Information Acquisition in a Limit Order Market^{*}

Ronald L. Goettler^{\dagger} Christine A. Parlour^{\ddagger}

Uday Rajan[§]

May 30, 2006

Abstract

We numerically solve a model with endogenous information acquisition in a limit order market.. The asset has a common value; in addition, each trader has a private value for it. Traders randomly arrive at the market, after choosing whether to purchase information about the common value. They may either post prices or accept posted prices. If a trader's order has not executed, he randomly reenters the market, and may change his previous order. The model is thus a dynamic stochastic game with asymmetric information. Agents with the lowest intrinsic benefit from trade have the highest value for information and also tend to supply liquidity. Agents' incentives to acquire information and subsequent equilibrium trading behavior changes systematically with the underlying volatility of the asset. This has two asset pricing implications. First, in equilibrium, asymmetric information creates a "volatility multiplier" (prices are more volatile than the fundamental value of the asset) that is higher, the higher the fundamental volatility. This is due to a change in the composition of trader types in the market at any given time. Second, changes in the difference between any transaction price and the fundamental value are negatively correlated with the fundamental returns. This correlation is more negative, the more volatile the asset. This is due to a change in trading strategies. We conclude that estimates of any factor model systematically underestimate the true coefficient, and this effect is more severe the higher the exposure.

Keywords: limit order market, adverse selection, informational efficiency.

JEL Codes: G14, C63, C73, D82.

^{*}We have benefitted greatly from comments provided by Kerry Back, Ekkehart Boehmer, Bill Lovejoy, Ben Van Roy, and seminar participants at Aladdin (2004), CFS (Eltville), Michigan, Carnegie Mellon, Texas A&M, the 2004 Oxford Finance Summer Symposium, the 2005 Utah WFC, and the 2005 INFORMS conference. The current version of this paper is maintained at http://webuser.bus.umich.edu/urajan/research/inform.pdf.

[†]Tepper School of Business, Carnegie Mellon University. Tel: (412) 268-7058, E-mail: ronald.goettler@cmu.edu

[‡]Tepper School of Business, Carnegie Mellon University, and Haas School of Business, University of California, Berkeley. Tel: 510-643-9391, E-mail: parlour@haas.berkeley.edu

[§]Ross School of Business, University of Michigan, Ann Arbor. Tel: (734) 764-2310, E-mail: urajan@umich.edu

1 Introduction

Financial markets are inherently dynamic: traders choose when to trade. As a result, the set of potential traders in the market, and hence transaction prices, can vary across time. Explicitly modeling this is important because, in a dynamic market, the price of the asset at any instant is determined by the last transaction. More precisely, as there are two sides to any transaction, prices depend on the valuations of the most desperate buyer and seller at that time (the "marginal traders"), and how they split the difference in their valuations. If two assets differ in the distribution of agents who are ready to trade at any point in time, the characteristics of the price process generated by each asset, including price level and volatility, must differ. We develop a model of trading in a financial market and characterize how agents' trading strategies and therefore observed transaction prices vary systematically with the underlying cashflows. We relate volatility in the underlying cash flows to an underlying factor model and show how equilibrium trading behavior can affect estimation.

We analyse dynamic trading in a limit order market. Many financial markets around the world, including the Paris, Stockholm, Shanghai, Tokyo, and Toronto stock exchanges, are organized in this fashion. Aspects of a limit order book are also incorporated into markets such as Nasdaq and the NYSE. A pure limit order market has no intermediary or market maker. Instead, traders must trade directly with each other. The device that enables this is a limit order book, which contains prices and quantities of unfilled orders. A trader can either post an order to the limit order book (i.e., submit a limit order) or choose to trade against a previously posted order (i.e., submit a market order).

Limit order markets have at least three important frictions, including two sources of asymmetric information which could affect the distribution of traders. Typically, the benefits obtained from trade are privately known by traders. In addition, some traders may also have private information about the fundamental value (or common value) of the asset. Finally, the sequential arrival of traders is a friction since it creates local monopoly power and may also result in delayed trade. A further consequence is that the timing of trade is endogenous, since an agent may prefer to wait for a better offer. An analytic model that incorporates these frictions is prohibitively difficult at best, and intractable at worst. As a result, we use numerical methods to solve for equilibrium.

Our model has risk neutral agents who arrive randomly at the market for an asset that has both common and private components to its value. Agents have different information about the expected cash flows accruing to the owners of the asset (i.e., the common value). Each agent chooses either to buy or sell one share. If his order does not execute, he revisits the market and can revise his order. Thus, agents face a dynamic problem: the actions they take at any point in time incorporate the possibility of future reentry. In addition, there is adverse selection: prior to his first entry into the market, each agent chooses whether to buy information about the fundamental value of the asset. An informed agent views the current expected value of the cash flows on each entry, whereas an uninformed agent forms an estimate of this value based on market observables.

Conditional on agents' information acquisition strategies, we determine the equilibrium of the trading game. The trading game is a dynamic stochastic game with a multidimensional state space. Since it is analytically intractable, we numerically solve for an equilibrium.

We find that traders with a low intrinsic motive for trade are willing to pay the most for information. These traders also submit a significant proportion of limit orders to the market. That is, in our model, many limit orders come from informed agents. Further, asymmetric information leads to a change in the composition of agent types in the market at any point of time. On average, when the fundamental value is more volatile, there are more agents with extreme private valuations in the market. That is, the potential marginal traders for the next trade have higher valuations. This creates a "volatility multiplier": assets with high fundamental volatility also exhibit greater excess volatility (i.e. volatility in excess of that implied by the fundamental value).

The informational efficiency of the market unambiguously improves when there are a greater number of informed agents. There are two reasons for this. First agents with a low intrinsic motive for trade, who have the largest incentive to acquire information, tend to be liquidity suppliers. Hence, market observables beyond transaction prices, such as bid and ask quotes, are also informative about the common value. If there are many informed agents (so that the probability of trading against one is high) market observables lead to better estimates of the common value of the asset. By contrast, if there are few informed agents observables are less informative about the common value. Second, an uninformed agent who faces a price she does not like has the flexibility of either waiting for a better price or posting her own order. Thus, a trader sometimes prefers to incur the cost of waiting rather than trade at an unfavorable price. Competition among the informed agents, and the ability of all agents to undercut previously posted orders, ensures that prices eventually become close to the common value.

More broadly, the set of agents in the market changes across time, leading to transaction prices departing from the fundamental value of the asset. Prices are high (relative to fundamental value) when there are more agents with high positive private values in the market, and low (relative to fundamental value) when there are more agents with large negative private values.¹ Thus, agents' response to market frictions naturally creates a time variance in transaction prices. This suggests that one explanation for time-varying expected returns or betas may be changes in the composition of the types of agents wishing to trade at any particular point of time.

We also document a negative correlation between returns to the fundamental asset value and changes in the "microstructure noise," — deviations of the transaction price from the fundamental value. These become larger as the asset value is more volatile, potentially leading to a systematic misestimation of any factor model. Further, this suggests that other variables (such as transaction cost measures) may explain returns in part because of this correlation. The intuition for this effect is simple. Consider a change in an asset's fundamental value: informed agents will execute against any mispriced limit orders. This is the "picking off" or adverse selection risk inherent in posting a limit order. There is thus a difference between the transaction price and the true value. In equilibrium, assets which are more volatile induce traders to place more conservative orders. Or, higher adverse selection risk and higher asset volatility lead to smaller changes in microstructure noise. Any econometric specification that ignores this affect systematically underestimates the true factor loading.

A distinctive feature of our approach is that our traders are risk neutral, and the potential gains from trade are fixed across different information acquisition regimes. We model the market in this way because we are interested in trade on information about each firm's cash flows. Such information should not affect the gains to trade if agents already hold well-diversified portfolios; that is agents are locally risk neutral. Adjustments to an agent's holdings of a particular asset are then primarily motivated by liquidity needs, as opposed to risk-sharing ones. Thus, the size of the potential gains to trade does not depend on the degree of adverse selection.

Asset pricing theory typically assumes competitive financial markets. In reality, most financial markets are imperfectly competitive. Blume and Easley (1990) show that generically there is no game with the competitive rational expectations equilibrium as an outcome. Perry and Reny (2004) provide a model of a double auction that, under stringent regularity conditions, converges to the fully-revealing rational expectations equilibrium as the number of traders becomes large. The limit order market we model is a continuous time variant of a double auction with discrete prices.

In contrast, in much of the rational expectations literature, the gains to trade are tied to adverse selection. Hirshleifer (1971) observes that, in an exchange economy with risk-averse

¹Admati and Pfleiderer (1988) consider a model with market makers, and show that traders benefit from co-ordinating the time at which they trade.

traders, information has no social value: if all agents are informed, risk-sharing opportunities are eliminated and the market breaks down.² Grossman and Stiglitz (1980) note that if costly information is immediately impounded into price, agents should not acquire it. Clearly, the argument depends on how agents profit from their information, so the results are specific to a price formation mechanism.³ Thus, a model with endogenous information acquisition should include stylized representations of the most important trading frictions. Previous general equilibrium work with endogenous information acquisition considers noise in the aggregate asset supply (e.g., Verrecchia, 1982; Admati and Pfleiderer, 1987) or "noise" traders with exogenous demands (Barlevy and Veronesi, 2001) to ensure that prices are only partially revealing. Our market is inherently dynamic, with the common value of the asset changing over time. In a temporal sense, informed traders are local monopolists. Hence, there can at best be partial revelation.

The canonical strategic rational expectations model is Kyle (1985), which has an informed trader and many "noise" or "liquidity" traders. An equilibrium condition in this model is that the market maker's price is the expected value of the asset conditional on all public information, including the direction and magnitude of contemporaneous order flow.⁴ As all trades are consummated at the market maker's quoted prices, there is no distinction between quotes and transaction prices.

Spiegel and Subrahmanyam (1992) demonstrate that introducing rational uninformed traders with risk-sharing motives into the Kyle framework generates different comparative statics. In particular, the welfare of liquidity traders monotonically decreases in the number of informed traders. This is because risk-averse liquidity traders reduce the amount they trade in the presence of adverse selection. In contrast, we find that aggregate welfare is almost invariant to changes in the degree of asymmetric information. Our agents can postpone trade and return to the market, and thus do not have to accept the prices offered at any particular time.

Endogenous information acquisition is examined by Mendelson and Tunca (2004) in a model with strategic risk-averse noise traders. Since market prices are partially revealing, an informed insider reduces uncertainty. However, this also reduces the gains to risk sharing. The insider takes into account the effect of his actions on uninformed traders, and may

 $^{^{2}}$ Hakansson, Kunkel, and Ohlson (1982) demonstrate that information can have social value if the market is not allocationally efficient. Bernardo and Judd (1997) find that information acquisition reduces welfare both because uncertainty is resolved before trade (the Hirshleifer effect) and because rent-seeking trades by informed agents reduce optimal risk-sharing.

³For example, Jackson (1991) demonstrates that the price-taking assumption is critical in order to sustain the Grossman–Stiglitz paradox.

⁴Taub, Bernhardt, and Seiler (2004) consider the case of multiple informed agents and repeated information shocks and find that the properties of Kyle (1985) hold in a more complex model.

choose not to acquire information (even at a zero cost). In our model, traders are risk neutral, and the gains to trade are invariant to information acquisition. Hence, at a zero cost all traders will acquire information. When there are multiple informed traders in a Kyle model, Foster and Viswanathan (1996) show that the correlation between informed traders' signals is important in determining the speed of information revelation. If this correlation is low enough, the equilibrium is characterized by waiting in later periods, so less information is revealed than in the monopolist case. As a result, the market may become illiquid towards the end of the overall trading period. Their numerical results are confirmed by Back, Cao, and Willard (2000) in a continuous time model. In our model, informed traders know the common value on each entry into the market. However, the common value changes over time. Hence, though signals are imperfectly correlated, agents have an incentive to act on information before it becomes stale due to an exogenous change in the common value, and the market remains active.

Our paper links the literature on information acquisition to that on dynamic limit order markets.⁵ The latter includes Rosu (2004), who presents a continuous time private value model of a limit order market with continuous prices and instantaneous punishment strategies. Foucault, Kadan, and Kandel (2004) characterize equilibrium in a dynamic limit order book with private values and differences in time preferences. Goettler, Parlour, and Rajan (2004) numerically solve an infinite horizon model of a limit order market with symmetric information and exogenous cancellation. None of these models consider information acquisition, or allow agents to differ in what they observe upon entering (or reentering) the market. Back and Baruch (2006) consider a continuous time model with asymmetric information, and demonstrate that, in the absence of frictions, market design is irrelevant (every equilibrium in a limit order market can be sustained as an equilibrium on a floor exchange with competitive market-makers, and vice versa), lending some justifiction to our modelling frictions in a limit order market.

Our work adds to the computational economics literature by showing that the stochastic approximation algorithm of Pakes and McGuire (2001) extends naturally to models of asymmetric information, where agents have private state variables. Further, we introduce the notion of "trembles" to the algorithm to ensure accurate beliefs for actions off the equilibrium path as required by perfection.

In Section 2 we outline the model, and provide an overview of the algorithm. Specific details of the algorithm appear in Appendix A. As we are interested in the relationship between asset volatility and trading strategies, we first analyze agents' order submission

⁵Parlour (1998) characterizes a limit order market with no common value. Foucault (1999) models a common value, but truncates the book to one share.

strategies across different information acquisition structures in Section 3. We consider aggregate investor surplus across different information regimes in the limit order market in Section 5, and examine informational efficiency of the market in Section 5.1. In Section ??, we provide some comparative statics on informational and allocative efficiency. In Section 6 we use the computed payoffs from the trading game to solve for the equilibrium information structures. The asset pricing implications of our results are developed in Section 7. Section 8 concludes.

2 Model

We model endogenous information acquisition in a dynamic limit order market for a single asset. We computationally solve for equilibrium in the trading game and then determine endogenous information acquisition in Section 6. In philosophy the trading game is similar to that in Goettler, Parlour, and Rajan (2004). There is a common value to the asset, v, and each trader also has a private benefit to trade, α . On entry into the market, a trader observes market conditions and decides whether to submit a buy or a sell order. The equilibrium cannot be determined analytically in closed form, so we solve for it numerically.

This model of the trading game has two important differences with Goettler, Parlour, and Rajan (2004). First, there may be asymmetric information, so that a trader in the market may have inferior information about the fundamental value of the asset, compared to previous traders. Second, traders who have submitted limit orders are allowed to reenter the market and change or cancel their order. Each trader, therefore, plays a dynamically optimal strategy on each entry, and cancellations are endogenous.

Time is continuous, although events happen after discrete time intervals. There is a discrete (but infinite) set of prices, $\mathcal{P} = \{p^i\}_{i=-\infty}^{\infty}$, at which traders may submit orders. The distance between any two consecutive prices is normalized to one and we refer to it as the "tick size." Associated with each price $p^i \in \mathcal{P}$ at time t is a backlog of outstanding orders to buy or sell the asset, ℓ_t^i . This backlog represents the depth at price p^i . We adopt the convention that a positive quantity denotes buy orders and a negative quantity sell orders. The limit order book at time t, $L_t = \{\ell_t^i\}_{i=-\infty}^{\infty}$, is the vector of outstanding orders. We consider an open limit order book, with no hidden orders.⁶ Given a limit order book L, the bid price or quote is $B(L) = \max\{i \mid \ell^i > 0\}$, the highest price at which there is a limit buy order on the book, and the ask price or quote is $A(L) = \min\{i \mid \ell^i < 0\}$, the lowest price at which there is a limit sell order on the book. If the corresponding set of prices is empty, define $B(L) = -\infty$ and $A(L) = \infty$.

 $^{^{6}}$ On some exchanges, limit orders may be "hidden" (i.e., not revealed to other traders), by choice of the order submitter.

New traders arrive at the market according to a Poisson process with parameter λ . Each trader has a type denoted by $\theta = \{\rho, \alpha\}$, where, ρ is a discount rate and α a private value for the asset. The payoff a trader earns as a result of trading is discounted back to his first arrival time in the market at the rate ρ . The cost of delaying trade could include an opportunity cost (e.g., if a trader is executing a trading strategy across different assets and must delay trades in other assets) and a cost to monitoring the market before execution, rather than the time value of money. Traders in some financial markets appear to care about differences in seconds in the time to execution; the discount rate captures this desire to trade early.⁷

The private value α represents private benefits of trade, accruing to a trader as a result of liquidity shocks or private hedging needs. Its presence implies potential gains to trade among agents. Let F_{α} denote the distribution of α . The private value α is independently drawn across traders.

In addition to a private value for each trader, the asset at any instant t has a common value, denoted v_t . The common value is the expectation of the present value of future cash flows on the stock, and evolves as a random walk. Innovations in the common value occur according to a Poisson process with parameter μ . If an innovation occurs, the common value increases or decreases by k ticks, each with probability $\frac{1}{2}$. Changes in the common value reflect new information about the firm. In Section 7 we show how v can be interpreted in a factor model.

On his first entry to the market, an agent may choose to buy information by paying a cost $c \ge 0$. Incurring this cost gives an agent access to a service that reports the current value of v on this and each subsequent entry. The timing of the acquisition decision captures the idea that agents research an asset before deciding to participate in the market.

Since all investors have a chance to acquire the information, it is publicly available: for example, information reported in financial statements, SEC filings, or analyst reports, or prices of related assets such as options.⁸ Our acquisition cost can be interpreted as an explicit cost such as subscribing to a news service, or an opportunity cost in terms of time required to gather and process the information. In equilibrium, traders in our model consider this cost when they choose whether to acquire information.

Uninformed agents view v with a time lag, Δ_t , measured in units of time. That is, an

⁷The execution speeds for market orders for stocks on the NYSE and Nasdaq are routinely mentioned in the trade literature. A Google search for "execution speed nyse nasdaq" brings up pages on the NYSE and Nasdaq (each claiming better execution over the other, albeit for different order sizes), congressional testimony, and pages at various brokerage houses.

⁸In as much as privately informed agents trade in such related assets, whose prices are publicly observed, our model can also be interpreted as a model of how insider information may be incorporated into the price of an asset.

uninformed agent in the market at time t knows $v_{t-\Delta_t}$, whereas an informed agent in the market at time t knows the current value v_t . In addition, all agents observe the history of the game. Let t denote a time at which an agent has entered the market, and let h_t denote the history up to time t, before the agent takes an action. The history includes all actions in the game until time t as well as changes in the common value until time t (for informed traders) and time $(t - \Delta_t)$ (for uninformed traders).

For all agents, the limit order book L_t provides information about current trading opportunities. In addition, uninformed agents use their information set to update their expectation about the common value v. The history offers strategic information to informed agents as well: using the information available to an uninformed agent allows informed agents to better predict the actions of uninformed agents, and thus earn a higher payoff themselves.

Since traders choose whether to acquire information prior to observing market conditions, we can think of these strategies as being chosen at time 0, before the start of the trading game. Our model is therefore equivalent to a two-stage game. At the first stage all agents choose whether to acquire information. At the second stage, with information acquisition strategies held fixed, the continuation "trading game" is played. We consider symmetric equilibria (i.e., traders of a particular type all play the same strategy at each stage).

In the remainder of this section, we discuss the trading game in greater detail, holding information acquisition decisions fixed.

2.1 Continuation Trading Game

Each trader is allowed to trade at most one share of the asset: however, he may choose to buy or sell that share. A trader who previously entered the market, but whose share has not yet executed, reenters the market at some random time. On any particular entry a trader may choose to submit no order. Traders are potentially active until their order executes, at which time they leave the market forever. Thus, at any point of time, there will be a random number of agents who have not yet traded. Each unexecuted trader reenters the market according to a Poisson process with parameter λ_r . Reentry, therefore, is not instantaneous, and represents a friction agents must take into account when submitting an order. The reentry friction captures the idea that agents do not continuously monitor the market and also provides a way to determine a priority of order arrival among several agents who all wish to trade at the same time. The reentry times are independent across agents. At any particular instant there is at most one agent (either a new or returning trader) who chooses an action. When he is in the market at time t, a trader takes an action a = (p, q, x), where p denotes the price at which he submits an order, $q \ge 0$ the priority of his order among all orders at price p, and

$$x = \begin{cases} 1 & \text{if a buy order is submitted,} \\ -1 & \text{if a sell order is submitted,} \\ 0 & \text{if no order is submitted.} \end{cases}$$
(1)

If x = 0, the values of p and q are irrelevant to payoff.⁹

If there is an existing order at price p on the other side of the market, a submitted order executes instantaneously and is called a marketable limit order, or a market order for short. A buy order at a price $p \ge A(L)$ automatically executes at A(L), and similarly with a sell order at $p \le B(L)$. For such an order, we set q = 0. Alternatively, if there is no order on the other side of the market at that price, the order joins the existing orders on the same side at that price.

Limit orders are executed according to time and price priority: that is, orders submitted earlier are further ahead in the queue. Buy orders at higher prices and sell orders at lower ones are accorded priority. Therefore, an order executes if no other orders have priority, and a trader arrives who is willing to be a counterparty.

Upon reentry, a trader may leave an existing order on the book or cancel it and submit a new order. The benefit of retaining the existing order is that he maintains his time priority (his place in the queue). The cost is that the asset value may have moved in a manner that affects the expected payoff from the order. For example, if he submitted a buy order and the asset value has since fallen, his order may now be priced too high. Conversely, if the asset value has since risen, his order may be at too low a price, and there may be little chance of it executing. Further, a trader may also find that the priority of a previous order has changed by the time he reenters the market.

Traders are risk neutral and submit orders to maximize their expected discounted payoff. Utility is earned only if an order executes. For a particular trader $\theta = (\rho, \alpha)$, the instantaneous utility at time t is

$$u_t = \begin{cases} \alpha + v_t - p^i & \text{if he executes a buy order at price } p^i \text{ and time } t, \\ p^i - \alpha - v_t & \text{if he executes a sell order at price } p^i \text{ and time } t, \\ 0 & \text{if he does not execute an order at time } t. \end{cases}$$
(2)

The expected payoffs to different actions depend on a trader's information set. Each time a trader is in the market, he chooses a payoff-maximizing action. If a state is defined

⁹For a newly submitted order, L_t , p, and x determine q. However, q evolves over time for an order on the books, and may change before the trader reenters the market. It is used in determining the continuation payoff on reentry.

to be the history of events in the game observed by the trader, his decision is Markovian. At each entry into the market, he chooses an action that maximizes his value given the current state. If a trader chooses to not submit an order, we have $\tilde{x} = 0$, so $\pi = 0$. Since a trader is never forced to submit an order, and, in this model, there is no cost to reentering the market, the value of any state is bounded below by zero, given any previous order submitted by the trader. Hence, the overall value of any state is no lower than zero.

Formally, we have a Bayesian game. Traders have privately known utilities from trade (since a trader's α is unknown to other traders). As Maskin and Tirole (2001) point out, the proper solution concept here is Markov perfect Bayesian equilibrium, which requires traders to play optimal (state-dependent) strategies at every decision node (i.e., on each entry into the market), given their current posterior beliefs. In principle, these beliefs could relate to the current value of v (for an uninformed trader), the likelihood that any order currently in the book was submitted by an informed trader (who may reenter the market in the future), the private value α of each trader who had an order in the limit order book, or even over the number of traders who had entered the market and submitted no order, and may be submitting orders in the future.

In practice the numerical algorithm bypasses the issue of posterior beliefs by directly determining the expected utility of different actions in each state, allowing for a direct computation of the optimal state-dependent action. Numerically, we impose specific restrictions on the state space to make it computationally tractable.¹⁰ We solve for a stationary symmetric equilibrium given the state space. That is, each type of agent plays the same time-invariant strategy. The perfection requirement ensures that optimal actions are assigned to states that are off the equilibrium path of play and are hence never reached in equilibrium. Given information acquisition strategies, the equilibrium of the trading game appears to be computationally unique. In Section 6 we show that, despite this, there are cost ranges that lead to multiple equilibria in the information acquisition game.

2.2 Solving for Equilibrium in the Trading Game

Since the common value evolves as a random walk, the set of prices at which trade can feasibly occur is, in principle, unbounded (although it is finite in any finite simulation). However, given the payoff on execution in (2) above, a trader cares only about the relative price at which trade occurs (i.e., the price relative to the common value). Consider an informed trader who arrives at the market at time t when $v_t = 15$, and there is only one

¹⁰In particular, these restrictions may exclude some payoff-relevant variables, such as the exact time at which different events happened. In principle, the algorithm can handle any discretization of time; in practice, the size of the state space is limited by computational constraints.

order on the limit order book, a buy order at 16. Suppose there has been no change in the book over a substantial period of time (greater than Δ_t). Now consider an identical trader who arrives at t', with $v_{t'} = 21$, and only one order on the limit order book, a buy order at 22. Again, suppose there has been no change in the book over a substantial period of time. Our specification of payoffs implies that these two traders must take the same action relative to the respective common values they observe. That is, if submitting a sell order two ticks above v (at a price of 17) is optimal for the first trader, doing so (at a price of 23) must be optimal for the second trader.

Historical prices and lagged values of v can also be expressed in terms of the current common value for an informed trader, so the restriction in the previous paragraph to nonempty histories is for expositional purposes only. In the same manner, historical transaction prices and current books can be expressed relative to the last observed common value for an uninformed trader. This significantly reduces the size of the state space, to the point that the set of recurrent states in our simulations is finite (although still very large).¹¹

As discussed, we fix information acquisition strategies and solve for the equilibrium of the corresponding trading game. We simulate a market session and update beliefs (about the payoff of each action in each state) using the simulated outcomes until beliefs converge. Our algorithm follows Pakes and McGuire (2001) in that it uses simulation to asynchronously update values only for states in (or near) the recurrent class of states.¹² The advantage of this approach is two-fold. First, the updating of beliefs at a given state is computationally efficient, using the realized outcome from the simulation as a Monte Carlo estimator of the originating state's value. For example, when a limit order is executed in the simulated market, the value of the state at which the limit order was submitted is updated by averaging in the discounted payoff from this transaction with the previously held belief of the state's value. If a trader returns to the market before his limit order executes, the value of the state at which the order was submitted is instead updated with the perceived value of the trader's revised (or maintained) order, discounted by the elapsed time. As the simulation progresses, the states are "hit" repeatedly and their values, which are simple averages,

¹¹A recurrent class is a subset of states with the following properties: (i) regardless of the initial state, the system eventually enters the recurrent class; (ii) once entered, the probability of each state outside the recurrent class is zero; and (iii) each state in the recurrent class is visited infinitely often as t approaches infinity.

¹²Pakes and McGuire (2001) solve for equilibrium in a dynamic oligopoly, obtaining convergence in firms' value functions. Goettler, Parlour, and Rajan (2004) use a similar algorithm to solve a trading game in which traders take an action only when they initially enter the market. In that model all traders know the common value on entry and limit orders are cancelled according to an exogenous cancellation function. Forming posteriors about the private values of traders currently posting limit orders is unnecessary since these traders never return to revise their orders. The fixed point is therefore directly obtained for beliefs about execution probabilities and changes in the common value conditional on execution.

converge to their true means.

The second advantage of the stochastic algorithm is that values are only updated for states actually visited: that is, the fixed point is computed only for the recurrent class of states. Since the full state space of the trading game is huge, this feature of the algorithm is particularly important.

Perfection requires that agents' beliefs about payoffs to actions off the equilibrium path be correct, to rule out the possibility that incorrect beliefs at states outside the recurrent class may lead players to mistakenly avoid such states. Consider an extreme case in which all traders believe that limit orders never execute. Suppose the book is empty when the first trader enters. Given this belief, submitting no order is a best response for this trader and the book will be empty when the second trader arrives. This trader now faces a decision problem identical to the one faced by the first trader, and hence submits no order. Therefore, no orders are ever submitted in this market. Perfection rules out such situations.

Numerically, perfection requires the computation of payoffs to actions that are not chosen. To obtain these payoffs, we allow for trembles: On each entry, there is a small probability $\epsilon > 0$ that an agent takes a suboptimal action. To ensure that strategies along the equilibrium path are not affected, we take the following steps: (i) the probability of a tremble is small, declining to 0.5% as the algorithm converges (ii) updates to payoffs always use the optimal action in any state, even if the trader is randomly selected to tremble, and (iii) traders never tremble to market orders (such trembles would make limit orders at ridiculous prices more attractive).¹³

The algorithm is a natural extension of the stochastic approximation algorithm of Pakes and McGuire (2001) for complete information games. A transparent difference is that different agents have different state variables, and some payoff-relevant variables (such as α and possibly v) are privately known to agents. A substantive difference is the use of trembles to ensure perfection.¹⁴

Details of the algorithm appear in Appendix A, along with the convergence criteria that we use.

2.3 Implementation of the Numerical Algorithm

In principle, we would like agents to condition their strategies on the entire history of the game. In practice, of course, this is computationally infeasible. Hence, we impose

¹³As discussed in Appendix A, payoffs to market orders are updated regardless of actions taken, so trembling to market orders is unnecessary.

¹⁴While checking for convergence, Pakes and McGuire (2001) avoid pessimistic beliefs by directly computing the integral that defines each state's continuation value. This is not an option in our model because there are too many future states an agent could be in, given a current state and action.

some restrictions on the state space. We allow agents to condition their action on a set of variables related to the book and market. In addition, uninformed agents choose actions that can vary with $v_{t-\Delta_t}$, the fundamental value they observe, and informed agents can choose actions that vary with both v_t and $v_{t-\Delta_t}$. Therefore, the set of conditioning variables for an informed agent strictly includes all variables used by an uninformed agent.

The set of book-related variables used by agents consists of: (i) the current bid and ask prices, (B_t, A_t) ; (ii) the total depths at these prices, (ℓ_t^B, ℓ_t^A) ; (iii) the cumulative buy and sell depths in the book, $D_t^b = \sum_{i=0}^N \{\ell_i^t > 0\}$ and $D_t^s = \sum_{i=0}^N \{\ell_i^t < 0\}$; and (iv) the price closest to the bid (ask) at which buy (sell) depth is available, $\hat{B}_t(\hat{A}_t)$. Denote this 8-tuple of book-related conditioning variables as $\hat{L}_t = \{B_t, A_t, \ell_t^B, \ell_t^A, D_t^b, D_t^s, \hat{B}_t, \hat{A}_t\}$.

In addition, for a limit buy order at price \tilde{p} , agents also condition on $\{\ell_t^i\}_{i=\tilde{p}}^{B_t}$, the number of shares on the buy side of the book at prices least as aggressive as \tilde{p} . Similarly, for a limit sell order at price \tilde{p} , agents also condition on $\{\ell_t^i\}_{i=A_t}^{\tilde{p}}$, the number of shares on the sell side of the book at prices at least as aggressive as \tilde{p} . Though this implies different conditioning variables for different actions, in practice, less aggressive orders have a minimal effect on state values. Hence, including the omitted variables would not affect the equilibrium.

In addition, we allow agents to vary their actions based on the price \hat{p}_t of the most recent transaction, if this transaction occurred in the interval $[t - \Delta_t, t]$, and whether this transaction involved a market buy or sell. We investigated a model in which agents also observe the cumulative market buys and sells in the interval $[t - \Delta_t, t]$. The added conditioning variables are virtually ignored by traders in updating beliefs about v, and do not affect market outcomes.

2.4 Parameterization of the Trading Game

Time is continuous, with three types of Poisson events—new trader arrivals, returning traders, and changes in the common value. We normalize the mean time between new trader arrivals to 1, so that any time interval may also be interpreted in terms of the expected number of new trader arrivals in that period.

The other parameters that define our base case are as follows:

• The support of the discrete α distribution in ticks is $\{-8, -4, -0.1, 0.1, 4, 8\}$. The probabilities of -8, 8 are each 15%, while that of -4, 4 are 20% and the probabilities of -0.1, 0.1 are 15% each.

The traders with $\alpha \in \{-0.1, 0.1\}$ constitute traders who may be willing to buy or sell, depending on the state of the market when they arrive. We refer to these agents as "speculators," since they have a very low intrinsic motive to trade. The traders with $\alpha \in \{4, 8\}$ are likely to be buyers overall, and those with $\alpha \in \{-4, -8\}$ are likely to be sellers. These characterizations are borne out in our simulations.

Our private value distribution approximately corresponds to the findings of Hollifield et al. (2004), who estimate the distributions of private values for three stocks on the Vancouver exchange. Since they consider a world with symmetric information, their estimate of the gains to trade represents a likely upper bound (since some trading will occur for informational reasons). Our parameterization of F_{α} is based on their identification of three types of traders within these distributions. They find that, on average across the three stocks, 44% of traders have private values within 2.5% of the common value of the stock, 26% have values that differ from the common value by 2.5% to 5%, and 30% have values that differ from the common value by more than 5%. This corresponds approximately to the probabilities of our three kinds of traders.

In terms of ticks, on average across the three stocks, 2.5% of the common value translates to approximately 3.45 ticks, and 5% of the common value to approximately 6.9 ticks. Since 30% of our traders have private values within ± 0.1 ticks of the common value, we choose 4 ticks as the private value for our second kind of trader, and 8 ticks for our third kind.

Changes in v, the fundamental value of the asset, occur at times drawn from a Poisson distribution, so that the inter-arrival time of innovations is exponential. The expected time between changes in v is 8 units of real time. We consider two sets of models:
(i) Low volatility: whenever v changes, it increases or decreases by one tick, each with probability ¹/₂.

(ii) High volatility: whenever v changes, it increases or decreases by 8 ticks, each with probability $\frac{1}{2}$.

Our low volatility parameterization roughly corresponds to the findings of Hollifield et al. (2004). For the three stocks they consider, they report the average number of transactions during each ten-minute period and the volatility of the midpoint of the bid and ask quotes. Using these transaction frequencies, we infer the volatility of the midpoint of the quotes per transaction to be 0.20, 0.34, and 0.42 for the three stocks.

The midpoint of the bid and ask quotes is a rough proxy for the common value. In our model, new traders arrive at the rate of one per unit time. Since we consider stationary equilibria, and it takes two traders to make a transaction, transactions occur approximately every two units of time. Thus, we parameterize the volatility of the asset at 0.125 per new trader arrival, or 0.25 per transaction. This translates to an innovation occurring on average every 8 units of time.¹⁵

• We let $\Delta_t = 16$: this represents the average number of new trader arrivals that occur between the time an uninformed trader observes v and the time he takes an action.

Note that we use $\Delta_t = 16$ only for our base case. Since we have no empirical basis for this selection, we varied Δ_t in our simulations and found no significant differences for much larger values.

- ρ , the continuous discount rate, is the same for all agents and is set to 0.05. Recall that this is not the time value of money, but rather a preference parameter that reflects the cost of not executing immediately.¹⁶
- Agents reenter the market at an average rate of 4 units of time per reentry. Reentries are independent across traders and entries.

Numerically, the algorithm can trivially handle reentry rates that differ across types and trader information. In this paper, however, we are primarily interested in isolating the effects of differential information on market outcomes, so we keep the reentry rate the same across informed and uninformed traders. Conceptually, we think of reentry rates as depending on the cost of monitoring the market, with a lower monitoring cost implying a higher rate of reentry.

• Limit orders may be submitted at any feasible price that lies in a range between 2.5 ticks above and below an agent's expectation of v. For an informed trader, this expectation is just the current value of v. For traders who observe v with a lag Δ_t , this expectation is their best estimate given the lagged common value, the current book, and the observed market history.¹⁷

Market orders, of course, may be submitted at the current bid (market sells) or ask (market buys) regardless of an agent's expectation of v.

• Initially, we set the probability that an agent trembles to a suboptimal order at 0.05. As the algorithm converges, we reduce this probability, to a final value of 0.005.

 $^{^{15}}$ The variance of the innovation in an interval of time t is the expected number of innovations in the interval.

¹⁶We experimented with lower values of ρ , and found the results to be qualitatively similar. However, traders took longer to execute on average, and the state space was considerably larger.

¹⁷We simulated versions of the models in which limit orders could be submitted further away from the common value. Although orders were occasionally submitted at such ticks, these orders rarely executed, appearing to substitute for not submitting an order at all. There was no appreciable effect on market outcomes, either in the aggregate or for any particular type of trader.

3 Trading Behavior

We are interested in how agents' optimal trading behavior determines the relationship between the observed transaction price and the consensus value, v_t . We do this by comparing how each trader type changes actions when faced with assets with different fundamental volatility and how each trader type reacts to different asymmetric information regimes. For brevity, we report two information settings: all agents are informed and only speculators are informed.

There are three elements to an agent's payoff: his private value, the discount (or premium) that he obtains over the common value of the asset, and the length of time it takes him to execute. The overall discounted payoff to an agent with private value α who buys a share at price p when the common value is v and executes with a time delay (from first entry to the market) t is $e^{-\rho t} (\alpha - (p - v))$.

3.1 Time to execution and terms of trade

We first consider the time to execution for each type of agent. Changes in the execution time across agent types imply changes in the set of agents who are available to trade at any given point of time. In Table 1, we report for each trader type the average time from entry to the market until trade is consummated.

Volatility Regime	Information Regime	Absolute value of α		
		0.1	4	8
Low	All informed	19.64	2.09	0.59
	Speculators informed	17.22	2.19	0.62
High	All informed	8.87	3.19	1.07
	Speculators informed	12.34	5.44	2.33

Table 1: Average time to execution

Table 1 shows that the speculators take significantly longer to execute than any other traders: since they have a low intrinsic motive for trade, they are willing to wait longer for a better price. When the fundamental volatility is low, they execute fastest when they are the only ones informed. If they compete with another group of informed agents, it takes them longer to execute: competition among informed agents makes profitable trades more difficult to find. Although the extreme α agents always execute quickly, they take the longest when only the speculators are informed.

Conversely, when fundamental volatility is high, asymmetric information leads to an unambiguous increase in time to execution for each type of trader. Comparing the last two rows of Table 1, each trader type takes longer to execute when only speculators are informed, compared to the situation in which all agents are informed.

In Table 2, we report (p - v) for all executed buy orders (i.e., market and limit). Two comparisons are useful from the table. First, comparing across information regimes (i.e., comparing the first and second rows of numbers or the third and fourth rows), speculators experience an improvement in the terms of trade when only they are informed. When the volatility is low, buy orders of speculators with a positive α execute 1.21 ticks below the common value when only they are informed, as compared to 1.14 ticks below the common value when all agents are informed.¹⁸ The effects of adverse selection are exhibited by the increased cost (in terms of amount paid in excess of the common value) paid by agents with higher private values ($\alpha = 4$ or 8) in the case in which only the speculators are informed.

Volatility Regime	Information Regime	Value of α			
		-0.1	0.1	4	8
Low	All informed	-1.07	-1.14	0.10	0.48
	Speculators informed	-1.30	-1.21	0.21	0.57
High	All informed	-2.38	-3.12	-0.21	0.45
	Speculators informed	-4.78	-4.01	0.25	0.75

Table 2: Average of (price – fundamental value) for buy orders

Second, comparing across volatility regimes, speculators see a substantial improvement in their terms of trade when the fundamental volatility is high. For example, comparing the first row with the third row, speculators with a positive α improve their terms of trade by almost two ticks per share when the volatility is high. On the other hand, when only speculators are informed, extreme private value agents (i.e. those with $\alpha = 8$) on average execute at worse prices when the volatility is high.

A limit order market may broadly be seen as a dynamic bargaining mechanism. If agents do not like the posted prices, they can make their own offers. In the presence of adverse selection, the optimal posted orders are more conservative, since uninformed agents are concerned about being picked off by informed agents. The greater conservatism leads to greater time to execution.

¹⁸The notion that limit order submitters execute at favorable prices is consistent with the empirical work of Biais, Bisiere and Spatt (2003), who fail to reject the hypothesis that competing limit order submitters on Island (an electronic limit order book) make positive profits.

	Absolute value of α					
	0.1		4		8	
	% Lim	%Mkt	% Lim	%Mkt	% Lim	%Mkt
% of trader type	30		40		30	
Low volatility, all informed	68.1	13.7	23.8	40.4	8.1	45.8
Low volatility, speculators informed	66.6	12.7	25.1	41.0	8.3	46.4
High volatility, all informed	42.1	48.8	35.1	31.3	22.8	19.9
High volatility, speculators informed	53.8	39.2	26.9	36.1	19.3	24.7

Note: "% Lim" denotes the percentage of all limit orders submitted by a trader type. "% Mkt" denotes the percentage of all market orders submitted by a trader type.

Table 3: Proportion of limit and market orders submitted by each type

In Table 3, we report the proportion of limit and market orders that are submitted by each trader type.¹⁹

We find that speculators are the traders with the highest value for information, and also the agents who supply liquidity. That is, informed traders in our market tend to submit limit orders, a finding that corresponds to the experimental results of Bloomfield, O'Hara, and Saar (2004) and the empirical observations of Kaniel and Liu (2004) on the TORQ database of NYSE stocks. When the fundamental volatility is low, asymmetric information does not affect the distribution of orders across trader types. However, when the fundamental volatility is high, two effects are noticed. First, traders with more extreme private values tend to submit an increased proportion of limit orders, since the books are thinner on average. Second, asymmetric information leads speculators to increase their supply of liquidity. The book is slightly thicker on average, and, at the margin, uninformed traders substitute away from limit orders towards market orders.

4 Excess Volatility and Trading Pool Composition

In a frictionless market with all agents informed, all trades should execute at the fundamental value, v_t . Thus, the dispersion of prices around the fundamental value represents a measure of excess volatility generated by the frictions in the market. As Table 4 shows, excess volatility is higher when only speculators are informed, and higher yet when the volatility in the fundamental value is high. Thus, even with symmetric information, the market represents a volatility multiplier, exacerbating the effects of an uncertain fundamen-

¹⁹Note that agents who take longer to execute submit more orders than agents who trade more quickly, since the former are in the market more often on average.

Regime	Std. Dev.	Average number of		
	of $(p-v)$	traders in market		
		$\alpha = 0.1$	$\alpha = 4$	$\alpha = 8$
Low volatility, all informed	0.8	2.65	0.66	0.24
Low volatility, speculators informed	1.11	2.38	0.69	0.23
High volatility, all informed	2.59	1.42	1.17	0.69
High volatility, speculators informed	4.32	1.96	1.66	0.94

tal value. This effect is further enhanced when there is asymmetric information.

Table 4: Excess volatility and trading pool composition

Notes: (i) "Low" and "high" volatility under "Regime" refer to volatility in the fundamental asset value. (ii) "Average number of traders in market" is the average number of traders who have entered the market, but not traded, at any given instant. (iii) Since the model is symmetric, the number of traders in the market with negative α corresponds to the number with positive α .

As noted from the table, when the fundamental value is volatile, and there is asymmetric information (i.e., only speculators are informed), there is a shift in the distribution of traders in the market at any instant (or the "trading pool"), from low private values (i.e., closer to zero) to more extreme ones. Since, for any given trade, at least one of the two sides must constitute a trader who entered the market previously, this represents a shift in the marginal trader across regimes. When there is a greater number of extreme value traders, we expect to see prices that depart from the common value. Thus, excess volatility increases.

A further implication of the shift in the composition of the trader pool across regimes is that any attempt at aggregation across traders will lead to results that vary by market or asset. Suppose all else is equal, and consider two assets with different fundamental volatilities. On average across time, these assets will exhibit different marginal traders, and hence different "representative agents." This is surprising in our context, given that all agents are risk-neutral (so wealth effects are absent). Our results here complement those of Schlee (2001), who shows that, in an economy with some risk-averse agents, if there is a representative agent, then all agents are worse off when greater information is acquired. Thus, in such an economy, there is no incentive to acquire information. In our model, by comparison, agents have an incentive to acquire information, and aggregation fails.

5 Aggregate Allocative and Informational Efficiency

5.1 Informational Efficiency

As prices in markets can influence decisions regarding the allocation of real resources, the prices and outcomes in a particular market provide an important externality to agents outside the market as well. In this section, we consider the informational efficiency of the limit order market.

In our limit order market, information about the common value can be conveyed by several variables other than the transaction price, including bid and ask quotes and order depth in the book at various prices. Consider an uninformed agent in the market. This agent observes market conditions (in particular, the book and the price and direction of the most recent transaction) and forms an estimate about the common value. One measure of the informational efficiency of the market is the average absolute error in his estimate of v. In Table 5, we report the mean absolute difference between an uninformed agent's expectation about the fundamental value and the true value.²⁰ As the table shows, the greater the number of informed agents in the market, the better the estimates of the current common value. However, even when all agents are informed, market observables are only partially revealing.

Regime	Mean absolute error in uninformed			
	agents' belief about v			
Low volatility, all informed	0.31			
Low volatility, speculators informed	0.49			
High volatility, speculators informed	2.59			

Table 5: Errors in beliefs across different regimes

Notice from Table 5 that the market has a certain resilience. If the proportion of informed traders is low, an uninformed agent forms less precise estimates about the common value. However, in this case, adverse selection is less of an issue, since the probability of trading with an informed agent is low. Conversely, if there are a large number of informed agents, market observables are more revealing, mitigating the adverse selection problem.

²⁰In the model in which all agents are informed, we use the beliefs of agents who deviate at the information acquisition stage to determine an uninformed agent's expectation about the fundamental value.

When the fundamental value has high volatility, uninformed agents' beliefs are less precise, and the mean absolute error in estimating the common value increases to 2.59 ticks.

Recall that our traders observe both a snapshot of market history and a subset of the current book on each arrival to the market. The history consists of the price and sign of the previous transaction.²¹ To determine which conditioning variables are most important to uninformed traders in updating their beliefs about the common value, we regress an uninformed agent's expectation of the common value on these variables. Nothing in the model suggests that this functional relationship should be linear: our goal with these regressions is merely to illustrate broad rules of thumb that explain how market observables influence beliefs about the common value. We find that linear regressions perform surprisingly well in this context.

We generated data for the regressions in the following way: for each set of parameters, after the algorithm has converged to an equilibrium, we hold values fixed and run a new simulation. Each time an uninformed agent enters the market in the new simulation we determine his estimate of the current common value, given market observables and given the common value at the lag Δ_t . We restrict attention to books that are nonempty on both sides of the market. In the low volatility case with only speculators informed, these represent 84.4% of all books encountered in the simulation.

Details of the regressions are omitted for brevity. The constant is insignificantly zero (as expected), and the bid, ask, and last transaction prices contribute almost equally to the estimate of v. All else equal, every increase of a tick in any one of these prices increases the expectation of v by approximately 0.25 ticks.²² Depths at and away from the quotes are statistically significant, but have close to zero correlation with the dependent variable. An extra limit order at the bid (ask) price increases (decreases) the estimate of v by 0.13 ticks, and an extra limit order away from the bid (ask) increases (decreases) the estimate of v by 0.4 ticks. The R^2 of the regression is 0.81.

Similar regressions were run the other information structures with all agents informed and agents with $|\alpha| \in \{0.1, 4\}$ informed. The results were qualitatively similar, though the coefficients vary somewhat with information structure.

5.2 Allocative Efficiency

We now consider aggregate welfare generated by the limit order market, which we measure by the mean surplus per trader. The surplus of a particular trader is defined as his instan-

²¹In our simulations, at least one transaction occurs in each $[t - \Delta_t, t]$ interval.

²²Note, however, that in the simulation the transaction price, bid, and ask are all correlated, so all else is rarely equal.

taneous utility at execution, given by equation (2), discounted back to his first arrival at the market.

Eventual allocations in this market are approximately invariant to the information structure. Agents with $\alpha = 4$ or 8 buy a share, and those with $\alpha = -4$ or -8 sell a share. Among speculators, roughly $\frac{2}{3}$ of agents with $\alpha = 0.1$ buy a share, with the remainder (roughly $\frac{1}{3}$) selling a share. These fractions are reversed for the $\alpha = -0.1$ agents. Any effects of adverse selection on mean surplus, therefore, do not occur via a difference in allocations.²³ Further, as Table 1 indicates, execution times do not vary systematically across information regimes: when speculators take longer to execute, agents with $|\alpha| \in \{4, 8\}$ trade more quickly. What is the overall effect of adverse selection on mean surplus?

Recall that there are three key frictions present in our model. First, traders arrive over time and waiting is costly. Second, prices are discrete. Finally, traders have private information about type and the common value. As a first-best benchmark, we consider a frictionless world with all agents present in the market at the same time. Then, a price $p^* = v$ represents a competitive equilibrium, and the resulting allocation is Pareto-optimal.²⁴ If each agent trades instantaneously at v, he obtains a gross payoff $|\alpha|$. Thus, given the probability distribution of types, the aggregate frictionless surplus is 4.03.

Regime	Gross surplus	Net Surplus
Low volatility, all informed	3.73	3.73 - c
Low volatility, speculators informed	3.72	3.72 - 0.3c
High volatility, all informed	3.61	3.61 - c
High volatility, speculators informed	3.42	3.42 - 0.3c

Table 6: Mean surplus per trader, in ticks

The mean surplus per trader in each regime is shown in Table 6. For ease of comparison, we report both the gross and net (after deducting information acquisition costs) surplus numbers from the four regimes under consideration.

As expected, in all cases the gross surplus in the market is less than the frictionless benchmark 4.03. When the asset has low volatility, the gross surplus is relatively invariant to asymmetric information. Even when only speculators are informed, the market recovers

 $^{^{23}}$ Our results on allocations complement those of Blouin (2003), who considers a large decentralized economy in which a good with two different qualities is traded via bilateral bargaining. In his model, all units of the good are traded in equilibrium, so that adverse selection does not affect the eventual allocations, but does affect prices and the time taken to trade. This contrasts with the equilibrium when trading is centralized (so that all trades must occur at the same price and same point of time).

 $^{^{24}}$ All prices between -0.1 and 0.1 represent Walrasian equilibria given our parameters; allocations are invariant across these equilibria.

92.1% of the surplus generated in the frictionless case. This is very close to the calculations of Hollifield et al. (2004), who estimate that, for the stocks they study, the consummated gains from trade are approximately 90% of the maximum gains from trade.

As mentioned earlier, their methodology may over-estimate the gains from trade, since they do not consider asymmetric information. In a model with asymmetric information (or equivalently, one in which the gains to trade are small relative to the degree of adverse selection), the limit order book performs worse.

When the asset has high volatility, adverse selection is an important trading friction. The increase in fundamental volatility, even when all agents acquire information, leads to a reduction in aggregate gross surplus of approximately 3% (from 3.73 with low volatility to 3.61 with high volatility). Asymmetric information leads to a further fall in gross surplus of 0.19 ticks per trader, or approximately 5.25%.

6 Endogenous Information Acquisition

To understand how agents trading behavior changes with asset volatility, we characterize agents' willingness to pay for information. To do so, we fix agents' information acquisition strategies and solve for equilibrium in the trading game. We consider regimes in which all agents with a given α take the same action in the information acquisition game: all acquire information, or choose not to (i.e., for each θ , $\sigma_I(\theta) = 0$ or 1). As we show below, the amount agents are willing to pay for information declines in $|\alpha|$. Hence, we report results from four information acquisition regimes: all agents are informed, only agents with $|\alpha| \in \{0.1, 4\}$ are informed, only agents with $|\alpha| = 0.1$ are informed, and no agents are informed. That is, we ignore information structures in which the speculators are uninformed, but some other type is informed.

For each of the four information structures we consider, once the algorithm has converged, we simulate a further 100 million trader decisions and obtain the expected consumer surplus (i.e., the equilibrium payoff or expected utility) for each trader type. We use the equilibrium values and strategies of the corresponding trading game to determine the payoff to an agent who deviates in information acquisition. We allow a small mass of each type (2%) to deviate in information acquisition and then trade optimally. All other agents in the simulation play the equilibrium of the original trading game. The equilibrium strategies and payoffs for the deviators (and only the deviators) are determined afresh by the algorithm. We ensure that at most one deviator is present in the market at any given time, to preserve the spirit of unilateral deviation.

Consider the low volatility case. The gross payoff (i.e., ignoring the cost of acquiring

information) of each type in each of the four information structures is reported in Table 7. All payoffs are quoted in ticks. We exhibit the mean payoffs both to agents playing according to equilibrium and to deviators. The value of information to each agent is represented by the difference in payoffs between being informed and remaining uninformed. This value immediately translates to a maximum cost agents with a given α are willing to pay for information.

Information Structure		Value of $ \alpha $		
		0.1	4	8
None Informed	Equilibrium	0.403	3.515	7.333
	Deviation	1.178	3.652	7.353
	Value of Information	0.776	0.138	0.020
Speculators Informed	Equilibrium	0.628	3.413	7.204
	Deviation	0.413	3.499	7.228
	Value of Information	0.215	0.086	0.024
$\alpha \in \{.1, 4\}$ Informed	Equilibrium	0.495	3.508	7.234
	Deviation	0.287	3.414	7.279
	Value of Information	0.207	0.094	0.044
All Informed	Equilibrium	0.447	3.510	7.311
	Deviation	0.244	3.430	7.251
	Value of Information	0.203	0.080	0.060

Notes

(i) Reported means and standard deviations are averages and sample standard deviations from market simulations over 100 million arrivals (new and returning traders).

(ii) Reported numbers exclude agents who trembled to suboptimal actions.

(iii) Standard errors on means are less than .0005 for equilibrium strategies and less than .0020 for deviator strategies (for which only 2% of traders deviate).

(iv) Payoffs in *italics* indicate informed agents.

Table 7: Average gross payoffs (in ticks) in low volatility case

As seen from the table, the value of information decreases in the absolute value of α . Information is most valuable to the speculators (i.e., agents with $|\alpha| = 0.1$), who have little intrinsic benefit to trade. These agents are willing to take either side of the market, depending on the available payoff. Conversely, agents with $\alpha = 8$ are unlikely to switch from buyers to sellers, so information is less valuable to these agents.²⁵

 $^{^{25}}$ Radner and Stiglitz (1984) demonstrate that information is valuable to a single Bayesian decision-maker only if it induces a change in action.

The value of information to an agent, of course, equals her willingness to pay for information. Thus, speculators have the highest willingness to pay for information. Verrecchia (1982) shows a similar result in a general equilibrium rational expectations framework—the least risk-averse agents (i.e., those with the lowest intrinsic motive to trade) acquire costly information.

Observation 1 In any information regime, agents' willingness to pay for information is decreasing in $|\alpha|$.

Since information is most valuable to speculators, in any equilibrium in which some agents are informed, the speculators (or a subset of them) must be informed. That is, there will not exist any equilibria in which, for example, only the agents with $|\alpha| = 8$ are informed.

We also consider how agents' willingness to pay for information, changes as asset volatility changes. Specifically, if the underlying asset is more volatile, all agents are willing to pay more for information.

7 Cross-sectional Asset Pricing Implications

In our model, all our prices are relative to v_t , the fundamental value of the asset. If all agents were informed then they would all agree on this as the value. Our transaction price can be decomposed as

$$p_t = v_t + u_t,$$

where p_t is the price at time t and u_t is the deviation from the fundamental value. Typically, in econometric studies, u_{it} is assumed to be white noise.²⁶

Further, u_{it} is uncorrelated with v_t , the level of the fundamental value. However, changes in u_t are correlated with changes in the fundamental value at both short and long horizons. Let $\Delta(x_t) = x_t - x_{t-1}$. Table ?? exhibits the correlations between $\Delta(u_t)$ and $\Delta(v_t)$ at high and low frequencies. The correlations are negative and they are larger absolute numbers when the asset has high volatility. Further, they have a higher absolute value if there is asymmetric information: i.e., if the speculators are informed. As we have observed, in any regime the speculators are willing to pay the highest cost to become informed, therefore for any cost of acquiring information, they are most likely to be informed.

To see the cross-sectional implications of these correlations, consider a world in which the CAPM is true. To emphasize differences across securities, we adopt a subscript i to

²⁶Habrouck (2002) provides a comprehensive discussion of these decompositions.

identify each asset. Therefore, if the CAPM is true, in the absence of frictions, the return process for asset i is

$$r_{i,t}^* = r_f + \beta_i r_{m,t} + \epsilon_{i,t}. \tag{3}$$

Here, we use the obvious notation, in which r_f is the risk free rate and ϵ_i is a firm specific idiosyncratic shock.

This relationship could be uncovered from transaction data. Such an estimate would be based on the observed price process, or $r_{i,t} = \frac{p_{i,t}-p_{i,t-1}}{p_{i,t-1}}$. The estimated beta, $\hat{\beta}$ is determined as

$$\hat{\beta} = \frac{cov(r_i, r_m)}{\sigma_m^2}.$$

As the CAPM is true,

$$r_{m,t} = \frac{r_{i,t}^*}{\beta} - \frac{r_f + \epsilon_{i,t}}{\beta}.$$
(4)

Using the properties of covariance, $\hat{\beta}$ is

$$\hat{\beta} = \frac{1}{\beta} \frac{cov(r_{i,t}, r_{i,t}^* - \epsilon_{i,t})}{\sigma_m^2}.$$
(5)

This implies that if the observed price process was the true one, so that $r_{i,t}^* = r_{i,t}$, then equation 5 simplifies to

$$\hat{\beta} = \frac{1}{\beta} \frac{cov(r_{i,t}^*, r_{i,t}^* - \epsilon_{i,t})}{\sigma_m^2}$$
(6)

$$= \frac{var(r_{i,t}^*) - \sigma_{\epsilon}^2}{\beta \sigma_m^2} \tag{7}$$

$$= \frac{\beta^2 \sigma_m^2}{\beta \sigma_m^2} \tag{8}$$

$$= \beta,$$
 (9)

where $E(r_m - Er_m)\epsilon_{i,t} = 0.$

However, the observed price process is not the "true one." In the context of our model, the "true," price of the asset is v_{it} . Therefore,

$$r_{i,t}^* = \frac{v_{i,t} - v_{i,t-1}}{v_{i,t-1}}.$$
(10)

High fundamental volatility correspond to stocks with either high CAPM β 's, or with high idiosyncratic shocks.

Substituting in for $r_{i,t}^*$ and $r_{i,t}$ into equation 5 yields

$$\hat{\beta} = \frac{1}{\beta \sigma_m^2} cov \left(\frac{v_{i,t} - v_{i,t-1}}{v_{i,t-1}}, \frac{p_{i,t} - p_{i,t-1}}{p_{i,t-1}} \right) \\ = \frac{1}{\beta \sigma_m^2} cov \left(\frac{\Delta(v_{it})}{v_{i,t-1}}, \frac{\Delta(v_{i,t}) + \Delta(u_{i,t})}{p_{i,t-1}} \right)$$

Notice, that this implies that anything else that appears to be correlated with microstructure noise, such as bid/ask spreads should also be informative about returns as these variables will pick up the misestimation in the asset pricing model. We explore the implications of this in a companion paper. To get a sense of the economic magnitude of the missetimation, in Table ??, we demonstrate the amount by which β is overestimated in simple cross–sectional regressions.

To understand the intuition behind these effects consider two assets, one with a high volatility and one with a low volatility. There are two effects that work in the same direction.

First, the adverse selection risk (including both the likelihood and size of the jump) is larger when asset volatility is higher. All classes of agents change their behaviour. Specifically, speculators with little intrinsic value for trade supply less liquidity and prefer to submit market orders (against mispriced limit orders) or leave the market and return. Therefore, conservative investors (those with a high absolute value to trade) submit a higher proportion of limit orders. This means that when transactions do occur, they are most likely after a limit order has been "picked off." This will generate a time series in which changes in the common value are correlated with changes in the microstructure noise. The misestimation of β is therefore systematically related to the underlying risk.

Second, holding fixed the set of agents who have acquired information, the more volatile an asset the more valuable (the greater the willingness to pay for) information. Further, holding fixed the set of agents who have acquired information, information is always more valuable to the speculators. This effect exacerbates the adverse selection risk, as informed agents are more likely to pick off ex post mispriced limit orders and therefore not to supply liquidity but to demand it.

Our results are robust to different time horizons, and to different estimates of the "true," price process. For example, sometimes the midpoint of the bid ask spread is taken as a proxy for the true price process. In Table **??** we demonstrate that this measure, too introduces bias in the estimation of systematic risk.

8 Conclusion

We model how agents change their trading behavior in the fact of greater adverse selection. This can come about either because asset volatility is higher ore because of asymmetric information. Agents with low intrinsic benefits from trade have the highest value for information and are most likely to be informed.

The informational efficiency of market observables is directly related to the number of informed agents in the market—an uninformed agent forms more precise estimates of the common value when there are a greater number of informed agents. Transparency of the book somewhat mitigates adverse selection, since market observables are more informative when an uninformed trader faces greater adverse selection. Since liquidity suppliers in our market (i.e., agents with low private values) have the highest value for information, bid and ask quotes are as informative as transaction prices.

Conversely, when the asset volatility is high, sequential trade and asymmetric information are both important market frictions. Each of these leads to an increase in the excess volatility of prices over the fundamental value. The composition of potential traders at any given point of time changes, with more extreme value traders in the market. Intriguingly, with high fundamental volatility, the book is a little deeper on average when there is asymmetric information, since informed traders submit the bulk of the limit orders in the market.

The change in the marginal trader across regimes also suggests a problem with aggregating up to a representative agent. If two assets have different fundamental volatilities, they will have different pools of potential traders in the market at any time. Hence, the aggregate "representative agent" inferred from these two markets will be different. Although we have risk-neutral traders (and therefore no wealth effects in the utility function), the representative agent remains an elusive concept.

Finally, we have shown that because agents' trading strategies change in the presence of adverse selection, changes in microstructure noise are negatively correlated with fundamental volatility. Therefore, any econometric specification that assumes such noise is uncorrelated with fundamental volatility underestimates factor exposure. This leaves open the possibility that other variables that are correlated with microstructure noise (such as many transaction costs or liquidity measures), should be priced.

A Appendix

A.1 Model Description: Trading Game

In this section, we describe the trading game more formally. Recall that the trading game assumes that traders' information acquisition decisions are fixed. Let $I \in \{0, 1\}$ denote the action an agent takes with respect to information acquisition, where I = 1 if the agent chooses to become informed. Let $\sigma_I(\theta) \in [0, 1]$ denote the information acquisition strategy of an agent of type θ : this is the probability that type θ acquires information.

Let G denote the distribution over a trader's reentry time, with g being the associated density. With a slight abuse of notation, let ℓ^p denote the outstanding limit orders at price p before the agent submits an order. Then, for a new order, priority $\hat{q}(p, x)$ is:

$$\hat{q}(p,x) = \begin{cases} 0 & \text{if (i) } x = 0, \text{ or (ii) } x = 1, p \ge A(L), \text{ or (iii) } x = -1, p \le B(L)_{(11)} \\ |\ell^p + x| & \text{ otherwise.} \end{cases}$$

Let $s(\theta, I)$ be the state observed on a particular entry to the market at time t by a trader with type θ who has taken the information acquisition action I. Here, $s(\cdot)$ includes: (i) the history of play in the game, and the history of changes in v up to time $t - \Delta_t$ (if I = 0; i.e., the agent is uninformed) or time t (if I = 1; i.e., the agent is informed). If the trader had previously entered the market and taken an action, $s(\cdot)$ includes the status of the previous action, a = (p, q, x), where p is the price at which the previous order was submitted, q its current priority at price p, and x, which is defined in equation (1) of the text to take on the value +1 for a buy order, -1 for a sell order, and 0 if no action was taken.

(ii) a variable $z \in \{0, 1\}$ that denotes the number of shares the agent has available to trade. Each trader enters with z = 1. Once he has traded, z is set to zero. As we comment after the Bellman equation (14) below, this variable is used to conveniently set an agent's future payoff to zero once he has traded.

Consider the problem faced by a trader in the market at time t. Suppose this trader is reentering the market (the problem faced by a new trader is identical to the problem faced by a reentering trader who did not submit an order on his previous entry), and, on his previous entry (at some t' < t), he had submitted an order at price p that is still active. This order may have improved in priority at price p between times t' and t. The trader has the option of leaving the order unchanged and taking no further action.

Let $\mathcal{A}(s)$ denote the feasible action set of a trader in state s. Recall that s depends on type θ and information acquisition strategy I, and includes information on the status of the trader's previous action, a = (p, q, x). If the trader has arrived at the market for the first time, we set x = 0. For computational tractability, we restrict limit order submission to a finite set of prices within k ticks of an agent's expectation of the common value. We choose k to be sufficiently large that it does not affect the equilibrium. Denote the agent's expectation of common value as $\hat{v}(s) = E(v \mid s)$, where s denotes the state observed by the agent. The feasible action set is then defined as

$$\mathcal{A}(s) = \{ (p,q,x) \mid (i) \ x \in \{-1,0,1\}, \ (ii) \ q = \hat{q}(p,x), (iii) \ q \neq 0 \Longrightarrow p \in [\hat{v}(s) - k, \hat{v}(s) + k] \cap \mathcal{P} \}.$$
(12)

In the trading game, the information acquisition strategy for each trader type θ is fixed as $\sigma_I(\theta)$, so we define the type of a trader as (θ, I) . Let $\sigma_I = \{\sigma_I(\theta)\}_{\theta \in \Theta}$ denote the information acquisition strategy across all types. Let $\Theta^I(\sigma_I) = \{(\theta, I) \in \Theta \times \{0, 1\} \mid \operatorname{Prob}((\theta, I) \mid \sigma_I(\theta)) > 0\}$ be the set of feasible types in the trading game. In what follows, we suppress the dependence of this set on σ_I .

Let $S_{(\theta,I)}$ denote the set of feasible states a trader with type θ and information I may encounter. A mixed strategy for such a trader in the trading game is then a map $\sigma_{(\theta,I)}$: $S_{(\theta,I)} \to \prod_{s \in S_{(\theta,I)}} \Delta(\mathcal{A}(s))$, where $\Delta(\mathcal{A}(s))$ is the set of probability distributions over $\mathcal{A}(s)$. Let $\mathcal{S} = \bigcup_{(\theta,I) \in \Theta^I} S_{\theta,I}$ be the entire set of states for the game, and let $\sigma = \{\sigma_{(\theta,I)}\}_{(\theta,I) \in \Theta^I}$ denote a strategy in the trading game.

Consider a trader in the market at time t. Suppose he enters in a state s, and the status of his previous action is given by a. When the trader submits an order, he must consider the distribution over execution times for that order, as well as the distribution of his own reentry time into the market. Upon reentry, if his order is unexecuted, he has the option to cancel it and submit a new order. The trader, therefore, solves a dynamic program to determine the optimal order.

Consider the value to trader type θ , with information acquisition strategy $\sigma_I(\theta)$, of being in the state s, given that his previous order is a. On entry into the market, the trader has a finite action set, $\mathcal{A}(s)$. Each action \tilde{a} in this set gives rise to an expected payoff that consists of two components: first, a payoff conditional on the order executing before the trader reenters the market, and second, the value associated with reentering the market (without having executed in the interim) in some new state s.

The likelihood of a limit order executing clearly depends on the strategies of other players in the game. Since we consider only symmetric equilibria, consider a trader in the market, and let $\sigma = \{\sigma_{(\theta,I)}\}_{\theta \in \Theta^I}$ denote the strategy adopted by every other player. For convenience, normalize the trader's entry time to 0. Let $\phi(\tau, v; s, \tilde{a}, \sigma)$ be the probability that an action $\tilde{a} = (\tilde{p}, \tilde{q}, \tilde{x})$ taken in state s at time 0 leads to execution at time $\tau > 0$ when the common value is v, given that all other agents are playing σ , and let $f(v \mid s, t)$ denote the density function over v at time τ , given state s. For an informed trader, $f(\cdot)$ is purely exogenous (recall that the fundamental; for an uninformed one, it incorporates the trader's beliefs over v_0 .

Suppose the trader reenters the market at some time w > 0. His expected payoff due to execution prior to reentry is

$$\pi(s,\tilde{a},w,\sigma) = \int_0^w \int_{-\infty}^\infty \left(e^{-\rho t} \tilde{x} \left(\beta + v_t - \tilde{p} \right) \phi(\cdot) \right) f(v \mid s,t) \, dv \, dt.$$
(13)

This equation is derived as follows. Suppose the agent's order executes at a time $t \in [0, w]$. The payoff to the order depends on the common value at time t, which we denote v_t . As noted, the instantaneous payoff of this order at time t is $\tilde{x}(\beta + v_t - \tilde{p})$. This payoff must then be discounted back to time 0, at the rate ρ . The innermost integral of the first term is over the different common values that can obtain at time t. We expect $\phi(\cdot)$ to be higher when v has moved in an adverse direction (for example, v has decreased after a limit buy was submitted)—this is another manifestation of adverse selection in this model. For a market order, we have $\phi(0, \cdot) = 1$, since the order executes immediately. The outermost integral is over the possible times at which execution could occur.

Recall that the reentry time is random and exogenous, with probability distribution $G(\cdot)$. Let $\nu(s' \mid s, \tilde{a}, w, \sigma)$ denote the probability that the trader observes state s' on reentry, given action \tilde{a} , previous state s, elapsed time w since entry into the market, and strategy of other players σ . Finally, let J(s) denote the value to an agent of being in state s. The Bellman equation for the agent's problem is

$$J(s,\sigma) = \max_{\tilde{a}\in A(s)} \int_0^\infty \left\{ \pi(s,\tilde{a},w,\sigma) + e^{-\rho w} \int_{s'\in S_\theta} J(s',\sigma) \nu(s'\mid s,\tilde{a},w,\sigma) ds' \right\} dG(w).$$
(14)

The first term on the right-hand side (defined in (13) above) indicates the payoff from execution before reentry at the random time w. The second term captures the continuation payoff to the trader on reentry to the market at time w. If his order executes before he reenters, we have z' = 0 (i.e., he cannot trade any more shares). Define $J(s', \sigma) = 0$ for all s' such that z' = 0, to ensure that the continuation payoff is set to zero if an order executes before the trader reenters the market.

The agent reenters the market at the random time w in some state s' different from s. If his previous order is still unexecuted, he can choose instead to submit a new order at a price $\tilde{p} \neq p$, and possibly in a direction $\tilde{x} \neq x$. A new order implies cancellation of the previous order. Alternatively, he can choose to leave his previous order on the books by setting $\tilde{p} = p$ and $\tilde{x} = x$. Of course, market conditions may have changed since he first submitted the order, either due to exogenous reasons (e.g., a change in the common value) or due to actions taken by other agents. The latter could enhance the priority of this agent's order at the price p, or it could reduce the overall priority if other agents submitted limit orders at prices more aggressive than p. Hence, the action a taken at time 0 evolves to a' by the time the trader reenters at time w. The outermost integral is over the random reentry time.

Since the action set is finite on any entry, the maximum over all feasible actions exists and is well-defined. The value of a state and previous action pair is just the maximal expected payoff over all feasible actions the trader can take.

Fixing the strategies of all other agents, a given pure strategy $y^*_{(\theta,I)}$ for a trader with type $(\theta, I) \in \Theta^I$ is a best response if (and only if), for every $s \in S_{(\theta,I)}$,

$$y_{(\theta,I)}^*(s) \in \arg\max_{\tilde{a}\in A(s)} \int_{w=0}^{\infty} \left\{ \pi(s,\tilde{a},w,\sigma) + e^{-\rho w} \int_{s'\in S_{\theta}} J(s',\sigma) \nu(s'\mid s,\tilde{a},w,\sigma) ds' \right\} dG(w).$$
(15)

Note that some of these states may not be attained in equilibrium. Nevertheless, we require the trader to act optimally in these states as well. Also, the trader's optimal action in any state must take into account the possibility of future reentry (and that the trader will play optimally in the new state).

Finally, a strategy for each player is defined as $y = \{y_{\theta,I}\}_{\theta \in \Theta^I}$. A strategy $y^* = \{y^*_{(\theta,I)}\}_{\theta \in \Theta^I}$ represents a stationary Markov-perfect equilibrium of the trading game if, for each pair $(\theta, I) \in \Theta^I$, $y^*_{(\theta,I)}$ is a best response in every feasible state $s \in S_{(\theta,I)}$, given that all other agents are using the strategy y^* .

A.2 Details of the Numerical Algorithm

We fix information acquisition strategies $\sigma_I(\theta)$, and solve for the equilibrium of the corresponding trading game.²⁷ We use an asynchronous value function iteration procedure, similar to Pakes and McGuire (2001), to find a $J(s,\sigma)$ that satisfies the Bellman equation in (14).

For the numerical implementation of the model, we restrict the state space as follows. Let $m_t(I)$ denote the market conditions observed by an agent at time t (recall that informed agents, with I = 1, observe the current common value; uninformed agents, with I = 0, observe it with a lag). We use

$$m_t(0) = \{\hat{L}_t, v_{(t-\Delta_t)}, \hat{p}_t, b_t\},\$$

$$m_t(1) = m_t(0) \cup \{v_t\}.$$

where \hat{L}_t is a set of variables that depend L_t , on the book at time t, \hat{p}_t is the price of the most recent transaction, and b_t is a variable indicating whether the most recent transaction

 $^{^{27}}$ The overview of the algorithm appears in section 2.2.

was a buy $(b_t = 1)$ or sell $(b_t = -1)$.²⁸

Ideally, we would like agents to condition on the entire book; in practice, this becomes computationally intractable. Instead, the variables in \hat{L}_t consist of: (i) the current bid and ask prices, (B_t, A_t) ; (ii) the total depths at these prices, (ℓ_t^B, ℓ_t^A) ; and (iii) the cumulative buy and sell depths in the book, $D_t^b = \sum_{i=0}^N \{\ell_i^t > 0\}$ and $D_t^s = \sum_{i=0}^N \{\ell_i^t < 0\}$.²⁹

Given the market observables used, in the numerical algorithm, the state at time t for an agent of type θ who makes the information acquisition decision I is defined by s = $\{\theta, m_t(I), a, z\}$, where a = (p, q, x) denotes the status of his previous order, and $z \in \{0, 1\}$ denotes how many shares he can trade (z is set to zero once he has traded).

In the algorithm, at each time t, each action \tilde{a} in each state s encountered by the simulation has an associated payoff $U_t(\tilde{a} \mid s)$. This payoff is a real number, and is the expected discounted payoff from taking action \tilde{a} in state s. Hence, it may be interpreted as the current belief of an agent about the payoff from this action.³⁰

At any point of time, current beliefs $U_t(\cdot)$ imply an optimal strategy profile y_t , which assigns the payoff-maximizing action in each state. Let $\tilde{a}^*(s) \in \arg \max_{\tilde{a} \in \mathcal{A}(s)} U_t(\tilde{a} \mid s)$ denote the optimal action in state s. Then, given beliefs $U_t(\cdot)$, the value of state s is determined as $J(s, y_t) = U_t(\tilde{a}^*(s) \mid s)$.

Each action and state pair, (\tilde{a}, s) has an initial belief $U_0(\tilde{a} | s)$. These initial beliefs are set as follows. Consider a limit buy order at price p when the last observed common value is v. The initial belief for such an order is the payoff $\beta + v - p$ discounted by the expected time until the arrival of a new trader for whom being a counterparty yields a nonnegative payoff. This initial value is optimistic since (i) limit orders tend to execute when the common value moves in an adverse direction, and (ii) counterparties usually hold-out for a strictly positive payoff. The initial belief for market orders also assumes the common value is unchanged from its last observed value, but of course involves no discounting. Given that we allow traders to tremble, any $U_0(\cdot)$ can eventually lead to an equilibrium. The choice of initial beliefs is driven more by computational considerations (in particular, converging to

²⁸We investigated a model in which agents also observe the cumulative market buys and sells in the interval $[t - \Delta_t, t]$. The added conditioning variables are virtually ignored by traders in updating beliefs about v, and do not affect market outcomes. In our model, only recent history is relevant to traders, for two reasons. First, traders leave the market forever after execution. Therefore, any knowledge about traders who have already executed does not affect agents' beliefs about future play in the game. Second, for uninformed agents, events prior to $t - \Delta_t$ offer no information about changes in v since it was last observed (at time $t - \Delta_t$). Nevertheless, any particular snapshot of history is potentially restrictive. Computational reasons require us to impose such a restriction; without it, the state space is too large.

²⁹Of these variables, the current bid and ask prices are the most important in influencing agents' actions. In some simulations, we restrict the book-related variables further to just the current quotes.

³⁰Our $U(\cdot)$ corresponds to the Q function in the Q-learning literature begun by Watkins (1989). Qlearning and other neuro-dynamic programming techniques related to our simulation algorithm are described in Bertsekas and Tsitsiklis (1996).

equilibrium more quickly) than by a theoretical need.

Additional details of the algorithm are as follows.

1. Three types of exogenous events drive the simulation—changes in the common value, the arrival of new traders, and the reentry of old traders who have not yet executed. At each point in time, let t_v denote the additional time until v changes, t_n the additional time until a new trader arrives, and t_r a vector of additional times until each old trader returns to the market to possibly revise his order. Let $\underline{t}_r = \min\{t_r\}$ denote the earliest reentry time across all old traders.

Whenever an event occurs, we redraw the time until its next occurrence accordingly (recall that the time interval between events for a Poisson process has an exponential distribution). We also subtract the elapsed time from the other "time until" variables.

At time 0, we start with an empty book, new draws for t_v and t_n , and no existing traders (i.e., t_r is an empty vector).

In theory, the initial common value can be chosen arbitrarily. However, since v follows a random walk, the price grid would need to be infinite. To avoid this problem, the algorithm records all prices relative to the current v, and appropriately shifts all orders on the book whenever v changes.³¹ The number of ticks around v for which orders are tracked is chosen sufficiently high that orders never "fall off" the grid. That is, orders get revised by returning traders before becoming too unaggressive for the grid, or get picked-off before becoming too aggressive for the grid. We use an odd number of ticks, with v_t itself lying on a tick at all times.

2. At time $t = \min\{t_v, t_n, \underline{t}_r\}$, an exogenous event occurs. Suppose $t_v < t_n$ and $t_v < \underline{t}_r$. Then, the common value changes at time t_v ; with probability $\frac{1}{2}$ it increases by one tick, and with probability $\frac{1}{2}$ it decreases by a tick. As specified in 1. above, we adjust the times for the three events as follows. We set $t_n = t_n - t_v$ and $t_r = t_r - t_v$, and then draw a new time t_v for the next change in v.

Suppose, instead, $t_n < t_v$ and $t_v < \underline{t}_r$. A new trader arrives to the market. His type is denoted as $\theta = \{\rho, \beta\}$. The discount factor ρ is the same for all traders, and β is drawn independently from the distribution F_{β} . The times for the three events are adjusted as specified in 1.

³¹Importantly, uninformed traders observe the prices of orders on the book relative to $v_{t-\Delta_t}$, else they could directly infer v_t as the mid-tick in the book. That is, the algorithm tracks the book relative to v_t but presents it to traders relative to their last observed v.

A given trading game is used to obtain payoffs to either equilibrium strategies or to deviator strategies in the information acquisition game. When obtaining equilibrium payoffs, we set $I = \sigma_I(\theta)$ for the new trader. When obtaining payoffs to deviating, we classify a trader as a deviator with probability 0.01, as long as no other deviators are currently in the market (to preserve the notion of unilateral deviation). If the new trader is a regular trader, we set $I = \sigma_I(\theta)$. If he is a deviator, we set I = 1when $\sigma_I(\theta) = 0$ and I = 0 when $\sigma_I(\theta) = 1$ (i.e., a deviator acquires information only when regular traders of his type do not). Importantly, beliefs and trading strategies of non-deviators are held fixed throughout the algorithm when obtaining payoffs to deviating in the information acquisition game.

Since the trader is new, we set z to 1 and his previous action x to 0. The trader observes the state $s = \{\theta, m(I), a, z\}$ and takes an action \tilde{a} . If he submits a market order, he executes and leaves the market for ever. If he takes any other action, we draw his random return time and include it in the vector t_r . We also draw a new random time t_n before the arrival of the next new trader.

Finally, suppose $\underline{t}_r < t_v$ and $\underline{t}_r < t_n$. An old trader returns to the market. He observes the current state $s = \{\theta, m(I), a, z\}$ which includes the current status a of his previous action. He then takes some action (which could include retaining his previous order). If he submits a market order, he executes and leaves the market for ever. If he takes any other action, we draw his new return time in t_r , and adjust the times t_n and t_v as specified.

3. Suppose a trader of type θ is in the market at time t. The trader observes the current state $s = \{\theta, m(I), a, z\}$ and chooses a payoff-maximal action $\tilde{a}^*(s) \in \arg \max_{\tilde{a} \in \mathcal{A}(s)} U_t(\tilde{a} \mid s)$. If the trader is informed, he knows v_t , which determines $\mathcal{A}(s)$. If he is uninformed, his belief about v_t is used to determine $\mathcal{A}(s)$. Denote this belief as $E(v_t \mid m_t(0))$.

Beliefs about the current common value are updated in the following manner.³² Let $\alpha_t(m_t(0)) = E(v \mid m_t(0)) - v_{t-\Delta_t}$ denote the extent by which an uninformed agent at time t revises his belief about v_t , given a lagged value $v_{t-\Delta_t}$. Since we consider stationary equilibria, we drop the time subscript on market conditions. Start with an initial belief $\alpha_0 = 0$ for each market m(0). Let r(m(0)) be an integer denoting the number of times market conditions m(0) are encountered in the simulation. We drop the argument of r for notational convenience. Each time a trader observes market

 $^{^{32}}$ Note that these beliefs do not depend on an agent's type. Hence, this updating can be (and is) performed even when the trader in the market is informed about the current value of v.

conditions m(0), we increment r by 1, and set

$$\alpha_r(m(0)) = \frac{r-1}{r} \alpha_{r-1}(m(0)) + \frac{1}{r} (v_t - v_{t-\Delta_t}).$$
(16)

Recall that $\hat{v}(m) = E(v \mid m)$ denotes a trader's estimate of the common value. For an uninformed trader who enters in market m(0), this estimate is $\hat{v}(m(0)) = v_{t-\Delta_t} + \alpha_{r-1}(m(0))$.

Using this estimate \hat{v} , the action set for each trader is defined as in equation (12) of the text. Now, suppose the optimal action \tilde{a}^* does not represent a market order; that is, it is either a limit order or no order. Suppose further that, at some future point of time, t', the trader reenters the market. He finds that his action has evolved to \tilde{a}' , and the new market is m'. Denote $s' = \{\theta, m'(I), \tilde{a}', z\}$.

The action \tilde{a}^* thus generates a realized continuation value $J(s', y_{t'})$ on this visit, which is "averaged in" to the belief $U_t(\tilde{a}^* \mid s)$ in the following manner. We define

$$U_{t'}(\tilde{a}^* \mid s) = \frac{n}{n+1} U_t(\tilde{a}^* \mid s) + \frac{1}{n+1} e^{-\rho(t'-t)} J(s', y_{t'}).$$
(17)

Here, $n(\tilde{a}^*, s)$ is a positive integer that is incremented by one each time action \tilde{a}^* is chosen in state s (for notational brevity, the dependence of n on \tilde{a}^* and s is suppressed in equation (17)). We start with an initial positive integer n_0 for each action and state pair (\tilde{a}, s) . This integer affects the speed at which the algorithm converges, with larger values implying slower convergence. Periodically, during the simulation, we reset n to n_0 for some action and state pairs to obtain quicker convergence.

Similarly, suppose a trader submits a limit order (denoted by action \tilde{a}^*) at time t, and this order executes against a market order submitted by another trader at time t'. The actual payoff to that limit order in the simulation is $\tilde{x}(\beta + v_{t'} - \tilde{p}^*)$, where β denotes the private value of the trader. In this case, we update

$$U_{t'}(\tilde{a}^* \mid s) = \frac{n}{n+1} U_t(\tilde{a}^* \mid s) + \frac{1}{n+1} e^{-\rho(t'-t)} \tilde{x}(\beta + v_{t'} - \tilde{p}^*), \quad (18)$$

4. Whenever a trader takes an action, his belief about the payoff to a market order is updated in similar fashion. For example, let \tilde{a}_b denote the action that involves submitting a market buy order, given market m and previous action a. In the simulation, we (as modelers) know the payoff to a market order in every state, whether a trader is informed about the current value of v or not. Hence, these payoffs can be averaged in for market orders even when such orders are suboptimal for the trader. For this updating, we use equation (18), with t' = t and $v_{t'} = v_t$. In determining the payoff to agents who deviate at the information acquisition stage, we update beliefs for deviators along the same lines as in items 2 and 3. This allows us to determine the payoff to a deviator who plays optimally in the stage game, while holding strategies of other agents fixed at the equilibrium of the trading game that has no deviators.

5. In the simulation, most traders take the optimal action given current beliefs. If all traders did this, there is the possibility that the algorithm would be "stuck" at a non-equilibrium state—every trader of a given type would take the same action in that state, so these traders would never learn the payoffs to other actions in that state. If there is an error in beliefs, all traders of that type may play suboptimally.

To ensure that beliefs are updated for all actions in every state, we introduce trembles. Specifically, with probability ϵ a trader trembles over all suboptimal limit orders available to him. He chooses among suboptimal limit orders with equal probability. The algorithm will then naturally update the beliefs about payoffs to this action.³³

Convergence Criteria

We run the model for a few billion events until we check for convergence. Along the way, we evaluate the change in value functions every 100 million new trader arrivals, by computing $|U_{t_2}^{k_2}(\tilde{a} \mid s) - U_{t_1}^{k_1}(\tilde{a} \mid s)|$ for each pair (\tilde{a}, s) that occurs along the path of play in the simulation. Here, k_1 is the number of times the action \tilde{a} has been chosen in state s at the start of the current 100 million new trader arrivals, and $k_2 \ge k_1$ the number of times it has been chosen at the end of the current 100 million new trader arrivals. Further, t_1 and t_2 represent the actual time at the start and end of the 100 million arrivals.

Essentially, if this weighted absolute difference (weighted by $k_2 - k_1$) is small, that suggests the value functions have converged. When this weighted difference is below 0.01, we apply other convergence tests. At this point, we hold the beliefs $U(\cdot)$ fixed and simulate the model for a total 100 million more new trader arrivals (new and returning). Let $U^*(\tilde{a} \mid s)$ be the fixed beliefs. These imply an optimal strategy profile y^* . For each (\tilde{a}, s) , define $J(s, y^*) = \max_{\tilde{a} \in \mathcal{A}(s)} U^*(\tilde{a} \mid s)$.

We compare the empirical payoffs from different actions in the simulation to the fixed beliefs. This comparison is done at two levels. The first is a "one-step ahead" check based on the trader's next entry time or execution time, whichever is sooner. Suppose a trader

³³When a player trembles at t' > t, the payoff of the optimal action at t' is used to update $U_t^k(\cdot)$ to $U_{t'}^{k+1}(\cdot)$. Thus, traders do not anticipate behaving suboptimally in the future.

takes an action \tilde{a} at time t, and reenters at t' > t with a new state s'. His one-step ahead empirical payoff is taken to be $\tilde{J}_1(s, y^*) = e^{-\rho(t'-t)} J^*(s', y^*)$. If the trader takes an action \tilde{a} at t and executes at t' > t before he can reenter, his one-step ahead empirical payoff is $\tilde{J}_1(s, y^*) = e^{-\rho(t'-t)} \tilde{x}(\beta + v_{t'} - \tilde{p})$.

Second, eventually every trader in this model executes, and leaves the market. At the time he executes, he obtains a realized payoff. Suppose the trader enters at t, and eventually executes at t'. Let \tilde{a} denote his most recent action before execution. His realized payoff is then $\tilde{J}(s, y^*) = e^{-\rho(t'-t)} \tilde{x}(\beta + v_{t'} - \tilde{p})$.

We use three convergence criteria for each of the two comparisons above. The most stringent of these is a χ^2 test similar to that in Goettler, Parlour and Rajan (2004).³⁴ Suppose $J^*(\cdot)$ indeed represents equilibrium values. Since the computed values $J^k(\cdot)$ are averages, the central limit theorem implies that the empirical distribution of payoffs for each action in each state is approximately normal with mean J^* and a variance that is empirically determined from the simulation. Let $\eta(s) = \frac{\tilde{J}(s,y^*) - J^*(s,y^*)}{\sigma_s}$, where σ_s denotes the empirical standard deviation of payoff in state s (a similar variable is constructed for the one-step ahead payoffs). The variables $\eta(s)$ then have the standard normal distribution.

Let S be a set including all states encountered at least 100 times during the convergence check (this ensures that the central limit approximation is accurate). The test statistic $\gamma = \sum_{s \in S} \eta^2(s)$ sums the squares of the standard normal variables, and is distributed as a χ^2 with degrees of freedom equal to the number of states used in the summation, |S|. The algorithm has converged if the test statistic is less than the 1% critical value.

The other two tests are similar to those proposed by Pakes and McGuire (2001). First, we consider the correlation between beliefs $J^*(\cdot)$ and realized outcomes \tilde{J} or \tilde{J}_1 . This correlation exceeds 0.999. Second, we consider the mean absolute error in beliefs, weighted by the number of times the state and action are observed. This mean absolute error is less than 0.01.

 $^{^{34}\}mathrm{The}$ theoretical properties of this test were derived by den Haan and Marcet (1994).

References

- Admati, Anat R. (1985), "A Noisy Rational Expectations Equilibrium for Multi-Asset Securities Markets," *Econometrica* 53(3): 629–657.
- [2] Admati, Anat R. and Paul Pfleiderer (1987), "Viable Allocations of Information in Financial Markets," *Journal of Economic Theory* 43: 76–115.
- [3] Admati, Anat R. and Paul Pfleiderer (1988), "A Theory of Intraday Patterns: Volume and Price Variability," *Review of Financial Studies* 1(1): 3–40.
- [4] Back, Kerry and Shmuel Baruch (2006), "Working Orders in Limit-Order Markets and Floor Exchanges," forthcoming, *Journal of Finance*.
- [5] Back, Kerry, H. Henry Cao and Gregory A. Willard (2000), "Imperfect Competition among informed traders," *Journal of Finance* 55(5): 2117–2155.
- Barlevy, Gadi and Pietro Veronesi (2000), "Information Acquisition in Financial Markets," *Review of Economic Studies* 67(1):79–90.
- [7] Bassan, Bruno, Olivier Gossner, M. Scarsini and S. Zamir (2003), "Positive value of information in games," *International Journal of Game Theory* 32 p 17–31.
- [8] Berk, Jonathan B. (1997), "The Acquisition of Information in a Dynamic Market," Economic Theory 9: 441–451.
- [9] Bernardo, Antonio and Kenneth Judd (1997), "Efficiency of Asset Markets with Asymmetric Information," CMS working paper #16-97.
- [10] Bertsekas, D.P. and J.N. Tsitsiklis (1996) Neuro-Dynamic Programming, Athena Scientific, Belmont, MA.
- [11] Biais, Bruno, Christophe Bisiere, and Chester Spatt (2003), "Imperfect Competition in Financial Markets," Carnegie Mellon working paper.
- [12] Bloomfield, Robert, Maureen O'Hara and Gideon Saar (2004), "The Make or take decision in an electronic market: Evidence on the evolution of liquidity," *Journal of Financial Economics* forthcoming.
- Blouin, Max (2003), "Equilibrium in a Decentralized Market with Adverse Selection," Economic Theory 22: 245–262.

- [14] Blume, Lawrence and David Easley (1990), "Implementation of Walrasian Expectations Equilibria," Journal of Economic Theory 51: 207–227.
- [15] den Haan, Wouter J., and Albert Marcet (1994), "Accuracy in Simulations," Review of Economic Studies 61: 3–17.
- [16] Foucault, Thierry, Ohad Kadan and Eugene Kandel (2004), "The Limit Order Book as a Market for Liquidity," HEC working paper.
- [17] Foster, S. Douglas and S. Viswanathan (1996), "Strategic Trading when Agents Forecast the Forecasts of Others," *Journal of Finance* 51(4): 1437–1478.
- [18] Foucault, Thierry (1999), "Order Flow Composition and Trading Costs in a Dynamic Limit Order Market," Journal of Financial Markets 2: 99–134.
- [19] Glosten, Lawrence (1994), "Is the Electronic Open Limit Order Book Inevitable?" Journal of Finance 49(4): 1127–1161.
- [20] Goettler, Ronald, Christine A. Parlour and Uday Rajan (2004), "Equilibrium in a Dynamic Limit Order Market," forthcoming, *Journal of Finance*.
- [21] Grossman, Sanford and Joseph Stiglitz (1980), "On the impossibility of Informationally efficient markets," *American Economic Review* 70: 393–408.
- [22] Hakansson, Nils H., J. Gregory Kunkel and James A. Ohlson (1982), "Sufficient and Necessary Conditions for Information to have Social Value in Pure Exchange," *Journal* of Finance 37(5): 1169–1181.
- [23] Hayek, F. A. (1945), "The Use of Knowledge in Society," American Economic Review 35(4): 519–530.
- [24] Hirshleifer, Jack (1971), "The private and social value of information and the reward to inventive activity," *American Economic Review* 61(4): 561–574.
- [25] Holden, Craig and A. Subrahmanyam (1992), "Long-lived Private information and imperfect competition," *Journal of Finance*, 47(1): 247–270.
- [26] Hollifield, B., R. Miller, P. Sandås, and J. Slive (2004), "Estimating the gains from trade in limit order markets," Working Paper, University of Pennsylvania.
- [27] Jackson, Matthew O. (1991), "Equilibrium, Price Formation, and the Value of Private Information," *Review of Financial Studies* 4(1): 1–16.

- [28] Kaniel, Ron and Hong Liu, "So What Orders Do Informed Traders Use?," forthcoming, Journal of Business.
- [29] Kyle, Albert (1985), "Continuous Auctions and Insider Trading," Econometrica 53: 1315–1336.
- [30] Maskin, Eric and Jean Tirole (2001), "Markov Perfect Equilibrium I: Observable Actions," Journal of Economic Theory 100: 191–219.
- [31] Mendelson, Haim and Tunay Tunca (2004), "Strategic Trading, Liquidity, and Information Acquisition," *Review of Financial Studies* 17(2): 295–337.
- [32] Pakes, Ariel and Paul McGuire (2001), "Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'Curse' of dimensionality," *Econometrica* 69(5): 1261– 1281.
- [33] Parlour, Christine A. (1998), "Price Dynamics in Limit Order Markets," Review of Financial Studies 11: 789–816.
- [34] Perry, Motty and Philip Reny (2004), "Towards a Strategic Foundation for Rational Expectations Equilibrium," Working paper, University of Chicago.
- [35] Rosu, I. (2004), "A Dynamic Model of the Limit Order Book," Working paper, University of Chicago.
- [36] Schlee, Edward, E., 'The Value of Information in Efficient Risk Sharing Arrangements, American Economic Review 91(3): 501–524.
- [37] Spiegel, Matthew and Avanidhar Subrahmanyam (1992), "Informed Speculation and Hedging in a Noncompetitive Securities Market," *Review of Financial Studies* Vol 5., No 2, 307-329.
- [38] Taub, Bart, Dan Bernhardt and P. Seiler (2004), "Cladistic Asset Pricing," University of Illinois working paper.
- [39] Verrecchia, Robert (1982), "Information Acquisition in a Noisy Rational Expectations Economy," *Econometrica* 50: 1415–1430.
- [40] Watkins, C.J.C.H. (1989) "Learning from Delayed Rewards," Ph.D. thesis, Cambridge University.