1}) = \pi_{-2}(p_{-2}) = ... = \pi_{-B}(p_{-B}) = \max_p \pi_{-B}(p) \equiv \overline{\pi}_B \tag{10d}$$

Proposition 2. *Prices* $\{p_q\}_1^A, \{p_q\}_{-B}^{-1}$ *satisfying* (10a), (10b), (10c) *and* (10d) *exist for any belief b. Due to single-peakedness and ordering they are unique.*

For intuition, consider what should happen if any agent deviates from the equilibrium prices in the scenario of figure 1. Suppose the agent in position 1 tries to overbid by posting $p = p_1 + \epsilon < p_2$. At the strike price $p_1 + \epsilon$ the first-positioned agent would earn greater expected profits: $\pi_1(p_1 + \epsilon) > \overline{\pi}$. They will not last; the deviator has created a profit opportunity for the other agents, who can Bertrand undercut the deviator to capture its surplus. Due to continuous time the deviator is undercut instantaneously and does not enjoy a moment of the higher-than-equilibrium expected profits. Instantaneous undercutting brings the expected payoff of the subgame to zero. Last, agents never wish to quote prices that would earn profits less than $\overline{\pi}$ because they can always do better by quoting an equilibrium price.

An important assumption behind the equilibrium is that it is competitive—no agents communicate or coordinate. It is possible to imagine scenarios in which agents form coalitions to keep prices high and punish any deviators with Nash threats. Such coalitions would not last long on a book with free entry because oligopoly profits would attract entrants.

Last, the agents cannot do better by playing mixed strategies.

Proposition 3. *There exist a continuum of mixed-strategy equilibria. The expected profits of mixed-strategy equilibria are bounded by the competitive profits of the pure-strategy equilibrium.*

2.6 Learning

The book updates its beliefs during two moments: (i) after a market order is filled and (ii) during periods of no trade. When a market order is filled it updates using Bayes' law:

$$b' = \frac{b \ f_\theta(\theta^{-1}(q, \mathcal{B})}{b \ f_\theta(\theta^{-1}(q, \mathcal{B}) + (1 - b) \ f_\theta(\theta^{-1}(q, \mathcal{B})} \tag{11}$$

where $\theta^{-1}(p, q)$ denotes the interval of market agent types for which $q = \arg\max_q U_\theta$ at the prevailing prices on the book $\mathcal{B}$. It is possible to speak of the moment "after" a market order is filled because the market order arrives in layered time; see section 2.7. Having updated its priors the book changes its prices and proceeds in continuous time.

Proposition 4. *The posterior belief b' is increasing in the order size q.*

For a large enough market buy order the update increases, $b' > b$. That being so, depending on the state of the book a "small" market buy may lower the posterior belief. A small enough market buy can lower the posterior because a market agent who knows the asset is underpriced would post a large order not a small one. (The relation holds *vice versa* for a small market sell.)

Next, the book also updates beliefs during periods of "market silence," periods during which the book observes no market orders. During market silence the book does not receive any news in the form of market orders. Nevertheless, "no news is still news," because it is always possible the market received its Poisson opportunity to move yet declined to trade at the quoted prices. Due to this possibility the book continually updates its prior. The law the book uses to update beliefs during silence is the "good news bad news" Bayesian differential equation.[9]

The book's updating law during market silence is a function of the fill rate, the rate at which nonzero market orders arrive. To define the fill rate, first define the book's spread,

$$S = [p_{-1} \ p_1] \tag{12}$$

which is the distance between the best bid and best offer (or asking price). By convention the best bid is called the "bid," and the best offer is called the "ask." Define the fill rate as the probability the market agent draws a liquidity shock such that it is interested in trading given the bid-ask spread:

$$r(S, \nu, \alpha) = \mu \cdot (1 - \Pr[\theta \in S \pm \{\alpha\} \mid \nu, \alpha]) \tag{13}$$

The fill rate quantifies the likelihood the market agent is motivated to trade should it receive the opportunity to move. The market agent trades when its random type θ realizes a value sufficiently far from the spread. As the spread dilates, it takes a relatively more extreme liquidity shock for the market to find any prices agreeable, diminishing the fill rate. As the spread approaches zero, the fill rate increases. Note the fill rate is always less than market agent's arrival rate even if there is no spread, due to the α term:

$$r(0, \nu, \alpha) < \mu$$

In equilibrium the spread is a function of the book's beliefs, so it is possible to reduce the notation surrounding the fill rates by writing $r_g(b)$ for the equilibrium fill rate conditional on a good asset value and writing $r_b(b)$ for the equilibrium fill rate conditional on a bad asset value. The "good news bad news" updating law during silence is

$$\frac{\partial b}{\partial t} = \underbrace{b(1 - b)(r_b(b) - r_g(b))}_{\text{information gain}} + \underbrace{2(1/2 - b)\lambda_x}_{\text{information decay}} \tag{14}$$

---

[9]A derivation can be found at the end of the appendix.

I call the first term in the continuous-time update the "information gain" because it describes the flow of information the book draws from silence. I call the second term in equation (14) the "information decay" because it describes the decay of the prior due to the possibility the asset might switch.

2.7 Continuous-time game structure

This last model section may be skipped on a first reading.

The device of continuous time removes any momentary incentive to deviate from equilibrium. In discrete time such an incentive may still exist (depending on the structure of the payoffs) because a deviator would enjoy a moment of noncompetitive profits, and in general the momentary payoff could be high enough to justify deviation. Agents who take turns committing to prices might well play a kind of dynamic oligopoly game with Edgeworth price cycles (Maskin and Tirole 1989). Price cycles cannot be removed by limiting the discrete time increment to zero; the price cycle simply revolves faster and faster.

Economically, it is unlikely dynamic oligopoly profits would survive long on a book with free entry because oligopoly profits attract competitive entrants. Still the technical problem of defining the equilibrium remains. To sidestep the potential for price cycles this game is not defined as a simple limit of a discrete-time game. The limit order book equilibrium will be characterized as the limit of a game with inertial strategies, defined by Bergin and MacLeod (1993). Loosely, inertia preserves the useful limiting characteristics of discrete-time games (namely, a well-ordering property) without using the payoff structure of the discrete-time game.

The phenomenon of price cycles may be empirically relevant to market maker competition; the idea is pursued in Garriott (2010wp). The use of the inertial strategy was inspired by Roşu (2009) and its unpublished technical aside, Roşu (2006wp). This paper completes the equilibrium definition contained in Roşu by offering the sufficient condition required by Bergin and MacLeod (1993) to construct a limiting outcome.

Following Bergin and MacLeod briskly, let a book agent's stage game action space A_i contain all the measurable functions that map from beliefs to prices,

$$A_i : b \rightarrow P \subset \mathbb{R}_+ \tag{15}$$

The space P is compact to satisfy the conditions in Bergin and MacLeod. Notice actions map not to some particular price but to a full schedule of prices the agent would quote given beliefs. Let $A = \prod A_i$.

The game takes place at times $T = [0, \infty)$. Use the definition of layered time from Roşu (2006wp): let $\nu(t) \in \mathbb{N}_0$ be the layer defined on T, and call T^ν the space of layered time. The convention of layered time defines "when" multiple stage games are played if they need to occur at the same time t. It is convenient for stage games to be played more than once when a market

12

order arrives and when agents enter or exit the book, and layering provides a convention to index when they happen. Almost everywhere $\nu(t) = 0$. Layered time is still measurable. At layered time t there are $N(t) \in \mathbb{N}_0$ agents on the book. An outcome vector is the history of action for an agent, a measurable function $h_i : T^{\nu} \to A_i$. Let $H = \prod H_i$. Agents choose strategies to maximize $U_i(h) = \int_{T^{\nu}} u_i(h(t)) \, e^{-t} dt$[10] where the integral is taken using the definition of layered time. The utility function $u_i(h(t))$ equals the expected profit of agent i given the prices at $h(t)$.

Define a metric on the outcome at time t, $d_i(h, h', t)$ and let $d(h, h, t) = \sum d_i(h_i, h_i, t)$. Using d_i define a metric on the outcome paths,

$$D_i(h_i, h_i', T) = \int_T d_i(h_i, h_i', t) \, e^{-t} dt \tag{16}$$

and let $D = \sum D_i$. A strategy is a measurable function that maps the space of outcomes and times to an action, a map for which the future does not affect current decisions:

$$x_i : \; H \times T^{\nu} \to A_i \tag{17}$$
$$x_i(h, t) = x_i(h', t) \text{ for } h, h' \in H \text{ s.t. } D(h, h', [0, t)) = 0. \tag{18}$$

Last, notate the information at time t and history h as $\mathcal{I}(h, t)$ and define the equilibrium action for the n^{th}-positioned limit order submitter:

$$a_n^*(b) = p_n \text{ satisfying (10a), (10b), (10c) and (10d)} \tag{19}$$

The equilibrium action may be interpreted as "take the n^{th} positioned equilibrium price."

Restrict the strategy set to those that are closed on the left by imposing an inertia condition. Inertia forces agents to retain actions for at least nonzero measures of time and on intervals that are open on the right. Inertia is one resolution of the well-ordering problem; for a discussion see Simon and Stinchcombe (1989).

Definition: A strategy x_i exhibits inertia if for any arbitrary time and history (t, h) there exists an $\epsilon > 0$ and an action mapping $a(\mathcal{I})$ such that:

$$D_i(x_i(h'), a(\mathcal{I}), [t, t + \epsilon)) = 0$$
$$\text{for every } h' \in H \text{ such that } D(h', h, [0, t)) = 0.$$

Theorem 1 in Bergin and MacLeod shows a strategy space restricted by an inertia condition is consistent with a unique outcome (a.e.). Theorem 2 ensures that strategies in the completion of the set of strategies with infinitesimal inertia are still consistent with unique outcomes (a.e.). Theorem 3 shows strategies x in the completion are subgame perfect if and only if Cauchy sequences of strategies $x^m \to x$ constitute sequences of ϵ^m-subgame perfect Nash equilibria with

---

[10]The rate of time preference is immaterial so I simply write e^{-t}.

$\epsilon^m \to 0$.

It is necessary to give such a sequence in order to construct the limiting outcome $h(x, t)$. Such a sequence is required on a timeline that lacks the well-ordering property to make sense of how one agent can react to a deviator "before" the deviator can move again. Without such a formalism for "continuous-time timing" the strategies may not be consistent with any outcome. For example, suppose agent A plays the strategy "quote above the equilibrium price unless someone else undercuts, in which case quote the equilibrium price," and agent B plays the strategy "quote the equilibrium price unless someone else overquotes, in which case undercut." There is no fixed point—neither will be able to agree on a set of prices. Accordingly the formalism below creates an infinitesimal "mover delay" (that is open on the right, to preserve inertia). At the limit, if any agent changes action at time t, it is as if the agent cannot change the action between t and dt, opening a window in which another agents may react. This is similar to continuous-time finance, in which no trade may occur between t and dt. Here is such a restriction, which formalizes the notion that agents can undercut "instantaneously" relative to the deviator.

$$\forall t \; \exists a(\mathcal{I}) \text{ s.t. } D(x_i(h'), a(\mathcal{I}), [t, t+\epsilon)) = 0 \text{ and } D(x_i(h'), a(\mathcal{I}), [t-\epsilon, t)) \neq 0,$$
$$D(x_i(h'), a(\mathcal{I}), [t, t+\tau(\epsilon)]) = 0 \tag{20}$$
$$\text{for every } h' \in H \text{ such that } D(h', h, [0, t+\tau(\epsilon))) = 0.$$

Let agents take turns quoting prices in rounds of length $\tau(\epsilon)$. Agents who switch actions cannot do so again for time ϵ. The candidate strategy is:

$$x_i^\epsilon(q \mid q_{-i}) = \begin{cases} \text{Quote } p_q^* \text{ if agents } -i \text{ are quoting } p_{q_{-i}} = p_{q_{-i}}^* \\ \text{If any agent deviates, switch to that agent's position } q_{-i} \end{cases} \tag{21}$$

Define $\tau(\epsilon)$ to be large enough that deviators cannot make short-term excess profits by switching giving the payoff structure. It does not matter how large $\tau(\epsilon)$ is (and it can be quite large relative to ϵ) because it will be taken to zero with ϵ, though slower than ϵ.

$$\forall p', \quad \int_t^{t+\tau(\epsilon)} \pi\left(p' \mid \{p_i\}_1^{I-1}\right) + \int_{t+\tau(\epsilon)}^{t+\epsilon} \pi\left(p' \mid p_i', \{p_i\}_i^{I-2}\right) \leq 0 \tag{22}$$

The switching lag $\tau(\epsilon)$ gives other agents a chance to react to any switch in action "relatively fast enough" that the agent who changed cannot profit by the switch.

3 Information and liquidity

This section gives the main results for the benchmark model in which x is public information. In summary, the book reacts to large buy orders by raising the bid and ask prices, and it reacts

14

to large sell orders by lowering the bid and ask prices. The spread is widest when the book has a relatively uninformative prior ($b \approx 1/2$) and lowest when the book is more certain ($b \approx 0$ or $b \approx 1$), so information improves liquidity. The fill rate is highest when the spread is low, so liquidity improves information revelation.

The results on information aggregation use the belief update during silence as the main tool of analysis. I say the book can "sustain" a belief if the belief update during silence tends toward an absorbing state that is not $1/2$. If the book *cannot* sustain a belief, b tends toward the "uniform prior" of $1/2$. The distance away from $1/2$ of an absorbing state is a measure of the book's ability to aggregate information because an absorbing state forms a threshold below which the book does not lose information. In contrast if b always tends toward $1/2$ the book is always losing information and cannot be said to aggregate. These results formalize the idea of limitations on information aggregation and will create the possibility of a liquidity drought when section 4 relaxes the assumption that x is public information.

Any graphical illustrations in this section and others will use the parameters in table 2:

Table 2: Simulation parameter values

Asset:		Market:		Book:
$\nu_g = 20$	$\nu_b = 10$	$l \sim \mathcal{N}(0,5)$	$\mu = 4$	$c = 1/50$
$\lambda_s = 1/100$	$\lambda_v = 1$	$\alpha(x_s) = 1$	$\alpha(x_v) = 2$	

In a stable economy the asset NPV switches between values of 10 and 20 on average once in 100 units of time. In a volatile economy the asset switches on average once per unit time.

3.1 Comparative statics

Notate the equilibrium ask and bid as functions of b: $p_1(b)$ and $p_{-1}(b)$.

Proposition 5. *Character of the bid and ask prices. Under technical conditions $p_{-1}(b) \geq \nu_b$ and $p_1(b) \leq \nu_g$,*

- *The equilibrium bid and ask are increasing in b. At $b = 1/2$ they surround the mean value.*

$$\frac{\partial p_1(b)}{\partial b} > 0 \ and \ \frac{\partial p_{-1}(b)}{\partial b} > 0$$
$$and \ at \ b = 1/2, \ \frac{p_1(b) + p_{-1}(b)}{2} = \frac{v_g + v_b}{2}.$$

The technical conditions $p_{-1}(b) \geq \nu_b$ and $p_1(b) \leq \nu_g$ sign derivatives that otherwise could not be signed without an explicit closed-form for the equilibrium bid and ask prices. A sufficient condition for $p_1(b) \leq \nu_g$ for at least low b is that every book agent on the ask side would prefer to quote prices lower than ν_g at $b = 0$, or $\pi_1'(\nu_g) < 0$. Symmetrically, $p_{-1}(b) \geq \nu_b$ for at least high b is that $\pi_{-1}'(\nu_b) > 0$. These conditions will be enough to sign the fill rates.

Intuitively, a more valuable asset trades at higher prices. If the agent quoting the ask did not raise its price with b, it would face increasing information risk. If the agent quoting the bid did not raise its price with b, it would face a vanishing fill rate. The spread grows wide at the uniform prior of $b = 1/2$ because at such a prior that limit order submitters believe they are most uninformed relative to the market. The less confidence they have in what they know about the asset, the more they fear trading with potentially informed traders. Book agents anticipate their vulnerability by raising the price of ask-side orders and lowering the price of bid-side orders.

The essential comparative static in the model gives signs on the fill rates.

Proposition 6. *Fill rates in equilibrium:*

- *Fill rates are decreasing in the spread's size. For $S_1 \supset S_2$, $r(S_1, \nu, \alpha) < r(S_2, \nu, \alpha)$.*

- *Denote the equilibrium spread at the book's current beliefs by $S(b)$. The market agent is most likely to trade when the asset is mispriced:*

$$r(S(b), \nu_b, \alpha) > r(S(b), \nu_g, \alpha) \qquad \text{for } b > 1/2 \qquad (23a)$$

$$r(S(b), \nu_b, \alpha) = r(S(b), \nu_g, \alpha) \qquad \text{for } b = 1/2 \qquad (23b)$$

$$r(S(b), \nu_b, \alpha) < r(S(b), \nu_g, \alpha) \qquad \text{for } b < 1/2 \qquad (23c)$$

$$r(S, \nu, \alpha(x_s)) > r(S, \nu, \alpha(x_v)) \qquad \text{for all } S, \nu \qquad (23d)$$

The set of equations (23a)–(23d) compares the fill rates in equilibrium. The fill rate comparisons justify many of the results concerning the book's updating program. The intuition behind them is that the market agent is most likely to trade when the asset is mispriced—when the spread is far away from the true value of the asset. Should the bid and ask lie near ν_g when the asset is actually worth ν_b, the market can exploit the prices by selling an asset worth ν_g for

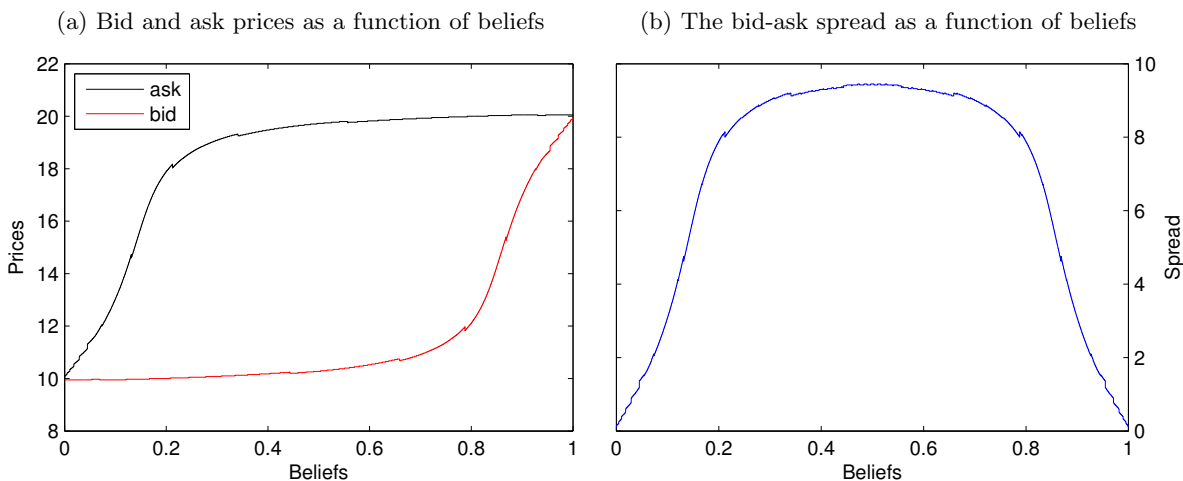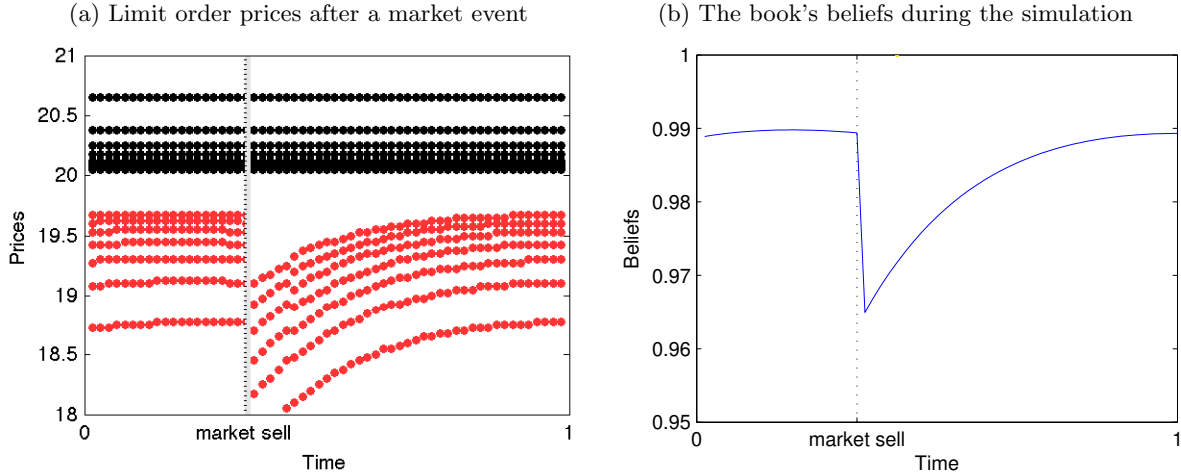Figure 1: Illustrations of the bid and ask prices for beliefs $b \in [0, 1]$

(a) Bid and ask prices as a function of beliefs

(b) The bid-ask spread as a function of beliefs

16

Figure 2: The book reacts to a single market event.

(a) Limit order prices after a market event   (b) The book's beliefs during the simulation

Black dots signify ask-side prices; red dots signify bid-side prices.

much more than it is worth. The incentive to take advantage of a mispricing vanishes as the equilibrium spread approaches the true value of the asset.

## 3.2 Comparative dynamics during market silence

As discussed in the section on learning, the book updates after market orders and during silence. To provide some preliminary intuition, figure 2 illustrates the book's response after a single market order. Its left-hand chart shows a simulation of the book's prices, and its right-hand chart graphs the book's beliefs. The collection of strike prices all near $\nu_g = 20$ is reflected in the book's strong belief the asset is good, $b \approx 0.99$. One-third of the way through the simulation, a market sell order arrives and disturbs the book from its belief the asset is good. Seeing the market sold, the book considers the possibility the asset might have switched value. Hence the book lowers its belief in accordance with proposition 4. Notice the spread increases afterward. The spread is a reflection of the book's uncertainty, yet it increases after an informative event. An informative event provokes uncertainty because it indicates the asset might have switched in value, which would make the information the book has gathered from history outdated. To the extent the book believes the asset might have changed in value, it discounts the information it has gathered.

After the market sell event, the book's priors slowly recover and asymptote near $b = 0.99$. The next proposition gives the structure of the book's updating law during such a period.

Proposition 7. *The information gain and the information decay in the two economic states:*

- *Due to proposition 6 the information gain has two stable roots $b = 0, 1$ and one unstable*

17

Figure 3: The velocity of beliefs during market silence $\partial b/\partial t$, by economic state

(a) In a stable economy x_s | (b) In an unstable economy x_v
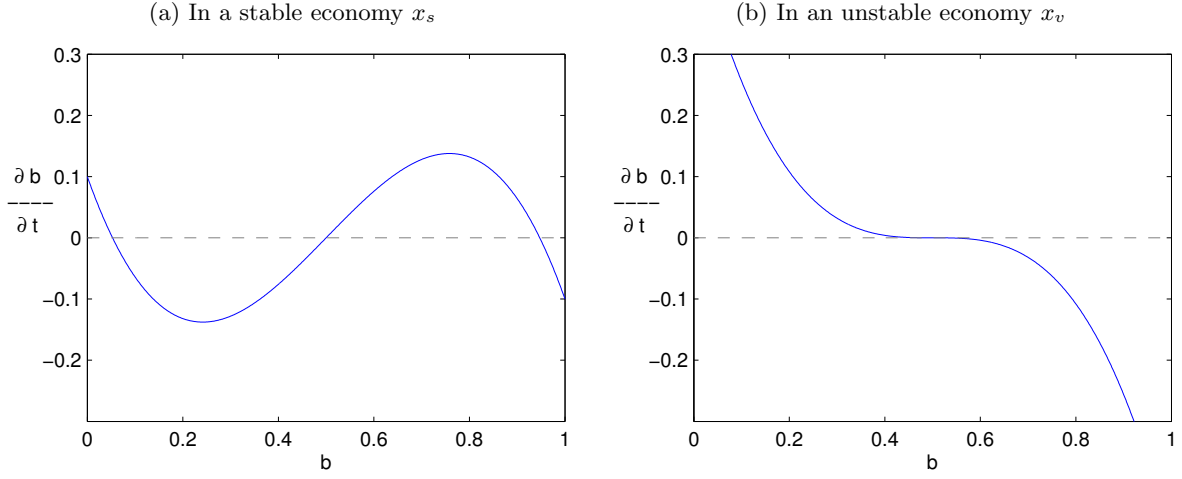
root $b = 1/2$.

$$for\ b > 1/2,\ r_b(b) > r_g(b),\ hence\ b(1-b)(r_b(b) - r_g(b)) = 0\ at\ b = 1$$
$$for\ b < 1/2,\ r_b(b) < r_g(b),\ hence\ b(1-b)(r_b(b) - r_g(b)) = 0\ at\ b = 0$$
$$for\ b = 1/2,\ r_b(b) = r_g(b),\ hence\ b(1-b)(r_b(b) - r_g(b)) = 0$$

*The information decay term has one root and it is stable:* $b = 1/2$.

- *For* λ_v *is sufficiently large, in the volatile economy the Bayesian update* (14) *is mean-reverting conditional on periods of market silence.*

Due to proposition 6 the information gain term is positive when $b > 1/2$ and negative when $b < 1/2$. The information gain term tends to confirm whatever beliefs the book holds because the most likely reason the book would observe no trade is that the spread is surrounding the true value of the asset; otherwise the market agent would be eager to trade. The decay term in contrast is a line crossing $1/2$—it exerts a mean-reverting influence. The combination of a linear decay function with the information gain pushes its stable roots in toward $1/2$, so the decay term limits aggregation.

Furthermore, if the asset's switching rate is too high the book cannot gather enough information from market silence to stem the decay of its store of signals. For an illustration see figure 4. Mathematically, the information decay is a straight line passing through $1/2$ with negative slope λ_x. Should the slope become sufficiently steep the absolute magnitude of the information decay term is so great it overwhelms the stable roots of the information gain and, since it passes through the root at $b = 1/2$, the middle root turns stable and unique.

For a high enough information decay any agent who observes a staggered signal would not

Figure 4: The book learns from market orders the asset value switched to 10
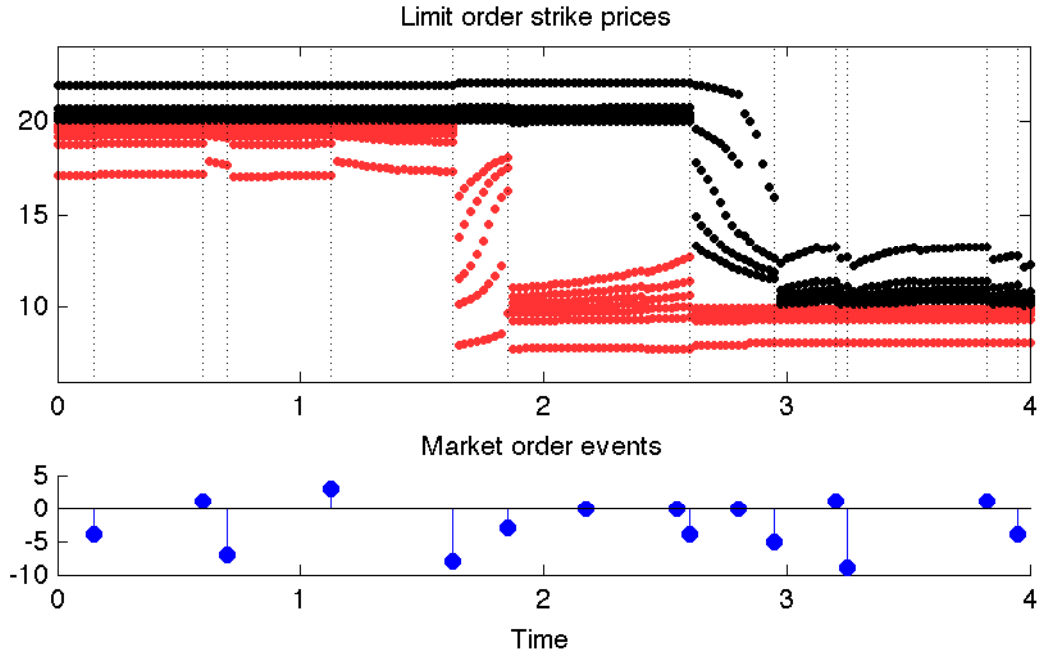
be able to sustain beliefs. On a competitive order book, the presence of profits aggravates the learning problem. Profits dilate the spread beyond the efficient distance, which results in less than efficient fill rates. Beliefs during silence can mean-revert on a competitive book for lower λ_v than for efficient books.

Proposition 8. *The efficient ask prices are less than the competitive prices: $p_q^e(b) < p_q^*(b)$ for $q > 0$, and the efficient bid prices are greater than the competitive prices $p_q^e(b) > p_q^*(b)$ for $q < 0$.*

To cement intuition on the learning process, figure 3 presents a second illustration, a session of trade lasting four time units in which the economy is stable. In its scenario the book believes the asset value is good when in fact $\nu = \nu_b$. The book learns the truth from the order flow. Its top graph displays the book's prices in time; the bottom graph gives the sizes of the market orders. In the simulation the market agent received 15 opportunities to move and posted an order during ten of them. Most of the orders were sell orders, and as the market posted them the book learned the asset had switched value and adjusted its prices downward. It is edifying to note the effect on the spread of small orders, for example the fourth-to-last and second-to-last market orders. At both points the book believed the asset was bad, yet the spread decreased after market buys. The buy orders were both for quantity one. A small buy order such as these can actually confirm the book's belief the asset is bad because at the prevailing prices a market buy of quantity 1 was more likely if the asset was $\nu = \nu_b$. A market agent would likely buy a great deal more than 1 if it were the case that $\nu = \nu_g$.

19

4   Endogenous liquidity droughts

This section relaxes the assumption the state of the economy is public information in order to generate an endogenous liquidity drought. Hereafter the book must use the market order flow to infer a hidden state with two dimensions: whether the asset is good, and whether the economy is stable. The drought occurs when the book loses confidence in economic stability. It can occur randomly (albeit rarely) as a result of variation in the order flow.

It is possible to prove that as the book's belief the economy is volatile grows, the absorbing states of its belief the asset is good move toward 1/2. In other words, if the book believes the economy is basically volatile the velocity of its asset-value beliefs look like figure 3b. If the book believes the economy is basically stable the velocity of its asset-value beliefs look more like figure 3a. To understand the drought, suppose the book has a "strong" belief the economy is volatile, meaning the absorbing state is "close" to 1/2. Then its belief the asset is good attenuates toward something "close" to the uninformative prior. As the belief attenuates the spread dilates, creating the drought. The drought is self-reinforcing because a dilated spread stems the flow of market orders and makes silence uninformative. A numerical example at the end illustrates.

The economic state is a hidden state that it follows a continuous-time Markov process with a Poisson switching rate. While the market agent does continue to maximize its utility (5) taking the asset value and economic state as given, the book must now infer both variables from the history of trade. The economic state's switching rates are notated using η_x,

Table 3: Switching rates for x and ν

	economy switches at rate	asset switches at rate
while $\begin{cases} x = x_s \\ x = x_v \end{cases}$	η_s	λ_s
	η_v	λ_v

Relaxing the assumption opens up a second dimension of uncertainty.[11] The second dimension increases the number of hidden states from two (ν_g or ν_b) to four:

Table 4: States of the game

state	asset ν:	economy x:	
1	ν_g	x_s	"good-stable"
2	ν_b	x_s	"bad-stable"
3	ν_g	x_v	"good-volatile"
4	ν_b	x_v	"bad-volatile"

To distribute probability mass over the four hidden states of nature the book uses three prior belief variables. As before the book learns from market orders using Bayes' rule, which takes

---

[11]The use of two dimensions of uncertainty has been used to study asset market failures before. Avery and Zemsky (1998) employ the assumption to study asset price bubbles; here it is applied to liquidity droughts.

the three prior variables as its arguments. As before the book uses the Bayesian differential equation from "good news bad news" during market silence, which is specified in the appendix. The relative fill rates still play a pivotal role in the velocity of beliefs during silence.

Proposition 9 proves a limiting result about the velocity of linear combinations of the belief variables during silence. Call these linear combinations the summary statistic variables b_ν and b_x. They are not sufficient statistics for the three prior belief variables; the proof and the numerical exercise both use the full space of prior belief variables. They are a notational and expositional convenience.

Table 5: Summary beliefs

$b_\nu = \Pr(\nu = \nu_g, x = x_s) + \Pr(\nu = \nu_g, x = x_v)$	"The belief the asset is good"
$b_x = \Pr(x = x_s, \nu = \nu_g) + \Pr(x = x_s, \nu = \nu_b)$	"The belief the economy is stable"

The value of b_x can be used to bound the absorbing states of $\partial b_\nu / \partial t$ during silence. For b_x low enough, b_ν tends toward a number arbitrarily close to $1/2$.

Proposition 9. *Roots of the belief updating law during silence for the four-state case. Let λ_v be high enough such that the belief update during silence in the volatile economy has one stable root at $1/2$. In the four-state model, for any ϵ-neighborhood around $1/2$ there exists a δ such that for $b_x < \delta$ the stable roots of the update of b_ν conditional on silence lie within $[1/2 - \epsilon, \ 1/2 + \epsilon]$.*
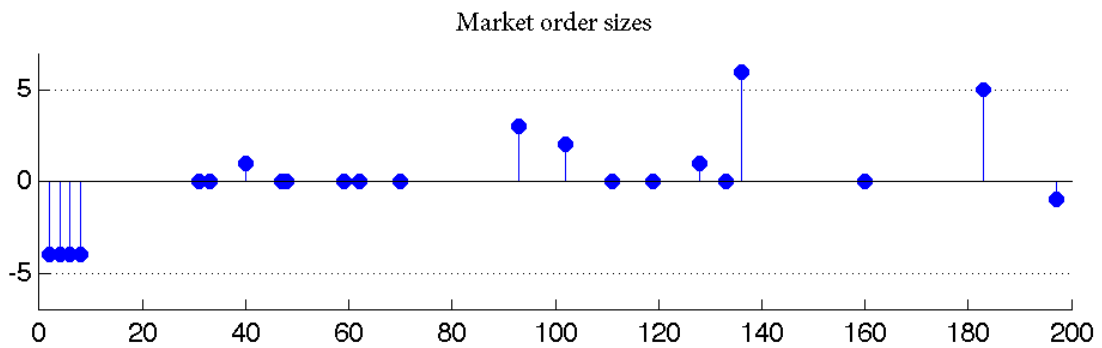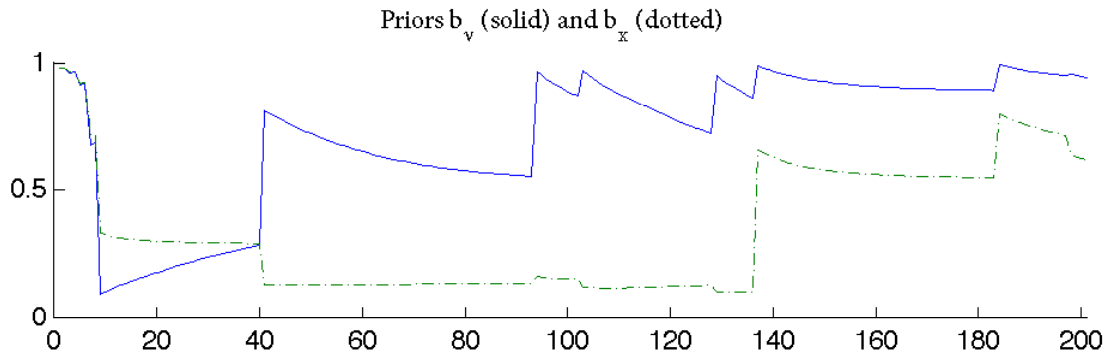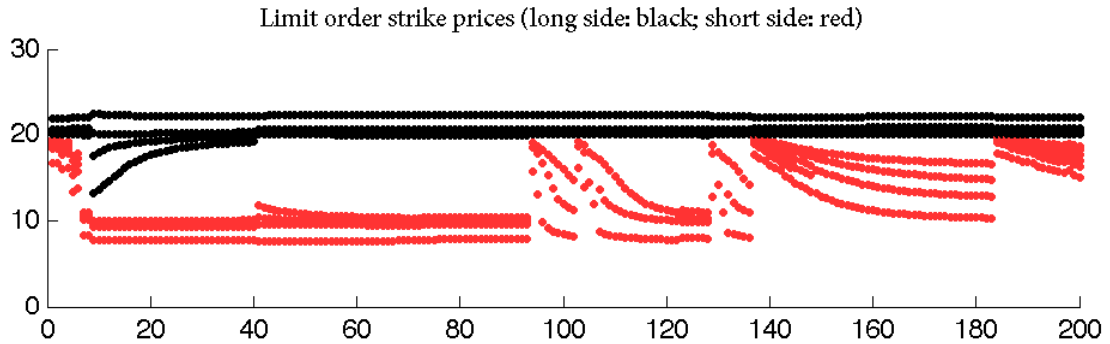
Should the book begin to doubt the economy is stable, the absorbing states of its asset-value belief attenuate. In other words, the book moves to a regime of updating in which it discards relatively more of its stock of signals as time passes. In the benchmark model the information decay was only a function of the fill rates. Here it is also a function of beliefs, so the rate of information decay has become endogenous; it changes with b_x.

The book can place probability mass on the volatile economic state even when the economy is stable. It can happen after a draw of market orders that look like they come from an unstable economy—a low probability event, but not impossible. For example, suppose the book believes the asset is good, and take a sequence of ten buy and sell market orders, alternating between market buy and market sell. If the order sizes are not "large," such a sequence can indicate stability because the market agent draws types from a symmetric distribution. It is normal for there to be a steady flow of varying orders. However, rearrange the same sequence of orders into a series of five successive market sells followed by five successive market buys, and the same set of orders tells a different story. Such a series of market orders would most likely come in a world in which the value changed to bad and then returned to good. The sequence now indicates volatility.

Since the prior b_x can fall even when the true state of the economy is stable, the book will sometimes discard valuable information. I conjecture the Borel-Cantelli lemma can be used here to show b_x falls arbitrarily low infinitely often because for any $b_x > 0$ it should be possible to

Table 6: Simulation parameter values

Asset:		Market:		Book:
$\nu_g = 20$	$\nu_b = 10$	$l \sim \mathcal{N}(0,5)$	$\mu = 4$	$c = {}^1/_{50}$
$\lambda_s = 1/100$	$\lambda_v = 1$	$\alpha(x_s) = 1$	$\alpha(x_v) = 2$	
$\eta_s = 1/1000$	$\eta_v = 1/10$			

Limit order strike prices (long side: black; short side: red)

Priors b_v (solid) and b_x (dotted)

Market order sizes

construct a finite sequence of market orders after which the book's Bayesian update would be b_x or lower.

The liquidity drought has a quality of stability. It is again due to an interaction between the spread and the fill rate. When the spread is wide the fill rate is low, so the book is unlikely to receive signals through market orders. Moreover the book cannot significantly revise its beliefs during silence because it expects few market orders to arrive—silence is uninformative when silence is expected. Hence the book can get "stuck" in a state of high spreads and no learning. Figure 5 presents a numerical example of such an endogenous drought. The simulation began with four market sell orders of quantity -4. The orders were not randomly generated but were placed there in order to "spook" the book into believing it was trading in the volatile economy.

5 Conclusions

Classical information aggregation results employ the assumption of a static underlying value. These results often obtain through limiting laws of large numbers. In contrast if the underlying value may shift an infinite series of staggered signals does not suffice for almost-sure convergence. The value of old signals decays. The margin between the rate at which information can be collected and the rate at which it is lost creates an upper bound on the sustainable precision of beliefs as time passes. The model order book in this paper gives a concrete example of such a bound, represented by the roots of the belief updating law during silence. As the example of the volatile economy illustrates, if the rate of information decay is too high it is not possible to sustain any nonuniform beliefs between signal draws. Because recent signals matter more than old ones, the process of forming beliefs about a moving underlying resembles the statistical technique of filtration more than the technique of taking an average.

The model book's informational problem is aggravated by the presence of profits. The competitive limit order book sets ask prices higher and bid prices lower than efficient prices. Since the competitive spread is wider it elicits fewer market orders than the efficient spread. Also, since the information gained from silence is greater when spreads are lower, the competitive book gains less information from silence. For certain values of the asset's switching rate, the markup that competitive book agents charge can make a difference between sustainable information aggregation and the complete mean-reversion of the belief during silence.

It is possible to speak of a "conjugacy" between information aggregation and liquidity provision. First, the book's ability to collect information is mediated by its ability to provide liquidity. When the spread is wide the book elicits few trades because few liquidity demanders would wish to trade at such prices. Since the model book learns through the order flow, by not providing liquidity it is also forgoing information. Moreover, the book's ability to collect information during silence is again mediated by liquidity supply. Market silence tends to confirm beliefs because the most likely reason the book does not observe trade is that the spread is surrounding the

true value of the asset. If the spread is wide, market silence is not particularly informative.

The converse is also true. The book's ability to provide liquidity is mediated by its success at information aggregation, which is the lesson of the endogenously generated liquidity drought. If some exogenous shock should cause the book to doubt economic stability, liquidity dries up and the book will cease to collect much information due to the dilated spread. Since liquidity departed for informational reasons, the book may not be able to collect enough information at the new spread to motivate its reentry. When capital stays out and the spread dilates, the book cannot learn much and remains stuck until the market agent draws a sufficiently high liquidity need. The droughts are a limitation on information aggregation because an adverse order flow can arise with positive probability even in the stable economy. Hence the book will eventually move to throw away its stock of signals.

The model conditions for good liquidity provision are thus related to the model conditions for information aggregation. In the benchmark two-state case, it is that the aggregator must be able to collect information faster than its information decays. In the four-state case, it is that the aggregator must have enough confidence in the worth of its stock of signals. As the book must have this confidence to fulfill both of its social functions, in this model context the liquidity provision business is the information aggregation business.

References

R.J. Aumann. Agreeing to disagree. *The annals of statistics*, 4(6):1236–1239, 1976. ISSN 0090-5364.

C. Avery and P. Zemsky. Multidimensional uncertainty and herd behavior in financial markets. *American Economic Review*, 88(4):724–748, 1998. ISSN 0002-8282.

K. Back and S. Baruch. Information in securities markets: Kyle meets Glosten and Milgrom. *Econometrica*, 72(2):433–465, 2004. ISSN 1468-0262.

J. Bergin and W.B. MacLeod. Continuous time repeated games. *International Economic Review*, 34(1):21–37, 1993. ISSN 0020-6598.

B. Biais, P. Hillion, and C. Spatt. An empirical analysis of the limit order book and the order flow in the Paris Bourse. *Journal of Finance*, 50(5):1655–1689, 1995. ISSN 0022-1082.

B. Biais, D. Martimort, and J.C. Rochet. Competing mechanisms in a common value environment. *Econometrica*, 68(4):799–837, 2000. ISSN 1468-0262.

S. Bikhchandani, D. Hirshleifer, and I. Welch. A theory of fads, fashion, custom, and cultural change as informational cascades. *Journal of political Economy*, 100(5):992–1026, 1992. ISSN 0022-3808.

M.K. Brunnermeier and L.H. Pedersen. Market liquidity and funding liquidity. *Review of Financial Studies*, 2008. ISSN 0893-9454.

T.E. Copeland and D. Galai. Information effects on the bid-ask spread. *Journal of finance*, 38 (5):1457–1469, 1983. ISSN 0022-1082.

M. Coppejans, I. Domowitz, and A. Madhavan. Resiliency in an automated auction. *Unpublished working paper. Barclays Global Investors, ITG Inc*, 2003.

H. Degryse, F.D. Jong, M.V. Ravenswaaij, and G. Wuyts. Aggressive orders and the resiliency of a limit order market. *Review of Finance*, 9(2):201, 2005. ISSN 1572-3097.

D. Easley and M. O'Hara. Price, trade size, and information in securities markets* 1. *Journal of Financial economics*, 19(1):69–90, 1987. ISSN 0304-405X.

D. Easley and M. O'Hara. Time and the process of security price adjustment. *Journal of Finance*, 47(2):577–605, 1992. ISSN 0022-1082.

R.E.A. Farmer and J.T. Guo. Real business cycles and the animal spirits hypothesis. *Journal of Economic Theory*, 63(1):42–72, 1994. ISSN 0022-0531.

F.D. Foster and S. Viswanathan. Strategic trading when agents forecast the forecasts of others. *Journal of Finance*, 51(4):1437–1478, 1996. ISSN 0022-1082.

T. Foucault. Order flow composition and trading costs in a dynamic limit order market. *Journal of Financial markets*, 2(2):99–134, 1999. ISSN 1386-4181.

T. Foucault, O. Kadan, and E. Kandel. Limit order book as a market for liquidity. *Review of Financial Studies*, 18(4):1171, 2005. ISSN 0893-9454.

L.R. Glosten. Insider trading, liquidity, and the role of the monopolist specialist. *Journal of Business*, 62(2):211–235, 1989. ISSN 0021-9398.

L.R. Glosten. Is the electronic open limit order book inevitable? *Journal of Finance*, 49(4): 1127–1161, 1994. ISSN 0022-1082.

L.R. Glosten and P. Milgrom. Bid, ask and transaction prices in a market-maker market with heterogeneously informed traders. *Journal of Financial Economics*, 14:71–100, 1985.

R.L. Goettler, C.A. Parlour, and U. Rajan. Equilibrium in a dynamic limit order market. *The Journal of Finance*, 60(5):2149–2192, 2005. ISSN 1540-6261.

M.A. Goldstein and K.A. Kavajecz. Trading strategies during circuit breakers and extreme market movements. *Journal of Financial Markets*, 7(3):301–333, 2004. ISSN 1386-4181.

M. Golosov, G. Lorenzoni, and A. Tsyvinski. *Decentralized trading with private information.* Citeseer, 2009.

J. Greene and S. Smart. Liquidity provision and noise trading: evidence from the investment dartboard column. *The Journal of Finance*, 54(5):1885–1899, 1999. ISSN 1540-6261.

M.F. Hellwig. Rational expectations equilibrium with conditioning on past prices: A mean-variance example* 1. *Journal of Economic Theory*, 26(2):279–312, 1982. ISSN 0022-0531.

A.S. Kyle. Continuous auctions and insider trading. *Econometrica: Journal of the Econometric Society*, 53(6):1315–1335, 1985. ISSN 0012-9682.

L. Makowski and J.M. Ostroy. Perfect Competition and the Creativity of the Market. *Journal of Economic Literature*, 39(2):479–535, 2001. ISSN 0022-0515.

E. Maskin and J. Tirole. A theory of dynamic oligopoly, II: Price competition, kinked demand curves, and Edgeworth cycles. *Econometrica: Journal of the Econometric Society*, 56(3): 571–599, 1988. ISSN 0012-9682.

S. Morris and H.S. Shin. Liquidity black holes. *Review of Finance*, 8(1):1, 2004. ISSN 1572-3097.

M. Ostrovsky. Information aggregation in dynamic markets with strategic traders. In *Proceedings of the tenth ACM conference on Electronic commerce*, pages 253–254. ACM, 2009.

C.A. Parlour. Price dynamics in limit order markets. *Review of Financial Studies*, 11(4):789, 1998. ISSN 0893-9454.

I. Rosu. Multi-Stage Game Theory in Continuous Time. Technical report, Working Paper, University of Chicago, January, 2006.

I. Roşu. A dynamic model of the limit order book. *Review of Financial Studies*, 22(11):4601, 2009. ISSN 0893-9454.

L.K. Simon and M.B. Stinchcombe. Extensive form games in continuous time: Pure strategies. *Econometrica: Journal of the Econometric Society*, 57(5):1171–1214, 1989. ISSN 0012-9682.

D. Vayanos and P.O. WEILL. A Search-Based Theory of the On-the-Run Phenomenon. *The Journal of Finance*, 63(3):1361–1398, 2008. ISSN 1540-6261.

X. Vives. Short-term investment and the informational efficiency of the market. *Review of Financial Studies*, 8(1):125, 1995. ISSN 0893-9454.

X. Vives. *Information and learning in markets.* Citeseer, 2008.

Appendix

Without loss of generality consider the ask side of the book. The profit function is symmetric so it is easy to rewrite proofs for the bid side.

Most of the proofs use the following decomposition of the profit function,

$$
\begin{aligned}
\pi_q(p) &= \int_{-\infty}^{\infty} (p - \mathbb{E}[\nu \mid \theta, b]) \, \mathbb{I}(p, q) \, d\Phi(\theta \mid \nu) \\
&= \int_{\alpha q + \frac{\alpha}{2} + p}^{\infty} p\phi(\theta \mid \nu) - \nu_g b\phi(\theta \mid \nu_g) - \nu_b(1-b)\phi(\theta \mid \nu_b) \, d\theta \\
&= b(p - \nu_g)(1 - \Phi_g(p, q)) + (1 - b)(p - \nu_b)(1 - \Phi_b(p, q))
\end{aligned}
\tag{24}
$$

where $\Phi_g(p, q) = \Phi(\alpha q + \frac{a}{2} + p \mid \nu_g)$ and $\Phi_b(p, q) = \Phi(\alpha q - \frac{a}{2} + p \mid \nu_b)$.

Proposition 1: Single-peakedness and ordering

Single-peakedness

The function $\pi_q(p)$ is single-peaked for $\sigma_l \geq \frac{\nu_g + \nu_b}{2}$.

Proof: For notational convenience, fix q and define

- $F(p) = b\Phi_g(p, q) + (1 - b)\Phi_b(p, q)$, and $f(p) = F'(p)$.

- $G(p) = b\nu_g\Phi_g(p, q) + (1 - b)\nu_b\Phi_b(p, q)$, and $g(p) = G'(p)$.

Then in order for $\pi_i'(p) > 0$,

$$
p < \frac{1 - F(p)}{f(p)} + \frac{g(p)}{f(p)}
\tag{25}
$$

and the opposite inequality holds for $\pi_i'(p) < 0$. Single-peakness means p crosses $\frac{1-F(p)}{f(p)} + \frac{g(p)}{f(p)}$ once and only once. *Sufficient*[12] conditions for single crossing are for (i) $\frac{1-F(p)}{f(p)}$ to be monotone decreasing and (ii) for $\frac{g(p)}{f(p)}$ not to increase at a rate faster than the linear rate of 1. By theorem 1 of Block, Li and Savits (2005) the first condition is satisfied for $\sigma \geq \frac{\nu_g + \nu_b}{2}$. The second condition is satisfied because $\arg\max_p \frac{\partial}{\partial p} \frac{g(p)}{f(p)} = (\nu_b + \nu_g)/2 - \log[(1-b)/b]\sigma^2/(\nu_b - \nu_g) = p^*$ and $\frac{\partial}{\partial p} \frac{g(p^*)}{f(p^*)} \leq 1$ again for the same range, $\sigma \geq \frac{\nu_g + \nu_b}{2}$.

Ordering

For $p \geq 0$, if $\pi_i(p) > 0$ then $\pi_i(p) > \pi_{i+1}(p)$. Graphically the $\pi_i(p)$ functions can be "stacked."

Proof: By the properties of the normal distribution,

- $1 - \Phi_g(p, q) > 1 - \Phi_g(p, q + 1)$ and $1 - \Phi_b(p, q) > 1 - \Phi_b(p, q + 1)$

- $\dfrac{1 - \Phi_b(p, q)}{1 - \Phi_g(p, q)} > \dfrac{1 - \Phi_b(p, q + 1)}{1 - \Phi_g(p, q + 1)}$

---

[12]Proposition 1 gives a sufficient condition bounding σ; in practice a lower σ can be used so long as p crosses $\frac{1-F(p)}{f(p)} + \frac{g(p)}{f(p)}$ once $\forall b$. If one were to use a lower σ it would be necessary to check this is true.

If $p \leq \nu_b$ then $\pi_i(p) \leq 0$, which contradicts the hypothesis of the proposition, so let $p > \nu_b$. If $p > \nu_g$ then the proposition is trivial, so let $p < \nu_g$. Last, if $\pi_{i+1}(p) < 0$ then the proposition is again trivial, so assume $\pi_{i+1}(p) \geq 0$.

$$
\begin{aligned}
\pi_i(p) &= b(p - \nu_g)(1 - \Phi_g(p, q)) + (1 - b)(p - \nu_b)(1 - \Phi_b(p, q)) \\
&= (1 - \Phi_g(p, q))\left(b(p - \nu_g) + (1 - b)(p - \nu_b)\frac{1 - \Phi_b(p, q)}{1 - \Phi_g(p, q)}\right) \\
&> (1 - \Phi_g(p|q+1))\left(b(p - \nu_g) + (1 - b)(p - \nu_b)\frac{1 - \Phi_b(p, q+1)}{1 - \Phi_g(p, q+1)}\right) \\
&= \pi_{i+1}(p)
\end{aligned}
\tag{26}
$$

The inequality (26) requires $b(p - \nu_g) + (1 - b)(p - \nu_b)\frac{1-\Phi_b(p|q+1)}{1-\Phi_g(p|q+1)} > 0$, which is assured because of the assumption $\pi_{i+1}(p) \geq 0$.

Proposition 2: Existence and uniqueness

Prices $\{p_q\}_1^A$ and $\{p_q\}_{-B}^1$ exist satisfying (10a), (10b), (10c) and (10d). Without loss of generality consider the ask side. To obtain (10a):

$$
\begin{aligned}
\lim_{q \to \infty} \pi_q(p) = \lim_{q \to \infty} &b(p - \nu_g)(1 - \Phi(\alpha q + {}^a\!/_2 + p) \mid \nu_g) \\
&+ (1 - b)(p - \nu_b)(1 - \Phi(\alpha q + {}^a\!/_2 + p) \mid \nu_b) = 0
\end{aligned}
\tag{27}
$$

The profit function converges pointwise to 0, so for any $c > 0$ there exists an integer $A \geq 0$ for which $c > \max_p \pi_A(p)$. (If this is true for $A = 0$ then there are no limit orders on the ask side.) Due to the properties of the CDF Φ the convergence is monotone, so there exists a unique first A for which any $a \geq A$ satisfies $\pi_a(p) < c$.

To obtain (10c), notice profit functions take negative values for $p < \nu_b$ and the profit is positive for $p > \nu_g$, so by the single-peaked property profits cross the x-axis only once. Take $q > q' > 1$. Again by the ordering property for any p there exists at least one p' for which $\pi_q(p) = \pi_{q'}(p')$. This p' must be unique or else either $pi_{q'}(p')$ must cross the x-axis again in order to preserve single-peakedness and ordering.

Proposition 3: Mixed strategies

Agents cannot do better than the equilibrium on the limit order book by playing mixed strategies. Suppose $N - 1$ agents make positive profits in the pure-strategy equilibrium. The model contains no mixed-strategy equilibrium in which N agents expect more profit than $\max_p \pi_N(p)$.

Proof: The proof is by contradiction. Suppose the N agents do play mixed strategies. Since the agents are identical look for symmetric mixed-strategy equilibria, *i.e.* agents draw prices from the same distribution.

For notational convenience, define

- $P(p, i)$, the probability that one trader's price p falls i^{th} in ascending order of value (its "position" is i) among the $N - 1$ other prices.

28

- $G(p) = P(p, N)$, the probability that p falls in the last position.

- $H(p, i) = P(p, i)/(1 - G(p))$ for $i \neq N$, the probability that p falls in one of the remaining positions conditional on not falling in the last position.

The function G is a CDF, and $\sum_{i=1}^{N-1} H(p, i) = 1$. The expected profit of posting p can be written

$$
\begin{aligned}
\pi(p) &= \sum_{i=1}^{N} \pi_i(p) P(p, i) \\
&= \sum_{i=1}^{N-1} \pi_i(p) H(p, i)(1 - G(p)) + \pi_N(p) G(p) \\
&= \tilde{\pi}(p)(1 - G(p)) + \pi_N(p) G(p)
\end{aligned}
\tag{28}
$$

where $\tilde{\pi}(p)$ is the convex combination of $\pi_i(p)$ for $i \neq N$ using $H(p, i)$ as the weights.

In a mixed-strategy equilibrium each agent must be indifferent to posting any price in the support of the mixing distribution. That is, the expected profit $\pi(p)$ must be equal over all p in the support. Label the expected profit of playing the mixed strategy $\overline{\pi}$. The proof shows a contradiction arises if $\overline{\pi}$ is greater than $\max_p \pi_N(p)$. Label the maximum $\pi_N^* = \max_p \pi_N(p)$.

The profits in excess of the max are

$$
\pi(p) - \pi_N^* = \overline{\pi} - \pi_N^*
$$

Substituting in (28) rearrange to solve for $G(p)$:

$$
G(p) = \frac{1}{1 - \frac{\pi_N(p) - \pi_N^*}{\tilde{\pi}(p) - \pi_N^*}} - \frac{\overline{\pi} - \pi_N^*}{\tilde{\pi}(p) - \pi_N(p)}
$$

Now exhaust the cases to show there cannot exist a support region for the mixed strategy $G(p)$ if $\overline{\pi} > \pi_N^*$.

1. Take any interval of p in which $\pi_N(p) < \tilde{\pi}(p) < \pi_N^*$. On such an interval

$$
\frac{1}{1 - \frac{\pi_N(p) - \pi_N^*}{\tilde{\pi}(p) - \pi_N^*}} > 1
$$

which threatens the ability of $G(p)$ to represent a probability. To acquire a $G(p) \in [0, 1]$ over such an interval it would be necessary to set $\overline{\pi} < \pi_N^*$.

2. Take any interval of p in which $\tilde{\pi}(p) < \pi_N(p) < \pi_N^*$. On such an interval

$$
\frac{1}{1 - \frac{\pi_N(p) - \pi_N^*}{\tilde{\pi}(p) - \pi_N^*}} < 0
$$

which again threatens the ability of $G(p)$ to represent a probability. Again it would be necessary to set $\overline{\pi} < \pi_N^*$ to acquire a $G(p) \in [0, 1]$ (because the second term is now positive).

3. Last take any interval of p in which $\pi_N^* \leq \tilde{\pi}(p)$. On such an interval the first term in $G(p)$ is bounded above by 1, which seems promising, but if in addition $\overline{\pi} > \pi_N^*$, the second term is negative and

$G(p) < 1 - \epsilon$ for some $\epsilon > 0$. Then $G(p)$ is not a CDF.

From the diagram it should be clear case 3 is the "interesting" case. The lesson of case 3 is that the mixing distribution cannot possibly assign enough probability to the other profit functions $\pi_i(p)$ for $i < N$ to satisfy $\overline{\pi} > \pi_N^*$.

So at most $\overline{\pi} = \pi_N^*$.

Proposition 4: Character of the bid and ask prices

Monotonicity

On a book with A limit orders on the ask side, the ask price p_1 satisfies the condition

$$\pi_1(p_1) = \max_p \pi_A(p) \equiv \pi_A(p_A^*)$$

Drawing inspiration from the implicit function theorem, take the first derivative with respect to beliefs

$$\frac{\partial \pi_1(p_1)}{\partial p}\frac{\partial p_1}{\partial b} + \frac{\partial \pi_1(p_1)}{\partial b} = \frac{\partial \pi_A(p_A^*)}{\partial p}\frac{\partial p_A^*}{\partial b} + \frac{\partial \pi_A(p_A^*)}{\partial b}$$

Noting $\partial \pi_A(p_A^*)/\partial p = 0$ at the maximum (or, use the envelope theorem) and rearranging,

$$\frac{\partial p_1}{\partial b} = \frac{\partial \pi_A(p_A^*)/\partial b - \partial \pi_1(p_1)/\partial b}{\partial \pi_1(p_1)/\partial p} \tag{29}$$

Using $\partial \pi_q(p)/\partial b = (p - \nu_g)(1 - \Phi_g(p, q)) - (p - \nu_b)(1 - \Phi_b(p, q))$, the numerator equals

$$(p_1 - \nu_b)(\Phi_b(p_A, A) - \Phi_b(p_1, 1)) - (p_1 - \nu_g)(\Phi_g(p_A, A) - \Phi(p_1, 1))$$
$$+ (p_A - p_1)(\Phi_b(p_A, A) - \Phi_g(p_A, A)) \tag{30}$$

a series of terms all greater than zero at least when $p_1 \leq \nu_g$. The denominator is always greater than zero because $p_1 < \arg\max_p \pi_1(p)$; the agent posting the ask price could profit by raising the price. (It does not because of the dynamic consequences.)

Symmetry when $b = 1/2$

When $b = 1/2$, it is easy to check (i) $\pi_q(p) = \pi_{-q}(\nu_g + \nu_b - p)$ and so (ii) $\max_p \pi_q(p) = \max_p \pi_{-q}(p)$. By (ii) the free entry condition is satisfied at $A = B$, and competitive profits on both sides of the book are equal, so the ask profit and the bid profit are equal, which by (i) means the bid and the ask must be symmetric.

Proposition 5: Fill rates

The first statement of the proposition should be evident from the definition of S. The probability is weakly increasing as the interval increases, so one minus the probability is decreasing.

Equations (23a)–(23d) obtain through proposition 3 and the symmetry properties of the normal distribution. Begin with equation (23b). By proposition 3 the prices straddle $\frac{\nu_g + \nu_b}{2}$ at $b = 1/2$. Because

the normal distribution is symmetric,

$$\Phi_g(p_1, 1) = 1 - \Phi_b(\nu_g + \nu_b - p_1, 1)$$

meaning at a spread that straddles the mean value the market agent who knows the asset is bad is as likely to trade as a market agent who knows the asset is good,

$$\Pr(\theta \in [p_1 \ p_{-1}] \pm \{\alpha\} \mid \nu_g, x) = \Pr(\theta \in [p_1 \ p_{-1}] \pm \{\alpha\} \mid \nu_b, x)$$
$$\Phi_g(p_1 + \alpha) - \Phi_g(p_{-1} - \alpha) = \Phi_b(p_{-1} + \alpha) - \Phi_b(p_{-1} - \alpha) \tag{31}$$

which establishes equation (23b). To reach the others, notice for any $\epsilon > 0$,

$$\Pr(\theta \in [p_1 + \epsilon \ p_{-1} + \epsilon] \pm \{\alpha\} \mid \nu_g, x) > \Pr(\theta \in [p_1 + \epsilon \ p_{-1} + \epsilon] \pm \{\alpha\} \mid \nu_b, x)$$
$$\Rightarrow r([p_1 + \epsilon \ p_{-1} + \epsilon], \nu_g, x) < r([p_1 + \epsilon \ p_{-1} + \epsilon], \nu_b) \tag{32}$$

so if the prices increase from the bid and ask at $b = 1/2$, market agents are less likely to trade a good asset than a bad asset. By the monotonicity of prices this is exactly what happens as b increases. The reverse also holds.

The last fill rate relation obtains because

$$\Pr(\theta \in [p_1 \ p_{-1}] \pm \{\alpha(s)\} \mid \nu_g, x_s) > \Pr(\theta \in [p_1 \ p_{-1}] \pm \{\alpha(v)\} \mid \nu_b, x_v) \tag{33}$$

Proposition 6: Bayesian update increases in q

The Bayesian update is always increasing in q if and only if the hazard ratio ordering holds for any $\theta_2 > \theta_1$:

$$\frac{\phi_g(\theta_2)}{\Phi_g(\theta_2) - \Phi_g(\theta_1)} > \frac{\phi_b(\theta_2)}{\Phi_b(\theta_2) - \Phi_b(\theta_1)}$$

It holds by the properties of the normal distribution.

Proposition 7: Continuous-time Bayesian update during silence

The first statement of the proposition contains its proof. The second statement follows because a linear combination of any bounded continuous function that crosses the x-axis at least once at $1/2$ and a sufficiently steep line crossing the x-axis at $1/2$ has only one root, $1/2$. The line may have to be very steep compared to the bounded continuous function.

The information decay is a linear function of b that crosses the x-axis at $b = 1/2$. The information gain is another function that crosses the x-axis at $b = 1/2$ among other places. The information gain is bounded by $b(1 - b)r_g(b)$ on the top and $-b(1 - b)r_b(b)$ on the bottom. The fill rates are bounded by μ, so the information gain has an absolute bound. It is a numerical matter to find an information decay line of sufficient steepness (the slope is λ).

Proposition 8: Efficient prices when x is public information

A mechanism designer in incentive compatibility in a quasilinear setting interested would set a discrete pricing scheme to maximize surplus: $\max_{q(\theta)\in\mathbb{N}} \theta q - \alpha(x)q^2 - q\cdot\nu$. Due to the delivery problem (Makowski and Ostroy 2001) there are no prices that are both surplus-maximizing and that can elicit truthful revelation of ν. The natural "second-best" option used by Biais, Martimort and Rochet (2000) is to interpret the market agent as knowing not ν but a noisy signal on ν, namely the demand parameter θ. The agent can reveal its demand incentive-compatibly, so the designer identifies $p(q)$ solving

$$\max_{\substack{q(\theta)\in\mathbb{N} \\ p(q)}} U_\theta(q(\theta),p) - q(\theta)\cdot\mathbb{E}[\nu\mid\theta,b] \tag{34}$$

So long as $\partial\mathbb{E}/\partial\theta < 1$ a discrete schedule of prices exists and is unique:

$$p^e(q) = \mathbb{E}\big[\nu\mid\min\theta^{-1}(p,q),b\big] \tag{35}$$

which equals the expected value of the asset given the type wants to buy q at the efficient prices. Due to the discreteness in q the types who want to buy q are not unique; they lie in a series of almost-disjoint intervals $\theta^{-1}(p,q)$, and prices make the threshold types $\min_\theta \theta^{-1}(p,q)$ indifferent.

The prices are efficient because they maximize surplus. They also result in (second-best) information revelation. Typically the mechanism designer implements a schedule of prices that is zero-profit.

Proposition 9: Four-state model update of b_v during silence

Because the updating regime is not a tractable one, the proof will study its limiting behavior. It is necessary to define state-specific beliefs and fill rates:

Table A: New state-specific variables

state	1 good stable	2 bad stable	3 good volatile	4 bad volatile
belief	b_1	b_2	b_3	b_4
fill rate	r_1	r_2	r_3	r_4

The fill rates are contingent on the equilibrium spread at the book's beliefs, but for notational ease I drop the argument. The argument will not be important because the proof will substitute for the rates their upper bound μ. Also, the b_4 variable equals $1 - b_1 - b_2 - b_3$ but I write it separately again for notational ease.

In addition, it will be notationally convenient to define conditional prior variables:

Table B: New state-conditional beliefs

$b_{i\mid s}$	the belief the state is i conditional on the economy being stable
$b_{i\mid v}$	the belief the state is i conditional on the economy being volatile

The differential equations governing b_1 and b_3 during silence are

$$\frac{\partial b_1}{\partial t} = b_1 \cdot \sum(b_i r_i) - b_1 r_1 + b_2 \lambda_s + b_3 \eta_v - b_1(\lambda_s + \eta_s)$$

$$\frac{\partial b_3}{\partial t} = \underbrace{b_3 \cdot \sum(b_i r_i) - b_3 r_3}_{\text{gain}} + \underbrace{b_1 \eta_s + b_4 \lambda_v - b_3(\lambda_v + \eta_v)}_{\text{decay}}$$

Adding them, rearranging and strategically renotating produces

$$\frac{\partial b_\nu}{\partial t} = \frac{\partial b_1 + \partial b_3}{\partial t} = b_x[b_x b_{1|s}(b_{1|s}r_1 + b_{2|s}r_2) - b_{1|s}r_1 + \lambda_s(b_{2|s} - b_{1|s})] \tag{36a}$$

$$+ (1 - b_x)[b_{3|v}(1 - b_x)(b_{3|v}r_3 + b_{4|v}r_4) - b_{3|v}r_3 + \lambda_v(b_{4|v} - b_{3|v})] \tag{36b}$$

$$+ b_x(1 - b_x)(b_{1|s}(b_{3|v}r_3 + b_{4|v}r_3) \tag{36c}$$

$$+ (1 - b_x)b_x(b_{3|v}(b_{1|s}r_1 + b_{2|s}r_2) \tag{36d}$$

The update is not very tractable so one would like to avoid working with it except at limiting values. When $b_x = 1$ the system reduces to the differential equation governing $b(t)$ during silence in the stable economy (14), an expression that is contained in (36a). Conversely when $b_x = 0$ it reduces to the belief update during silence in the volatile economy, an expession similarly couched in (36b).

A natural strategy is to argue that as $b_x \to 0$ the roots of $\partial b_\nu/\partial t$ collapse to that of the volatile economy belief update during silence. It can be formalized using bounds that vanish as $b_x \to 0$. Here are bounds for equations (36a), (36c) and (36d):

$$-b_x(\mu/4 + \lambda_s) < b_x[b_x b_{1|s}(b_{1|s}r_1 + b_{2|s}r_2) - b_{1|s}r_1 + \lambda_s(b_{2|s} - b_{1|s})] < b_x(\mu/4 + \lambda_s) \tag{37a}$$

$$0 \le b_x(1 - b_x)(b_{1|s}(b_{3|v}r_3 + b_{4|v}r_3)) < b_x\mu \tag{37b}$$

$$0 \le (1 - b_x)b_x(b_{3|v}(b_{1|s}r_1 + b_{2|s}r_2)) < b_x\mu \tag{37c}$$

The conditional belief $b_{3|v}$ can be bounded,

$$\frac{b_\nu - b_x}{1 - b_x} < b_{3|v} < \frac{b_\nu}{1 - b_x}$$

The bounds on $b_{3|v}$ can be used to bound (36b):

$$(b_\nu - b_x)((b_\nu - b_x)r_3 + (1 - b_\nu - b_x)r_4)) - b_\nu r_3 + \lambda_v \frac{1 - 2b_\nu - 2b_x}{1 - b_x} \tag{38a}$$

$$< \ (36b) \ <$$

$$b_\nu(b_\nu r_3 + (1 - b_\nu)r_4) - (b_\nu - b_x)r_3 + \lambda_v \frac{1 - 2b_\nu - b_x}{1 - b_x} \tag{38b}$$

Combining (37) and (38), the lower and upper bounds on $\partial b_\nu/\partial t$ are

$$(b_\nu - b_x)((b_\nu - b_x)r_3 + (1 - b_\nu - b_x)r_4)) - b_\nu r_3 + \lambda_v \frac{1 - 2b_\nu - 2b_x}{1 - b_x} - b_x(\mu/4 + \lambda_s) \tag{39a}$$

$$b_\nu(b_\nu r_3 + (1 - b_\nu)r_4) - (b_\nu - b_x)r_3 + \lambda_v \frac{1 - 2b_\nu - b_x}{1 - b_x} + b_x(9\mu/4 + \lambda_s) \tag{39b}$$

The boundaries both contain the belief update during silence in the volatile economy except for the addition or subtraction of constant terms that vanish as $b_x \to 0$. As before $r_3 < r_4$ when $b_\nu > 1/2$ and $r_3 > r_4$ when $b_\nu < 1/2$. By the same argument as in proposition 7, equation (39a) has at least one stable root but shifted down to $b = 1/2 - \delta_1$ for some $\delta_1 > 0$ due to the subtraction of some constants. Equation (39b) also has at least one stable root but shifted up to $b_\nu = 1/2 - \delta_2$ for some $\delta_2 > 0$ due to the addition of some constants. If there are multiple stable roots choose δ_1 and δ_2 that correspond to the stable roots furthest away from $b_\nu = 1/2$. Then choose $\delta = \max\{\delta_1, \delta_2\}$. The upper and lower bounds are continuous because the fill rates are continuous functions of the bid and ask prices, which are continuous functions of beliefs. As $b_x \to 0$, the magnitude of the constant terms shifting the Bayesian update is decreasing, so δ is weakly decreasing. (Notice eventually the upper and lower bounds sandwich the Bayesian update during silence in the volatile economy.) Since the upper and lower bounds are continuous functions, the Bayesian update $\partial b_\nu / \partial t$ has all roots within $b_\nu \in [1/2 - \delta \; 1/2 + \delta]$ and at least one is stable.