Information Acquisition in a Limit Order Market*

Ronald L. Goettler† Christine A. Parlour‡ Uday Rajan§

Tepper School of Business
Carnegie Mellon University
Pittsburgh, PA 15213

Ross School of Business
University of Michigan
Ann Arbor, MI 48109

December 21, 2004

*We have benefitted from the comments provided by Ekkehart Boehmer and seminar participants at Aladdin (2004), CFS (Eltville), Ross, Tepper, Texas A&M, and the 2004 Oxford Finance Summer Symposium. The current version of this paper is maintained at http://ozymandias.tepper.cmu.edu

†Tel: (412) 268-7058, E-mail: ronald.goettler@cmu.edu

‡Tel: (412) 268-5806, E-mail: parlourc@andrew.cmu.edu

§Tel: (734) 764-2310, E-mail: urajan@umich.edu

Information Acquisition in a Limit Order Market

Abstract

We model endogenous information acquisition in a limit order market for a single financial asset. 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. We numerically solve for the equilibrium of the trading game, and characterize equilibria with endogenous information acquisition. Over a range of information acquisition costs, the game exhibits a prisoner's dilemma—all agents, including those who acquire information, are worse off. Agents with the lowest intrinsic benefit from trade have the highest value for information and also tend to supply liquidity. As a result, market observables such as bid and ask quotes, in addition to transaction prices, are informative about the common value of the asset. Adverse selection is important for individuals (agents have lower payoffs when uninformed), but in the aggregate it has little effect on investor surplus, unless gains to trade are small. Comparisons to a frictionless benchmark show that the limit order market is effective at consummating trade and generating consumer surplus, even in the presence of asymmetric information.

1 Introduction

Hayek (1945) stresses the role of the price system as a "mechanism for communicating information" about the values of different goods. But how is information incorporated into price? Consider the market for a financial asset: For the price of the asset to convey all available information about its value, some traders must voluntarily acquire this information and then trade in a manner that reveals the information to the market. However, if some traders have superior information about the asset's value, the other agents in the market are faced with adverse selection. Such adverse selection may impede trade, leading to a fall in aggregate welfare. Therefore, a tradeoff may exist between the allocative and informational efficiency of a market.

To address the effects of adverse selection on informational and allocative efficiency, it is important to model a market explicitly. The rules of the market determine how an agent can benefit from superior information, as well as how an uninformed agent can learn from market observables and respond to adverse selection. Traders' strategies, in turn, affect the degree to which information acquired by a subset of traders is reflected in prices and other market observables. Finally, agents' allocations and payoffs also depend on the market form.

We model information acquisition by strategic traders in a dynamic limit order market for a single financial asset. 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 orders to the limit order book or choose to trade against previously posted orders. We consider a market with an open limit order book; that is, traders observe every order in the book.

We choose this market form for both theoretical and practical reasons. Theoretical asset pricing theory typically assumes competitive financial markets. Blume and Easley (1990), however, 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.

On the practical front 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. The electronic format is especially conducive to limit order books (since it is easy to keep track of all orders and display the book and history). Indeed, many recent electronic exchanges maintain and display a limit order book. Glosten (1994) finds that a limit order market is resilient in that it deters entry by competing market forms.

These markets have at least three important frictions, including two sources of asymmetric information. 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.

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. Prior to his first entry into the market, each agent chooses whether to buy information. 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 a Markov-perfect equilibrium. We then use the payoffs from the trading game to determine the value of information to each trader type and identify equilibria in the overall game with endogenous information acquisition. We find that traders with a low intrinsic motive for trade are willing to pay the most for information. For some range of costs, information acquisition represents a prisoner's dilemma: all agents in the market are worse off when some set of agents optimally acquire information.

We consider the allocative and informational efficiency of the market. Asymmetric information is important for individual traders: It transfers surplus from uninformed agents to informed ones. However, it has a negligible impact on aggregate welfare: The gross surplus per trader in the trading game (i.e., without accounting for information acquisition cost) is approximately the same, regardless of the number of informed agents. More broadly, we find the limit order market is an effective incentive compatible mechanism for consummating trade: it achieves 92.1% of a frictionless benchmark (which ignores all frictions, including the sequential arrival of agents).

The informational efficiency of the market unambiguously improves when there are a greater number of informed agents. 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. Uninformed agents' beliefs about the common value are less precise when fewer agents are informed.

Two features of the open limit order market aid its success as a trading mechanism. First, uninformed traders can learn from current and past market observables. 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 (so the probability of trading with one is low) observables are less informative about the common value. Adverse selection itself, however, is less of an issue in this case. Thus, the transparency of the limit order market mitigates the adverse selection risk faced by the uninformed.

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.

Hirshleifer (1971) observes that, in an exchange economy with risk-averse traders, information has no social value: if all agents are informed, risk-sharing opportunities are eliminated and the market breaks down.[1] A key difference in our approach is that our traders are risk neutral, and the potential gains from trade are fixed across different information acquisition regimes. We do this because information pertaining to idiosyncratic variation in cash flows should not affect the gains to trade if agents already hold well-diversified portfolios. Adjustments to an agent's holdings of a particular asset, therefore, are primarily motivated by liquidity needs, as opposed to risk-sharing ones. Thus, the size of the potential gains to trade should not depend on the degree of adverse selection. We show that the realized aggregate surplus in the limit order market depends on the size of the potential gains to trade relative to the volatility in common value: when gains to trade are small, adverse selection can affect aggregate welfare.

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.[2] Thus,

[1] Hakansson, Kunkel, and Ohlson (1982) demonstrate that if (i) the market is not allocationally efficient and (ii) information structures are not homogeneous across investors, information can have social value. 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.

[2] For example, Jackson (1991) demonstrates that the price-taking assumption is critical in order to sustain

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.

Our framework is a fully strategic trading game. 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.[3] Thus, this framework is inappropriate for examining how public information about (say) earnings gets incorporated into price. Further, as all trades are consummated at the market maker's quoted prices, there is no distinction between quotes and transaction prices.

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.

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.

the Grossman–Stiglitz paradox.

[3]Taub, 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.

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 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.

Our paper links the literature on information acquisition to that on dynamic limit order markets. The latter includes Rosu (2004), who presents a continuous time model of a limit order market with no common value. His solution technique requires 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 private and common values. They assume that cancellations are exogenous: agents do not revisit the market and thus do not solve a truly dynamic problem. Further, their model is in discrete time. By contrast, the current framework is in continuous time and agents may revisit the market.[4]

Our work also adds to the computational literature. First, we show that the stochastic approximation algorithm of Pakes and McGuire (2001) extends naturally to models of asymmetric information, where agents have private state variables. Second, 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. In Section 3 we use the computed payoffs from the trading game to solve for the equilibrium information structures. We then analyze agents' order submission strategies across different information acquisition structures (Section 4). 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 6. In Section 7, we provide some comparative statics on informational and allocational efficiency. 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

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

endogenous information acquisition in Section 3. 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, β. 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 countable set of discrete prices, $\mathcal{P} = \{p^i\}_{i=-\infty}^{\infty}$, at which traders may submit orders. The distance between any two consecutive prices is a constant, d, and we refer to it as "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. We sometimes refer to this as 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 or transparent limit order book, with no hidden orders.[5] 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 $A(L) = \infty$.

New traders arrive at the market according to a Poisson process with parameter λ_N. Hence, the actual time between trader arrivals is random. Although arrival at the market is exogenous, participation in the market is completely endogenous—a trader may choose to not submit an order.

Each trader has a type denoted by $\theta = \{\rho, \beta\}$, with Θ being the set of types. Here, ρ is a continuous discount rate. The payoff a trader earns as a result of trading is discounted back to his first arrival time in the market at this rate. This captures the notion that traders prefer to execute sooner rather than later, and prevents a trader from infinitely postponing trade. The discount rate is interpreted as the cost of delaying trade, which 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,

[5] In practice, a small proportion of limit orders on many limit order markets are "hidden" (i.e., not revealed to other traders), by choice of the order submitter.

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.[6]

Each trader has a private value for the asset, β. The private value represents private benefits to trade as a result of liquidity shocks or private hedging needs. Its presence implies potential gains to trade among agents. We consider β distributions that are symmetric, have a mean of zero, and have finite support. Let F_β denote the distribution of β. The private value β 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 interpreted as 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 one tick each with probability $\frac{1}{2}$. Changes in the common value reflect new information about the firm.

On his first entry to the market, an agent may choose to buy information by paying a cost $c \geq 0$. 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.[7] 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.

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. Informed agents know the current value of v on each entry into the market. Uninformed agents view v with a time lag, Δ_t, measured in units of time. That is, an 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.

An agent in the market at time t observes the limit order book, L_t, and the appropriate value of v (either v_t or $v_{t-\Delta_t}$). He also observes some information about market events in the

[6] The 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.

[7] In 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.

interval $[t-\Delta_t, t]$, specifically 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, $b_t \in \{-1, 1\}$. Let $m_t(I)$ denote the set of market-related variables observed by an agent (we sometimes refer to m as "the market m"). Then,

$$
\begin{aligned}
m_t(0) &= \{L_t, v_{(t-\Delta_t)}, \hat{p}_t, b_t\}, \\
m_t(1) &= m_t(0) \cup \{v_t\}.
\end{aligned}
$$

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.[8] The variables in $m_t(0)$ offer 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). 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.

Given the payoffs in the trading game for different information acquisition strategies, we can construct a payoff matrix for the information acquisition game for a given value of information acquisition cost, c. We then compute the change in expected payoff for each type from deviating at the information acquisition stage. This provides bounds for c for which the assumed information acquisition strategies constitute an equilibrium.

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

[8] 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. 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.

2.1 Continuation Trading Game

Each trader is allowed to trade exactly 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. Let G denote the distribution over reentry time with g being the associated density. 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, $q \geq 0$ the priority of his order among all orders at price p, and

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

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 market order. 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. All 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

[9]For 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.

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 $\theta = (\rho, \beta)$, the instantaneous utility at time t is

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

The expected payoffs to different actions depend on a trader's information set. Let $s = \{\theta, m_t(I), a, z\}$ be the state a trader observes on a particular entry to the market at time t. Here, $a = (p, q, x)$ denotes the status of the trader's previous order. Note that the priority of his previous order (q) may have changed from the time it was submitted. The state includes z, the number of shares the agent may continue to trade (this is either 0 or 1 by assumption). When $z = 0$ (i.e., after he has traded one share), the trader's continuation payoff is set to zero. Recall from the definition of m that informed and uninformed agents see different market variables (indeed, a given m indicates whether the agent is informed). Hence, these traders have different state variables, and the state space for an informed trader is a strict superset of the state space for an uninformed one.

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 a state is denoted as a 4-tuple $s = \{\theta, m(I), a, z\}$. Thus, $\mathcal{A}(s)$ denotes the feasible action set for a trader with type θ who can trade z shares when the market is given by $m(I)$ and $a = (p, q, x)$ is the status of his previous action (if he is a new trader, we set $x = 0$). With a slight abuse of notation, for any state s let ℓ^p denote the outstanding limit orders at price p. For a new order, its priority $\hat{q}(p, x)$ is determined as follows:

$$\hat{q}(p, x) = \begin{cases} 0 & \text{if (i) } x = 0, \text{or (ii) } x = 1, p \geq A(L), \text{or (iii) } x = -1, p \leq B(L) \\ |\ell^p + x| & \text{otherwise.} \end{cases} \quad (3)$$

A buy order at a price $p \geq A(L)$ automatically executes at $A(L)$, and similarly with a sell order at $p \leq B(L)$.

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. Denote this expectation as $\hat{v}(m) = E(v \mid m)$, where m denotes the market conditions observed by the agent. The feasible action set is then defined as

$$\begin{aligned} \mathcal{A}(s) \quad = \quad \{ \ (p,q,x) \ \mid \ \ & \text{(i) } x \in \{-1,0,1\}, \ \text{(ii) } q = \hat{q}(p,x), \\ & \text{(iii) } q \neq 0 \Longrightarrow p \in [\hat{v}(m)-k, \hat{v}(m)+k] \cap \mathcal{P} \ \}. \end{aligned} \tag{4}$$

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 $\Theta^I(\sigma_I) = \{(\theta, I) \in \Theta \times \{0,1\} \mid \text{Prob}((\theta, I) \mid \sigma_I(\theta)) > 0\}$ 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 faces the market conditions given by m, 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 payoff-maximizing order depends on both these outcomes. 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 market m, 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 market conditions m'.

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 $\phi(t, v; m, \tilde{a}, \sigma)$ be the probability that an action $\tilde{a} = (\tilde{p}, \tilde{q}, \tilde{x})$ taken in market m at time 0 leads to execution at time $t > 0$ when the common value is v, given that all other agents are playing σ, and let $f(v \mid m, t)$ denote the density function over v at time t, given market m. For an informed trader, $f(\cdot)$ is purely exogenous; 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

$$\pi(s, \tilde{a}, w, \sigma) = \int_0^w \int_{-\infty}^{\infty} \left(e^{-\rho t} \tilde{x} \; (\beta + v_t - \tilde{p}) \; \phi(\cdot)\right) f(v \mid m, t) \; dv \; dt. \tag{5}$$

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 $h(s' \mid s, \tilde{a}, w, \sigma)$ 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

$$J(s, \sigma) = \max_{\tilde{a} \in A(s)} \int_0^{\infty} \left\{ \pi(s, \tilde{a}, w, \sigma) \; + \; e^{-\rho w} \int_{s' \in S_\theta} J(s', \sigma) \; h(s' \mid s, \tilde{a}, w, \sigma) ds' \right\} dG(w). \tag{6}$$

The first term on the right-hand side 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' possibly 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, the market conditions may have changed to m' 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. Recall that G is the distribution of reentry time; the outermost integral is over this random reentry time.

Each time a trader is in the market, he chooses a payoff-maximizing action; that is, 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.

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 $s \in S_{(\theta,I)}$,

$$y^*_{(\theta,I)}(s) \in \arg\max_{\tilde{a} \in A(s)} \int_{w=0}^{\infty} \left\{ \pi(s, \tilde{a}, w, \sigma) \; + \; e^{-\rho w} \int_{s' \in S_\theta} J(s', \sigma) \; h(s' \mid s, \tilde{a}, w, \sigma) ds' \right\} dG(w). \quad (7)$$

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^*.

Formally, we have a Bayesian game. Traders have privately known utilities from trade (since a trader's β 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 β 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. Therefore, even though this is a Bayesian game with some unobservable actions, we refer to our equilibrium as a Markov-perfect equilibrium.[10] We solve for a stationary, symmetric Markov-perfect equilibrium. 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.

2.2 Existence of Equilibrium in the Trading Game

Given the state space we have defined (and subject to the earlier caveat of states and possibly state spaces differing across agent types), the equilibrium concept we use in the trading game is stationary Markov-perfect equilibrium. The existence of a Markov perfect equilibrium follows from standard results. The set of players (i.e., new traders) is countable. On each entry the action space for a trader is finite. Further, the state space is countable. The state changes as a result of either changes in the common value or actions taken by traders; each occurs at most a countable number of times over an infinite horizon. It then follows from the theorem of Rieder (1979) that a Markov-perfect equilibrium exists. Since the time at which a trader enters the market is unimportant, given his state and the status of his previous action, this equilibrium is stationary.

Given information acquisition strategies, the equilibrium of the trading game appears to be computationally unique. In Section 3 we show that, despite this, there are cost ranges that lead to multiple equilibria in the information acquisition game.

2.3 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 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

[10] It should be understood that we solve for optimal Markov strategies conditional on the state space we have defined. The state space incorporates specific restrictions that 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.

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).

As discussed, we fix information acquisition strategies $\sigma_I(\theta)$ and solve for the equilibrium of the corresponding trading game. Equilibrium is obtained by finding a $J(s, \sigma)$ that satisfies the Bellman equation in (6). In equilibrium, when each trader plays his best response, $y^*_\theta(s)$, the means of the distributions of realized outcomes will indeed match the expected values for these outcomes, as specified by $J(s, \sigma)$. That is, traders' beliefs about the expected discounted value of each action are consistent with the actual distribution of outcomes.

To find this fixed point using traditional value function iteration is computationally prohibitive given the integrals in equations (5) and (6). In particular, the probability and density functions ϕ, f (for uninformed agents), and h that determine the evolution of the state are complicated posteriors that depend on endogenous strategies.

Instead, we simulate a market session and update beliefs (about the value of each action) 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.[11] 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

[11] Pakes 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.

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, 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.[12] 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. Of course, the probability of trembles must be small enough not to affect the strategies along the equilibrium path. To retain the notion that $J(\cdot)$ corresponds to optimal future play, updates to $J(\cdot)$ always use the expected utility of the optimal action available, even if the trader is randomly selected to tremble. Of course, traders will respond to the possibility that other agents may tremble. To minimize this effect, 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 β and possibly v) are privately known to agents. A substantive difference is the use of trembles to ensure perfection.[14]

[12] A 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.

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

[14] While 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.

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

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 β 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 $\beta \in \{-0.1, 0.1\}$ 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 $\beta \in \{4, 8\}$ are likely to be buyers overall, and those with $\beta \in \{-4, -8\}$ 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. Our parameterization of F_β 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. Hence, we choose 4 ticks as the private value for our second kind of trader, and 8 ticks for our third kind.

- F_v, the distribution for changes in common value, is a Poisson distribution. The inter-arrival time of innovations to common value is therefore exponential. The expected time between changes in v is 8 units of real time.

  Again, our 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 vary Δ_t in our simulations and report the sensitivity of our results to changes in Δ_t.

- 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

[15]The variance of the innovation in an interval of time t is the expected number of innovations in the interval.

[16]We 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.

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.

In addition to the above parameterization, we also condense the limit order books into a smaller set of variables. In principle we would like agents to condition their strategies on the entire limit order book. In practice this yields too many states to be computationally tractable. Hence, as a simplification, we allow agents to use: (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 test runs with different parameterizations, where the state space using the full book was smaller, we verified that equilibria and market outcomes are essentially unaffected by this condensation. The algorithm's insensitivity to this condensation reflects the fact that the states being combined indeed have similar values for $J(\cdot)$.

3 Endogenous Information Acquisition

As a first step to determining equilibria with endogenous information acquisition, we fix agents' information acquisition strategies and solve for equilibrium in the trading game. We consider regimes in which all agents with a given β 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 $|\beta|$. Hence, we report results from four information acquisition regimes: all agents are informed,

[17] We 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.

only agents with $|\beta| \in \{0.1, 4\}$ are informed, only agents with $|\beta| = 0.1$ 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.

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 1. All payoffs are quoted in ticks. We exhibit the mean and sample standard deviation of 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 β are willing to pay for information.[18]

Figure 1 shows the value of information to each type across information structures. The value of being informed decreases in the absolute value of β. Information is most valuable to the speculators (i.e., agents with $|\beta| = 0.1$), 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 $\beta = 8$ are unlikely to switch from buyers to sellers, so information is less valuable to these agents.[19]

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.

[18] Note that we have a large sample of 100 million trader decisions (approximately 36 million traders). As mentioned in the notes to Table 1, standard errors on equilibrium payoffs are less than 0.0005. Since only 2% of simulated traders are deviators, we base deviator payoffs on a larger sample with 300 million trader decisions. This reduces the standard errors in deviator payoffs to less than 0.0020. Any differences in payoffs on the order of 0.01 or higher are statistically significant.

[19] Radner and Stiglitz (1984) demonstrate that information is valuable to a single Bayesian decision-maker only if it induces a change in action.

Information Structure			Value of $\|\beta\|$ 0.1	4	8
None Informed	Equilibrium	Mean	0.403	3.515	7.333
		Std. Dev	0.968	1.499	1.607
	Deviation	Mean	*1.178*	*3.652*	*7.353*
		Std. Dev	0.772	1.342	1.558
	Value of Information	Mean	0.776	0.138	0.020
Speculators Informed	Equilibrium	Mean	*0.628*	3.413	7.204
		Std. Dev	0.662	1.040	1.137
	Deviation	Mean	0.413	*3.499*	*7.228*
		Std. Dev	0.779	0.955	1.072
	Value of Information	Mean	0.215	0.086	0.024
$\beta \in \{.1, 4\}$ Informed	Equilibrium	Mean	*0.495*	*3.508*	7.234
		Std. Dev	0.542	0.792	0.965
	Deviation	Mean	0.287	3.414	*7.279*
		Std. Dev	0.587	0.877	0.864
	Value of Information	Mean	0.207	0.094	0.044
All Informed	Equilibrium	Mean	*0.447*	*3.510*	*7.311*
		Std. Dev	0.469	0.766	0.855
	Deviation	Mean	0.244	3.430	7.251
		Std. Dev	0.459	0.804	0.895
	Value of Information	Mean	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 1: **Payoffs (in ticks) in different information structures**

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

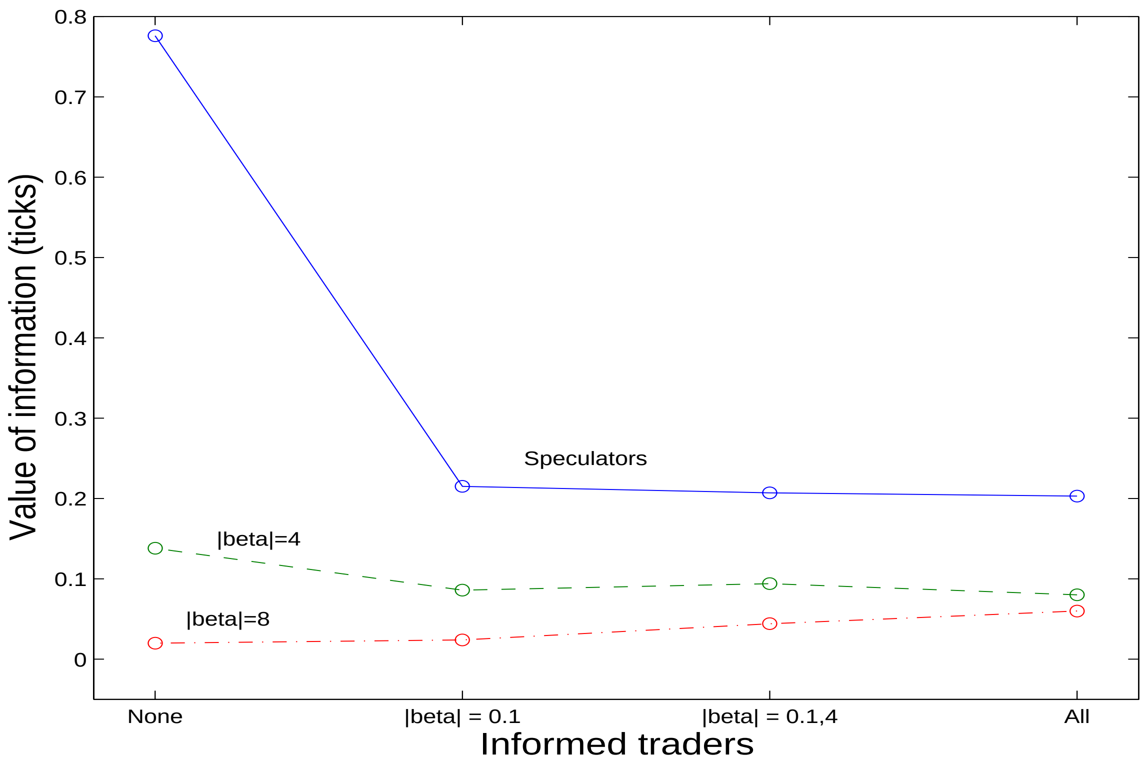

Figure 1: **Value of information**

Figure 1 also shows that the value of information depends on which other agents are informed. One may expect that increased competition among informed agents (in the sense of a larger number of agents who are informed) reduces the benefit of information. This is true for the speculators (most strikingly, a speculator can improve her payoff by 0.776 if she acquires information when no one else does), and in some cases for the traders with $|\beta| = 4$. The extreme β traders (those with $|\beta| = 8$), however, are willing to pay more for information when other agents are also informed.

Admati and Pfleiderer (1987), in a multi-asset model, and Barlevy and Veronesi (2000), in a single-asset framework with exponentially distributed signals, both show that the value of an additional signal to agents in a noisy rational expectations exchange economy can be increasing in the number of other signals they have. In their respective settings, signals can be substitutes or complements in equilibrium, depending on the underlying parameters. The value of information (i.e., additional signals) to agents is increasing when signals are complementary. In such a case, the price in turn is a poorer predictor of the true value of the asset.

In our framework, signals of different agents are substitutes. However, since only one agent is in the market at any given point of time, agents have local monopoly power. The monopoly power of an informed agent depends on the proportion of other agents who are informed. The degree of substitutability of signals also varies according to agents' trading strategies. In our model, as we show below, prices (and more broadly, market observables) are unambiguously more revealing when more agents are informed. Nevertheless, the value of information to agents may increase in the number of other agents who are informed, depending on the degree of substitutatibility of signals.

Now, consider endogenous information acquisition. 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 $|\beta| = 8$ are informed.

Different information acquisition equilibria arise, depending on the cost of acquiring information.[20] As mentioned, we only characterize equilibria in which all agents with the same (ρ, β) take the same action in the information acquisition game. For some cost ranges (e.g., $c \in (0.215, 0.776)$), the equilibrium of the information acquisition game does not exhibit this feature.

Observation 2 *The following are equilibria in the information acquisition game:*

$$c \in \begin{cases} [0, 0.060) & \text{all agents acquire information} \\ (0.044, 0.094) & |\beta| \in \{0.1, 4\} \ \text{ acquire information} \\ (0.086, 0.215) & |\beta| = 0.1 \ \text{ acquire information} \\ (0.776, \infty) & \text{no agent acquires information} \end{cases}$$

Consider the information structure in which only the speculators are informed: their average gross payoff in this case is 0.628. Thus, a cost of 0.086 (which allows this regime to be an equilibrium) represents 13.7% of this average payoff, and a cost of 0.215 (the maximum cost that supports this equilibrium) is 34.2% of the average payoff. In this information structure, the average gross payoff of the $|\beta| = 8$ traders is 7.204. Thus, a cost of 0.086 is 1.2% of their average payoff, and a cost of 0.215 represents 3.0% of their average payoff.

It is immediate from the previous observation that the cost ranges that support different equilibria overlap. In other words, for the same cost, different equilibria are possible.

Observation 3

(a) For $c \in (0.044, 0.060)$ *there are two equilibria:*

(i) $|\beta| \in \{0.1, 4\}$ *acquire information.*

[20] One may reasonably suppose that this cost differs across securities. For example, some stocks are carefully followed by analysts, making pertinent information available at a low cost.

(ii) All agents acquire information.

(b) For $c \in (0.086, 0.094)$*, there are two equilibria:*

(i) Only speculators acquire information.

(ii) $|\beta| \in \{0.1, 4\}$ *acquire information.*

Multiple equilibria in the information acquisition game occur because agents behave differently in the continuation trading game as the number of informed agents changes. Consider an information acquisition cost of 0.05. For this cost, there is an equilibrium in which all agents acquire information. There is also an equilibrium in which only the agents with $|\beta| = \{0.1, 4\}$ acquire information. In the latter equilibrium, the agents with $|\beta| = 8$ are uninformed. However, they know that the other agents in the market are informed, and choose their trading strategies accordingly. To some extent, the change in strategy is a substitute for acquiring information, allowing for multiple equilibria.

The existence of multiple equilibria potentially presents problems for cross-sectional comparisons across assets. For example, two stocks may have similar information acquisition costs, but operate in markets with different equilibria, so that they exhibit different degrees of asymmetric information and thus different trading patterns. Conversely, the intuition that stocks with markedly different information acquisition costs should exhibit different trading patterns may not be borne out in practice if the cost range that supports a given equilibrium is wide.

In the information acquisition game, agents face the prisoner's dilemma for some cost ranges: relative to the regime in which no agents are informed, all agents (including those who acquire information in equilibrium) are strictly worse off. It is intuitive that an uninformed agent in a market in which others have acquired information will be worse off, since he faces a higher degree of adverse selection. However, given a positive cost of information, it is also possible for the informed agents to be worse off in equilibrium (though it is a best response for these agents to acquire information).[21]

Observation 4 *For* $c \in (0.044, 0.060)$*, the equilibrium in which all agents are informed reflects a prisoner's dilemma—each type of agent would prefer to be in the regime in which no agents are informed.*

[21] Bassan, Gossner, Scarsini, and Zamir (2003) provide conditions under which all agents prefer more information. If a game with a particular information structure has unique Pareto efficient payoffs, then in all coarser information partitions more information is preferred. Pareto efficient payoffs in our trading game are not unique (since small redistributions from extreme β agents to speculators, for example, also lead to Pareto efficiency). Further, the cost of acquiring information is important in ensuring that agents who purchase information are also worse off as a result.

The dissipative nature of information acquisition in our model differs from the celebrated Hirshleifer effect. Hirshleifer (1971) shows that risk-averse agents may be worse off with more information (i.e., information about the realization of the state).[22] Essentially, adverse selection can lead to possible market breakdown: the gains to risk sharing disappear if some agents are fully informed.[23] By contrast, in our model the gains to trade are represented by the private value distribution, which is held constant across different regimes. The volume of trade remains approximately the same, regardless of information acquisition.[24] However, uninformed traders' strategies change in response to adverse selection, limiting the increase in gross profits to the informed traders. Once the cost of information acquisition is factored in, informed traders also have a lower payoff.

4 Effect of Information on Trading Behavior

In this section, we consider the impact of information and adverse selection on each trader type separately. Two of our information structures (no agents informed and all agents informed) represent cases with no adverse selection. Since it takes a relatively large cost (greater than 0.776 ticks) to sustain an equilibrium in which no agents are informed, we omit further discussion of this case.

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 payoff to an agent with private value β 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 $e^{-\rho t}\,(\beta - (p - v))$.

First, consider the time to execution. In Table 2, we report for each trader type the average time from entry to the market until trade is consummated. For the case in which only speculators informed, we also report the average time to execution for agents who deviate at the information acquisition stage. Speculators who deviate in this case are uninformed, whereas agents with $|\beta| = 4, 8$ who deviate are informed.

Table 2 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. Across information regimes, they execute fastest when they are the only ones informed. If they compete with another group of informed agents, it takes them longer to

[22]Schlee (2000) extends this result to an economy in which some agents are risk neutral.

[23]Building on the Hirshleifer effect, Berk (1997) shows that if markets are dynamically complete, information acquisition can make all agents worse off.

[24]In a stationary equilibrium, agents must leave the market at approximately the same rate at which they arrive. Since a new trader arrives every period, and it takes two traders to transact, this implies that on average a transaction must occur every two periods.

Informed Agents	Value of $\|\beta\|$ 0.1	4	8
$\|\beta\| = 0.1$	17.22	2.19	0.62
Deviators	25.73	2.46	0.58
$\|\beta\| \in \{0.1, 4\}$	19.41	2.13	0.56
All	19.64	2.09	0.59

Table 2: **Average time to execution**

execute: competition among informed agents makes profitable trades more difficult to find. Although the extreme β agents always execute quickly, they take the longest when only the speculators are informed.

The effects of information acquisition on time to execution depend on agent type. Again, consider the case in which only speculators are informed (top two rows of Table 2). A speculator who remains uninformed (who is a deviator in this case) takes significantly longer to execute (25.73 time units) compared to informed ones (17.22 units). However, among the agents with $|\beta| = 4$, informed traders take longer to execute. When faced with greater adverse selection, uninformed traders with extreme private values prefer to trade quickly (0.58 time units, rather than 0.62 time units).

Very few agents in our simulation submit no order given the opportunity to do so. Therefore, agents who execute more slowly are the ones who submit more limit orders. In Table 3, we report the average number of new orders submitted by agents with a positive β in the informed speculators case. The orders submitted in other information regimes and by traders with a negative β display a similar pattern.

A trader who executes instantaneously on arrival will submit only one order. Thus, the fact that agents with $\beta = 8$ submit 1.06 orders on average indicates that they execute quickly.

Table 3 shows that each speculator submits more limit orders than each agent of any other type. As discussed in the previous section, these traders also have the highest value for information. Over and above new order submission, we find that speculators choose to retain an existing limit order on approximately 50% of their entries into the market. In terms of eventual execution, only about 20% of all speculator executions occur via market orders. The table also shows that speculators are willing to take either side of a trade, unlike the other two types.

Comparing the equilibrium and deviator columns for each type, we find that, all else

	Equilibrium			Deviators		
	Value of β			Value of β		
Order type	0.1	4	8	0.1	4	8
Limit buy	1.66	0.66	0.29	2.22	0.68	0.26
Market buy	0.11	0.51	0.77	0.12	0.54	0.81
Limit sell	0.67	0.00	0	1.19	0	0.00
Market sell	0.11	0	0	0.06	0	0
Total orders per trader	2.54	1.17	1.06	3.58	1.22	1.07

Table 3: **Average number of orders submitted when only speculators are informed**

equal, uninformed speculators submit a greater number of limit orders. That is, they appear to take longer to identify market opportunities when uninformed. However, if endogenous information acquisition is considered, the 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).

Next, we consider execution prices relative to the common value. In Table 4, we report $(p - v)$ for all executed orders submitted by agents with $\beta > 0$ (the pattern for agents with $\beta < 0$ is symmetric).[25]

Informed	$\beta = 0.1$		$\beta = 4$	$\beta = 8$
Agents	Buy Orders	Sell Orders	Buy Orders	Buy Orders
Speculators:				
Equilibrium	−1.21	1.30	0.21	0.57
Deviators	−1.30	1.18	0.00	0.62
$\lvert\beta\rvert \in \{0.1, 4\}$	−1.17	1.19	0.11	0.57
All agents	−1.14	1.07	0.10	0.48

Table 4: **Average of (price minus common value) for executed orders, in ticks**

Two comparisons are useful from the table. First, comparing across information regimes (i.e., the first, third, and fourth rows of numbers), speculators experience an improvement in the terms of trade when only they are informed. On average, their buy orders execute about 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. The effects of adverse selection

[25]Since our model is symmetric, the mean of $(p - v)$ across all trades is zero in all information regimes (though there is dispersion that varies with information structure).

are exhibited by the increased cost (in terms of amount paid in excess of the common value) paid by agents with $\beta = 4$ in the case in which only the speculators are informed. Similarly, the prices paid by the extreme β traders (i.e., those with $\beta = 8$) are highest in the first two cases, in which these agents are uninformed.

Another way to examine the effects of adverse selection is to look at agents who deviate on information acquisition in a given information structure. Examined in this way, the effects of information on terms of trade are mixed. Consider the first two rows of numbers in Table 4. Speculators who deviate (and are thus uninformed) execute orders at better prices than those who remain informed. For example, an uninformed speculator executes buy orders at an average of 1.30 ticks below the common value, as opposed to 1.21 ticks below common value for an informed speculator. Similarly, agents with $\beta = 8$ execute at better terms of trade when uninformed. It is odd to think that an uninformed agent can secure better terms of trade than an informed one. However, note that this is but one component of payoff. As shown before, speculators and extreme β agents take longer to execute when uninformed, so that their payoff is lower. Conversely, agents with $\beta = 4$ experience worse terms of trade when uninformed, trade more quickly than when they are informed.

Thus, overall we find that uninformed speculators take longer to execute, since it takes longer for them to identify profitable trading opportunities. However, the overall effect on their terms of trade is mixed. Information is thus valuable to them since it allows them to trade more quickly.

5 Overall Allocative Efficiency

We now consider overall investor surplus generated by the limit order market. Our measure of aggregate welfare in a given market is the mean surplus per trader.

Allocations in this market are approximately invariant to the information structure. Agents with $\beta = 4$ or 8 buy a share, and those with $\beta = -4$ or -8 sell a share. Among speculators, roughly $\frac{2}{3}$ of agents with $\beta = 0.1$ buy a share, with the remainder (roughly $\frac{1}{3}$) selling a share. These fractions are reversed for the $\beta = -0.1$ agents. Any effects of adverse selection on overall surplus, therefore, do not occur via a difference in allocations.[26] Further, as Table 2 indicates, execution times do not vary systematically across information regimes: when speculators take longer to execute, agents with $|\beta| \in \{4, 8\}$ trade more quickly. What

[26]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. This contrasts with the equilibrium when trading is centralized (so that all trades must occur at the same price and same point of time).

is the overall effect of adverse selection on aggregate welfare?

We consider two benchmarks for aggregate 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.[27] Let W_f be the surplus (or welfare improvement) per trader in a frictionless market. The surplus each agent gets if he trades instantaneously at v is $|\beta|$. Thus, given the probability distribution of types, $W_f = 4.03$.

A second-best benchmark for welfare must account for the friction introduced by incentive compatibility. In our framework, the optimal mechanism subject to incentive compatibility remains an open question. Instead, we consider one mechanism that is incentive compatible, that works according to a LIFO (last-in-first-out) rule. Suppose the planner executes all trades at $p^* = v$. However, the planner respects the arrival times of agents, and discounts accordingly. A simple mechanism is as follows: all traders with $\beta > 0$ are designated buyers, and those with $\beta < 0$ are designated sellers. As soon as at least one buyer and one seller are present, a trade takes place between the buyer and seller most recently arrived to the market.

Given a fixed arrival process, this mechanism is incentive compatible—no agent can gain by misreporting his β. The LIFO rule prevents excessive discounting; by comparison, a FIFO (or first-in-first-out) allocation rule performs much worse. Let W_ℓ be the surplus per trader per trader by such a LIFO mechanism. This measure is straightforward to compute given a particular arrival sequence of traders. We compute W_ℓ from a simulated sequence of traders that is sufficiently long for this welfare benchmark to have a negligible standard error.

	Frictionless: W_f	LIFO: W_ℓ	All Informed	$\lvert\beta\rvert = \{0.1, 4\}$ Informed	$\lvert\beta\rvert = 0.1$ Informed	None Informed
Gross surplus	4.030	3.482	3.734	3.724	3.718	3.730
Net surplus	4.030	3.482	$3.734-c$	$3.724-0.7c$	$3.718-0.3c$	3.730

Table 5: **Welfare gain per trader, and benchmarks**

The aggregate welfare per trader in each market, and the benchmarks, are shown in Table 5. Neither of the two benchmarks are affected by information acquisition (since there is no incentive to acquire information in either the frictionless or the LIFO mechanisms).

[27] All prices between -0.1 and 0.1 represent Walrasian equilibria given our parameters; allocations are invariant across these equilibria.

For ease of comparison, we report both the gross and net (after deducting information acquisition costs) surplus numbers from the four markets under consideration.

As expected, in all cases the gross surplus in the market is less than the frictionless benchmark W_f. However, even in the worst case (when only agents with $|\beta| = 0.1$ are informed), it recovers 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. Since their model has exogenous cancellations, the realized gains to trade may be a little underestimated in their empirics.

The market as a whole outperforms the LIFO benchmark in each case, by about 6.7%. Although not reported in the table, it is remarkable that, in each information structure, the gross surplus accruing to each trader type is higher than in the LIFO case. In our market, traders with extreme β values have an incentive to submit aggressive orders (and thereby ensure early execution), in order to improve their priority in the queue. In equilibrium, this leads to a greater welfare improvement than a scheme that does not fully account for traders' private values.

The gross surplus varies by less than 1% across the four markets we consider. Thus, although adverse selection is important in terms of the welfare of individual trader types, it has no appreciable impact on overall welfare. As we show in the next section, the limit order market is robust in the following sense: more information is revealed (via market observables) to uninformed traders when they are faced with a more severe adverse selection problem.

Finally, for any cost of information acquisition greater than 0.004, net surplus is highest when no agents are informed—while it is optimal for individual traders to acquire information, the overall effect is that too much information is acquired, relative to its social value.[28]

6 Informational Efficiency

Even if allocations are similar across different information regimes, the informativeness of market observables may differ. Since 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.

[28] It is important to emphasize that, in our model, the value of the asset (or v) is assumed exogenous. Clearly, in an environment in which agency issues between the shareholders and managers are important, information acquisition will have a greater social role.

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. In fact, as we show, uninformed traders in our model do learn from each of these. In a static rational expectations model, by contrast, the transaction price is perhaps the only device that can communicate information to the uninformed. Therefore, rather than focusing on how much (or how quickly) information is impounded into prices, we discuss more broadly the role of market observables (including transaction prices) in conveying information.

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 6, we report the mean absolute difference between an uninformed agent's expectation about the common value and the true consensus value.[29] 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. The table also reports the standard deviation of the price minus the true common value: this is increasing as the number of informed agents decreases. Thus, the transaction price is also less informative when fewer agents are informed.

	All Informed	$\vert\beta\vert \in \{4, 0.1\}$ Informed	Speculators Informed	None Informed
Standard deviation of $(p - v)$	0.80	0.91	1.11	1.61
Mean absolute error in uninformed agents' belief about v	0.31	0.39	0.49	1.05

Table 6: **Volatility in prices and errors in beliefs**

Notice from Table 6 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.

[29]In the model in which all agents are informed, we use the beliefs of agents who deviate at the information acquisition stage.

As shown in the previous section, gross investor surplus is approximately the same across information structures.

Recall that our traders observe both a snapshot of market history and the current book on each arrival to the market. The history consists of the price and sign of the previous transaction.[30] 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.

To obtain the observations in these regressions, we do the following. 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 base case (with only speculators informed), these represent 84.4% of all books encountered in the simulation.

The results are reported in Table 7. In the regressions, all prices and values are relative to the current common value, v_t. The dependent variable is thus $E(v_t \mid \hat{L}_t, v_{t-\Delta_t}, \hat{p}_t, b_t) - v_t$, and the lagged common value is similarly $v_{t-\Delta_t} - v_t$. The dependent variables include information from the book (bid, ask, depths at the bid and ask prices, and cumulative buy and sell depths), as well as the last transaction price $\hat{p}_t$ and the sign of the last order, b_t. The signed order variable has a value of 1 if the last transaction in the market was a buy order and -1 if the last transaction was a sell order. The first column in the table indicates the raw correlations between the explanatory and dependent variables.

We report two sets of regressions. From the first regression, the constant is basically 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 one-quarter of a tick.[31] The other variables are statistically significant, but have close to zero correlation with the dependent variable.

The last column reports the same regression with only prices (quotes and transaction price). The R^2 remains high, so that in both cases a linear regression captures the essence of the updating process. Again, a one tick increase in any of the bid, ask, or last transaction prices results in an increase of close to one-quarter of a tick in the agent's belief about v. In

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

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

Independent Variable	Correlation with dependent variable	Regression 1	Regression 2
Constant	0	-0.00 (0.61)	0.00 (1.87)
Lagged common value, $v_{t-\Delta_t}$	0.39	0.06 (114.86)	0.08 (140.47)
Signed order, b_t	0.01	-0.06 (80.49)	
Last transaction price, $\hat{p}_t$	0.68	0.22 (231.73)	0.21 (202.57)
Bid price, B_t	0.72	0.25 (204.40)	0.22 (166.73)
Ask price, A_t	0.72	0.25 (203.65)	0.21 (164.78)
Bid depth	0.00	0.13 (120.86)	
Ask depth	-0.00	-0.13 (120.29)	
Total buy depth	0.01	-0.04 (71.38)	
Total sell depth	-0.00	0.04 (72.12)	
No. of observations		146,069	146,069
R^2		0.81	0.74

Note: t-statistics in parentheses

Table 7: **Regression of change in belief about v, when only speculators are informed**

this model, therefore, prices (both transaction prices and quotes) are extremely informative about the common value.

Similar regressions were run for the other information structures (all agents informed, agents with $|\beta| \in \{0.1, 4\}$ informed, and no agents informed). The results were qualitatively similar, though the coefficients vary somewhat with information structure.

Finally, agents' outcomes (in terms of payoffs) are also less dispersed when a larger number of agents are informed. Table 1 includes the sample standard deviations in payoffs across different information structures. The standard deviation of equilibrium payoffs for all β values is monotone in the extent of overall information acquisition, being highest in the regime with no informed agents, and lowest when all agents are informed. Though our agents are all risk neutral, it is nevertheless interesting to observe that even uninformed agents obtain less dispersed outcomes as the number of informed agents increases.

7 Comparative Statics on Degree of Adverse Selection

Section 5 shows that, in our base case, different degrees of adverse selection (caused by a different number of informed traders) have little impact on aggregate investor surplus. In this section, we consider two other ways to increase the degree of adverse selection: first, by changing the timeliness of information about the common value observed by our uninformed traders, and second by changing the size of the gains to trade. We focus on the case in which only speculators are informed, since the information acquisition cost range that supports this outcome in equilibrium is relatively large.

First, we consider varying Δ_t, the lag with which our uninformed traders observe the common value. In addition to providing a comparative static of interest, this exercise serves as a robustness check on the parameter Δ_t, the value of which is difficult to empirically obtain. We have already examined the cases in which $\Delta_t = 0$ (in which all agents are informed) and $\Delta_t = 16$ (our base case). We now consider the effect of increasing Δ_t to $32, 64$, and 128. When $\Delta_t = 128$, since v changes on average every 8 units of time, the common value will change an average of 16 times from the value observed by an uninformed trader to its true value when he arrives at the market. The potential for information asymmetry between uninformed and informed traders is highest in this model. Of course, the realized information asymmetry is endogenous, since uninformed traders learn from market observables.

Figure 2 displays the aggregate welfare across different values of Δ_t (measured against the left axis), as well as the mean absolute error in the beliefs an uninformed trader has about v (measured against the right axis). As the figure indicates, increasing the potential effect of asymmetric information in this manner has a minimal impact on aggregate surplus: surplus declines monotonically from 3.734 when $\Delta_t = 0$ to 3.705 when $\Delta_t = 128$. The mean absolute error in an uninformed agent's belief about v increases with Δ_t, from 0.30 when $\Delta_t = 0$ to 0.62 when $\Delta_t = 128$. Notice that this error levels off between $\Delta_t = 64$ and

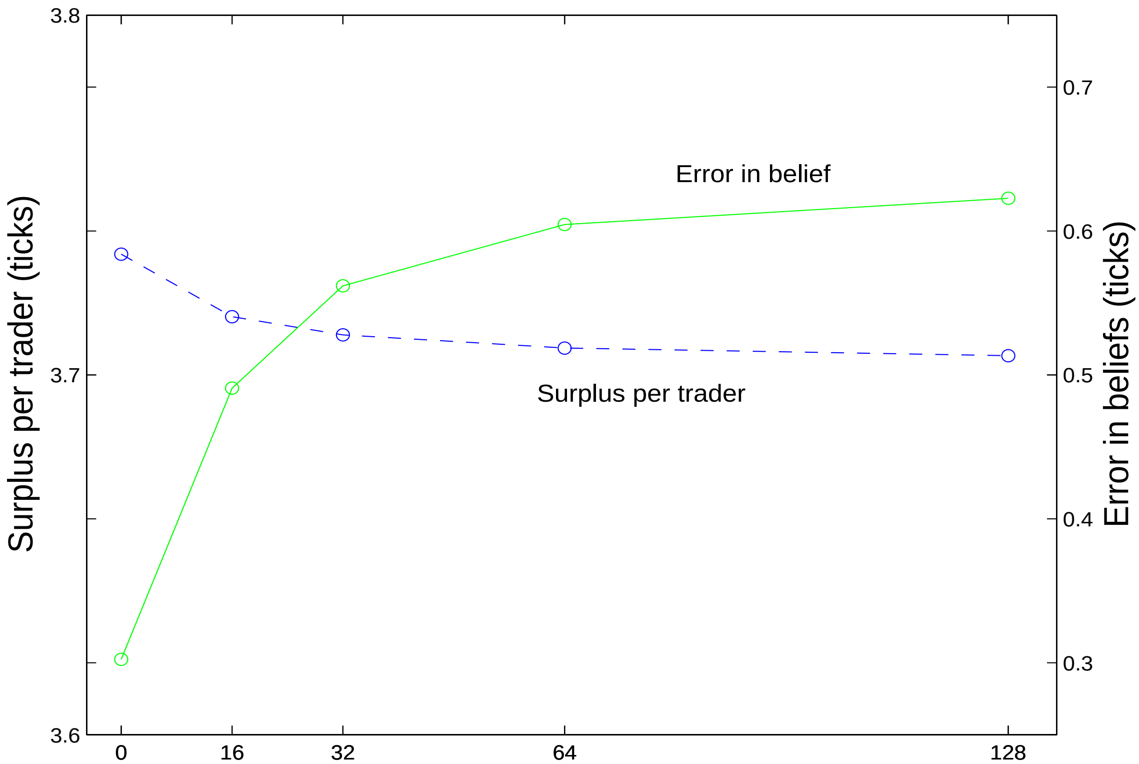

Figure 2: **Aggregate surplus and mean absolute error in belief, as Δ_t changes**

$\Delta_t = 128$.

Computationally, we cannot run the model with significantly higher values of Δ_t (since the state space grows too large). However, by the time $\Delta_t = 128$, the last observed value of v plays a minimal role in the updating process for an uninformed agent. In regressions similar to Regression 1 of Table 7, the coefficient on $v_{t-\Delta_t}$, the lagged common value, falls monotonically as Δ_t increases, from 0.059 in the base case ($\Delta_t = 16$) to 0.016 when $\Delta_t = 128$. These agents essentially look at market observables, such as the book and transaction history, to form their estimates of the common value. Since the informed traders compete with each other, market observables remain informative as Δ_t increases.

This comparative static suggests that increasing Δ_t beyond 128 will not hamper the market in any significant manner. Market observables will continue to be informative, limiting the degree of realized adverse selection, and trade will be consummated almost as efficiently as in the base case.

The robustness of the market in terms of aggregate surplus, however, depends on the size of the gains to trade among agents relative to the volatility in common value. If the private benefit to trade is high, an unknown common value is less likely to deter trade. Conversely, if the volatility in common value is low, trades can occur with high frequency even when the private benefit of trade is low. Thus, given the same proportion of informed agents in the population, adverse selection is more important when the idiosyncratic benefit of trade is lower. One important caveat is that another friction, the discreteness of prices at which orders can be submitted, is also more restrictive when the gains to trade are smaller. Thus this comparative static both changes the degree of adverse selection and the importance of other frictions in the model.

We simulate a model in which changes in v occur at the same frequency and magnitude. However, the distribution of private values is as follows: $\beta \in \{-1.2, -0.6, -0.1, 0.1, 0.6, 1.2\}$, with probability 0.3 of $\beta \in \{-0.1, 0.1\}$ (as before), probability 0.4 of $\beta \in \{-0.6, 0.6\}$ (instead of $\{-4, 4\}$) and probability 0.3 of $\beta \in \{-1.2, 1.2\}$ (instead of $\{-8, 8\}$). Note that the benchmark frictionless surplus with these parameters is 0.63, significantly lower than in our base case.

Consider two cases with this new private value distribution. First, suppose all agents are informed: the aggregate surplus in this case is 0.55, which represents 87.24% of the frictionless surplus. Next, suppose only the speculators (agents with $|\beta| = 0.1$) are informed. This leads to an aggregate surplus of 0.50, representing 79.37% of the frictionless surplus. In this latter case, the speculators earn a higher payoff, but the other types of agents suffer, and all agents take longer to execute.

Therefore, the effect of asymmetric information on aggregate welfare depends on the size of the gains to trade, relative to the volatility in the common value of the asset. In our base case, with gains to trade corresponding to known empirical results, adverse selection has a minimal effect on welfare. However, when the gains to trade are significantly smaller (and the relative friction due to discrete prices is thus larger), adverse selection does reduce aggregate welfare.

8 Conclusion

We model endogenous information acquisition in a dynamic limit order market, and demonstrate that the value of information to an agent depends both on his intrinsic motive for trade (i.e., his private value in our model) and other agents' information acquisition strategies. Agents with low intrinsic benefits from trade have the highest value for information. These agents are also the natural liquidity suppliers in the limit order market. Competition among these agents implies that the limit order book communicates information about the common value to uninformed investors. In particular, uninformed investors rely on the quotes to update their beliefs about the common value.

As a result, the transparent limit order market is robust in terms of consummating trade and generating consumer surplus. When gains to trade among agents are sufficiently large (and in accordance with empirical measurements), aggregate surplus generated in the market does not depend on the degree of adverse selection. If gains to trade are small, adverse selection can hamper trade. However, when gains to trade are small, the market itself is of limited interest.

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.

The open limit order book appears to be gaining in popularity around the world, with many financial markets organized on this basis. Encouragingly, we find that a limit order market is efficient at consummating trade. Competition between informed agents, transparency of the book, and the flexibility afforded to uninformed agents in terms of the ability to wait for, and post, better prices are all valuable features in this regard.

Appendix A: Details of the Algorithm

We fix information acquisition strategies $\sigma_I(\theta)$, and solve for the equilibrium of the corresponding trading game.[32] 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 (6).

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.

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 $U_0(\tilde{a} \mid s)$. 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 non-negative 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 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.

   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.

[32]The overview of the algorithm appears in section 2.3.

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.[33] 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 31 ticks on the price grid, with 15 ticks below v_t, 15 ticks above v_t, and one tick at v_t.

2. At time $t = \min\{t_v, t_n, t_r\}$, an exogenous event occurs. Suppose $t_v < t_n$ and $t_v < t_r$. 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 $t_v < t_r$. 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.

   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 label the trader as a deviator with probability 0.01.[34] If he is a deviator, we set $I = 1$ when $\sigma_I(\theta) = 0$ and $I = 0$ when $\sigma_I(\theta) = 1$. If the new trader is not a deviator, we set $I = \sigma_I(\theta)$. 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

[33] Importantly, 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.

[34] We also require that there be no other deviators currently in the market, to preserve the notion of unilateral deviation.

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 $t_r < t_v$ and $t_r < t_n$. 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.[35] 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 $m(0)$. If market conditions $m(0)$ are encountered in the simulation, set

$$\alpha(m(0)) = \frac{r}{r+1}\alpha(m(0)) + \frac{1}{r+1}(v_t - v_{t-\Delta_t}), \tag{8}$$

where $r(m(0))$ is a positive integer that is incremented by one each time the market conditions $m(0)$ are encountered in the simulation. An uninformed trader's estimate of common value at any point of time is then $\hat{v}(m(0)) = v_{t-\Delta_t} + \alpha(m(0))$.

Using this estimate of $\hat{v}$, the action set for each trader is defined as in equation (4) 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

$$U_{t'}(\tilde{a}^* \mid s) = \frac{n}{n+1} U_t(\tilde{a}^* \mid s) + \frac{1}{n+1} e^{-\rho(t'-t)} J(s', y_{t'}). \tag{9}$$

[35] 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.

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 (9)). 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

$$U_{t'}\left(\tilde{a}^* \mid s\right) \quad = \quad \frac{n}{n+1}\, U_t\left(\tilde{a}^* \mid s\right) + \frac{1}{n+1}\, e^{-\rho(t'-t)}\, \tilde{x}(\beta + v_{t'} - \tilde{p}^*), \tag{10}$$

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 (10), 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.[36]

[36] When a player trembles at $t' > t$, the payoff of the optimal action at t' is used to update $U_t^k(\cdot)$ to

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 $k_2 \geq k_1$ 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 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).[37] 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

$U_{t'}^{k+1}(\cdot)$. Thus, traders do not anticipate behaving suboptimally in the future.

[37] The theoretical properties of this test were derived by den Haan and Marcet (1994).

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.

References

[1] 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 p 76-115.

[3] Back, Kerry, H. Henry Cao and Gregory A. Willard (2000), "Imperfect Competition among informed traders," *Journal of Finance* Vol 55, No 5, p 2117 –2155.

[4] Barlevy, Gadi and Pietro Veronesi (2000), "Information Acquisition in Financial Markets," *Review of Economic Studies* 67(1):79–90.

[5] 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.

[6] Berk, Jonathan B. (1997), "The Acquisition of Information in a Dynamic Market," *Economic Theory* 9: 441–451.

[7] Bernardo, Antonio and Kenneth Judd (1997), "Efficiency of Asset Markets with Asymmetric Information," CMS working paper #16-97.

[8] 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.

[9] Blouin, Max (2003), "Equilibrium in a Decentralized Market with Adverse Selection," *Economic Theory* 22: 245–262.

[10] Blume, Lawrence and David Easley (1990), "Implementation of Walrasian Expectations Equilibria," *Journal of Economic Theory* 51: 207–227.

[11] den Haan, Wouter J., and Albert Marcet (1994), Accuracy in Simulations, *Review of Economic Studies*, 61, 3–17.

[12] Foster, S. Douglas and S. Viswanathan (1996), "Stratetic Trading when Agents Forecast the Forecasts of Others," *Journal of Finance* 51(4): 1437–1478.

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

[14] Glosten, Lawrence (1994), "Is the Electronic Open Limit Order Book Inevitable?" *Journal of Finance* 49(4): 1127–1161.

[15] Goettler, Ronald, Christine A. Parlour and Uday Rajan (2004), "Equilibrium in a Dynamic Limit Order Market," forthcoming, *Journal of Finance*.

[16] Grossman, Sanford and Joseph Stiglitz (1980), "On the impossibility of Informationally efficient markets," *American Economic Review* Vol 70 p 393–408.

[17] 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.

[18] Hayek, F. A. (1945), "The Use of Knowledge in Society," *American Economic Review* 35(4): 519–530.

[19] Hirshleifer, Jack (1971), "The private and social value of information and the reward to inventive activity," *American Economic Review* 61(4): 561–574.

[20] Holden, Craig and A. Subrahmanyam (1992), "Long–lived Private information and imperfect competition," *Journal of Finance*, Vol 47, No 1 p 247–270.

[21] Hollifield, B., R. Miller, P. Sandås, and J. Slive (2004), "Estimating the gains from trade in limit order markets," Working Paper, University of Pennsylvania.

[22] Jackson, Matthew O. (1991), "Equilibrium, Price Formation, and the Value of Private Information," *Review of Financial Studies* 4(1): 1–16.

[23] Kyle, Albert (1985), Continuous Auctions and Insider Trading, *Econometrica* 53 p 1315-1336.

[24] Maskin, Eric and Jean Tirole (2001), "Markov Perfect Equilibrium I: Observable Actions," *Journal of Economic Theory* 100: 191–219.

[25] Mendelson, Haim and Tunay Tunca (2004), "Strategic Trading, Liquidity, and Information Acquisition," *Review of Financial Studies* Vol 17 No 2. p 295-337.

[26] Pakes, Ariel and Paul McGuire (2001), 'Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the 'Curse' of dimensionality,' *Econometrica* Vol 69 No 5. 1261-1281.

[27] Parlour, Christine A. (1998), "Price Dynamics in Limit Order Markets," *Review of Financial Studies* 11: 789–816.

[28] Perry, Motty and Philip Reny (2004), "Towards a Strategic Foundation for Rational Expectations Equilibrium," Working paper, University of Chicago.

[29] Rieder, U. (1979), Equilibrium Plans for Non-Zero-Sum Markov Games, in *Game Theory and Related Topics*, eds. O. Moeschlin and D. Pallaschke, North Holland Publishers, 91–101.

[30] Rosu, I. (2004), "A Dynamic Model of the Limit Order Book," Working paper, University of Chicago.

[31] Schlee, Edward, E., 'The Value of Information in Efficient Risk Sharing Arrangements, *American Economic Review* 91 (3) p501-524.

[32] 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.

[33] Taub, Bart, Dan Bernhardt and P. Seiler (2004), "Cladistic Asset Pricing," University of Illinois working paper.

[34] Verrecchia, Robert (1982), "Information Acquisition in a Noisy Rational Expectations Economy," *Econometrica* 50: 1415–1430.