

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

May 15, 2004

Preliminary: Comments welcome

*We have benefitted from seminar participants at CMU. The current version of this paper is maintained at <http://ozymandias.gsia.cmu.edu>

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

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

[§]Tel: (412) 268-5744, E-mail: urajan@andrew.cmu.edu

Abstract

We model an infinite horizon trading game of a limit order market with informed traders. Agents with a private and common value motive for trade randomly arrive in a market and may either post prices (submit limit orders) or accept posted prices (submit market orders). If their orders have not executed, traders may reenter the market and thus solve a dynamic problem. We consider agents' incentive to acquire information. We characterize how information acquisition changes agents' strategies and demonstrate the effect of this on the efficiency of market prices. We demonstrate that for some costs of acquiring information, there are multiple equilibria in the information acquisition game. Finally, we demonstrate that information acquisition can make all agents worse off.

1 Introduction

How does information get into stock prices? The answer to this question is important for both practical and theoretical reasons. Theoretically, the notion of efficient markets requires that information is acquired voluntarily by agents and through their trading behavior is incorporated into price. Practically, if we know who acquires information and how this affects their trading strategies, better inferences can be drawn from transactions data.

The relationship between price and information pertinent to the common value was first explored by Hirschliefer (1971), who found that in equilibrium, if prices are fixed, then each agent has an incentive to acquire information. However, this can lead to market breakdown and has no social value (i.e., all agents are worse off).¹ Grossman and Stiglitz (1980) observed that if costly information is immediately impounded in price, then agents do not acquire it. Clearly, both arguments depend on how agents profit from their information. If they benefit through trading in the asset, then all results are specific to the price formation mechanism. Thus, to answer these questions robustly, a model should either explicitly model current real-world markets or include stylized representations of the most important trading frictions.

To date, important insights into how information is impounded in price have been generated in both competitive and strategic rational expectations models. By contrast, our framework is a fully strategic trading game and thus falls outside this paradigm. Understanding trading frictions in a competitive rational expectations equilibrium is complicated by the fact that generically, Blume and Easley (1990) show that there is no game that has the competitive rational expectations equilibrium as an outcome. Or, there is no reason to suppose that any functioning security market will display the properties of a rational expectations equilibrium. Interestingly, Reny and Perry (2004) provide an example of a double auction that under stringent regularity conditions converges to the rational expectations equilibrium. Thus, while the rational expectations equilibrium is not generically the limit of any strategic market game, the closest example to a game which could converge to this outcome has the flavor of a limit order market with discrete prices.

We model an infinite horizon asset market as an open electronic limit order market. Most equities are traded in a variant of a limit order market. Briefly, this is a continuous double auction, where investors may either post prices (submit limit orders) or accept a posted price (submit a market order). In our model, risk neutral agents who value the asset as the sum of a private and common value arrive randomly at the market. Agents have different information about the cash flows accruing to the owners of the asset i.e.,

¹Hakansson, Kukul and Ohlson (1982) demonstrate in an exchange economy that if the market is not allocationally efficient and posteriors are not the same then information can have social value.

the common value. Each agent has one share to trade, and if his order does not execute, revisits the market with fixed probability and can revise his order. Thus, agents' strategies are fully dynamic. Ex ante, before the start of the trading game, agents choose to acquire costly information. An informed agent views the current value of the cash flows, while an uninformed agent views the cash flows with a lag. Each time an agent revisits the market, his information set is updated.

We show how traders' strategies change depending on their beliefs about other's information sets. One measure we use to determine how much trading outcomes have changed is the allocative efficiency of the market. We analyze the informational efficiency of market prices: how quickly new information becomes impounded in price and how quickly agents update their beliefs about the common value. We demonstrate that there can be multiple equilibria in the information acquisition game. Finally, we find that all agents in the market can be made worse off when some optimally acquire information.

The key difference between our approach and that of the older literature is that in our model, the gains from trade are fixed. If traders in one particular stock are modelled as risk averse, then if some agents become informed, the risk sharing opportunities are reduced. This, is the basis of the Hirshleifer effect. By contrast, we consider a world in which there are gains to trade and information affects how agents split such gains. Thus, we focus on the role of adverse selection, and how it affects agents' market behavior. We interpret rational expectations results with risk averse traders as primarily about industry or factor returns. Trade in these models is driven by risk sharing. However, from a portfolio perspective, the only risk sharing that is relevant should be systematic risk. One thinks of insiders or privately informed having information that pertains to the idiosyncratic returns. If investors hold well-diversified portfolios then aggregate risk sharing should not be an important trade motivation for agents averse to idiosyncratic risk.²

Agents in financial markets can acquire different types of payoff relevant information. First, there can be information about the underlying cash flows of the asset. For example, agents might have insider information or information about earnings or sales. Second, agents may have information about trading opportunities. That is, agents may be differentially informed about quotes, depth and other variables that allow them to tip the terms of trade to their advantage. In this paper, we consider information about the common value of the asset. Thus, we primarily interpret information as general company information such as earnings reports or corporate balance sheets.

Our model can also be interpreted as a model of private information. Since trading

²Bernardo and Judd (1997), find that information acquisition leads to a reduction in welfare as uncertainty is resolved before trade (the Hirshleifer effect), and rent-seeking trades by informed reduce optimal risk sharing.

on inside information is illegal in the United States, it is not immediate that a privately informed investor will trade in the underlying stock. In the US, the SEC monitors stocks and aggressively prosecutes those who trade based on insider information. If an insider chooses a more circuitous route (either trading on options or through other instruments), the information content of those prices will become available to all market participants.

While all this of information is readily available, acquiring it is costly. There can be an explicit cost such as subscribing to a news service, or an opportunity cost of time. Information about trading opportunities also enters our analysis in that it determines the inference agents can make about the common value. In equilibrium, traders in our model must choose whether to pay the cost of information. All informed investors have the same information; i.e., their signals are perfectly correlated.

The canonical strategic rational expectations model is Kyle (1985). One 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. Thus, this framework is inappropriate for examining how potentially information about (say) earnings gets incorporated into price. Presumably, in this framework, the market makers' price instantaneously adjusts to reflect such public information. Further, as all trades are consummated at the market maker's quoted prices, there is no distinction between quotes and transaction prices. The only other observable characteristic of the market is the net order flow.

An agent can benefit from superior information if he can trade on it: anything that restricts his ability to do so will reduce his benefit from acquiring it. Holden and Subrahmanyam (1992) show in a Kyle setting with multiple informed traders that prices reveal information almost immediately: informed traders compete away trading profits. Such a result relies on the fact that noise traders do not adjust their trades. Spiegel and Subrahmanyam (1992) demonstrate that changing the Kyle assumptions, by allowing uninformed traders with risk sharing motives to choose trades, 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 that they trade in the presence of adverse selection. Endogenous information acquisition has been examined in a Kyle framework with endogenous liquidity traders by Mendelson and Tunca (2004). In a model with strategic risk averse noise traders, they consider the effect of three types of information: intractable that cannot be acquired, tractable that can be, and public. As liquidity traders are risk averse, they benefit from the existence of an insider (this reduces uncertainty). As the insider takes into account the effect of his actions on liquidity traders, they find that the insider may choose not to acquire information (even at a zero cost).

Finally, an insider may acquire more information than the welfare maximizing amount. In a continuous time setting, Back, Cao and Willard (2000) demonstrate competing traders reveal less information than a monopolist insider if their signals are not perfectly correlated and they find that the market becomes illiquid at the end of trading. 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.

This paper is conceptually related to the seminal work of Admati and Pfleiderer (1987a) who provide conditions under which information is concentrated among traders (i.e., agents' incentives to acquire information) in a competitive multi-asset noisy rational expectations framework. They find that because information becomes aggregated in price, and agents condition on price, signals about the payoff of an asset, may become more valuable.³ By contrast, we find that as more agents acquire information, the price they are willing to pay for it is lower (i.e., information is less valuable). Further, as we have a model of strategic trade, agents may trade at prices before information becomes revealed. Indeed, the true value of the asset changes frequently in our model, and trade occurs both before and after such changes. This is one significant difference with between a strategic model and the rational expectations literature: the welfare properties in a strategic model are necessarily different.

Our paper links the literature on information acquisition to that on dynamic limit order markets. Work on dynamic limit order markets includes Rosu (2004), who presents a continuous time model of a limit order market. 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.

There is an extensive literature that considers the effects of differential information about the limit order book or trading opportunities. Foucault et al (2003) consider the effect of trader identity on a limit order book through a natural experiment. Pagano and Roell (1993) consider the redistributive effects of transparency in a Kyle (1985) type setting. Madhavan (1995) finds that the price volatility is higher in a market with less transparency because less information gets impounded into price. Biais (1993) compares centralized and

³Barlevy and Veronesi (2000) in a single asset Grossman Stiglitz framework with supply shocks also find that learning can be a strategic complement.

fragmented markets in which traders only have a private motive for trade. From a natural experiment on the NYSE, Boehmer, Yu, and Saar (2004) find that agents condition on the limit order book and that access to the limit order book can affect market quality. We find that uninformed traders benefit from information in the limit book as it allows them to infer private signals.

There is a literature that considers purveyors of information. Admati and Plederer (1986, 1987b, 1990) in a rational expectations framework, demonstrate that if price is revealing, a monopolist seller of information may prefer to provide noisy signals, or to set up a mutual fund to trade directly on his information. Simonov (1999) shows in a rational expectations context that if signals are complements then the price of information in a duopoly is higher than in a monopolist setting. Leshchinskii (2000) demonstrates that if strategic traders differ in risk aversion, then the more risk averse may reveal his signal to allow the less risk averse as an incentive to provide more information. In contrast to this literature, we take the price of information as given and characterize the equilibrium allocations that can obtain.

In Section 2 we outline the general model. Details of the algorithm, convergence criteria and parameterizations follow. We proceed by analyzing agents' order submission strategies under different information acquisition assumptions (Section 3). As market efficiency has been identified as an important determinant of the demand for information, we consider how quickly information becomes impounded in price across different assumptions about information acquisitions (Section 4). Finally, in Section 5 we characterize the demand for information for different possible information acquisition strategies, which allows us to find subgame perfect equilibria of the endogenous information acquisition game. We then compare the welfare properties across the different equilibria and draw conclusions in Section 6.

2 Model

We model an infinite horizon limit order market for a single stock. In philosophy, the model is similar to that in Goettler, Parlour and Rajan (2004). There is a common value to the asset, v , and it also has a private benefit, β , to each trader. On entry into the market, a trader observes the limit order book, 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. One important difference that significantly enriches the model is that traders who have submitted limit orders are allowed to re-enter the market, and change or cancel their order. That is, cancellations in this paper are fully endogenous.

Time in the model is continuous, though events happen only after discrete time intervals. There is a large (possibly infinite) set of discrete prices, denoted $\{p^0, p^1, \dots, p^N\}$, 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 $p^i \in \{p^0, \dots, p^N\}$ is a backlog of outstanding limit orders, $\ell^i(\tau)$. We sometimes refer to this as the depth at price p^i . We adopt the convention that buy orders are denoted as a positive quantity, and sell orders as a negative quantity. The limit order book, $L(\tau)$, is the vector of outstanding orders, so that $L(\tau) = \{\ell(\tau)^i\}_{i=0}^N$. The bid price is defined by the highest price at which there is a limit buy order on the book, and the ask price by the lowest price at which there is a limit sell order.

New traders arrive at the market according to a Poisson process with parameter λ_N . Hence, the actual time between the arrivals of traders is random. Each trader has a type denoted by $\theta = \{\rho, \beta, I\}$. The first element of the trader’s type, ρ , is a continuous discount rate. The payoff he earns as a result of trading is discounted back to his first arrival time in the market at this rate. The discount rate captures the notion that traders would rather execute sooner than later. In the model, it prevents a trader from infinitely postponing trade.⁴ In this paper, we hold ρ fixed across all agents.

In addition, each trader has a private value for the asset, which we denote 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 τ has a common value, denoted $v(\tau)$. The common value is interpreted as the expectation of the present value of future cash flows on the stock. Innovations in the common value arrive according to a Poisson distribution with mean μ . If an innovation occurs, then with probability $\frac{1}{2}$ the common value increases by one tick and with the same probability decreases by one tick. Changes in the common value reflect new information about the firm or the economy.

Finally, I refers to the information the trader has about the common value, v . In this paper, we consider two kinds of agents. Informed agents know the current value of v at each instant. Uninformed agents view v with a lag Δ_τ . For computational convenience, this lag is measured in terms of the number of trader arrivals to the market. That is, an uninformed agent in the market at time t knows the true value of v at the time of the previous Δ_τ^{th} trader arrival. An informed agent in the market at time t knows the current value $v(t)$.

⁴Discount rates are also present in the models of Foucault, Kadan and Kandel (2004) and Rosu (2004). These models fix the gains from trade and follow Demsetz (1968) in that differential waiting costs imply particular patterns of trade.

Thus $I = \{0, \Delta\tau\}$.

We are interested in endogenous information acquisition by the agents. That is, suppose there is a cost c to acquiring information about v . Before they first enter the market, agents may choose whether to subscribe to a service that reports the current value of v . Since the state of the market changes more rapidly (on average) than the common value of the asset, our traders do not have the option of choosing whether to buy information based on the observed situation in the market. Since the information is potentially available to all investors, there is a sense in which this is public information.

If they choose to not buy information, they observe v with a lag of $\Delta\tau$ trader arrivals. If they do subscribe to the service, they observe the current value of v when they are in the market. Since traders make this choice before entering the market, in practice, each (ρ, β) pair will make the same choice, given c and given the information acquisition strategies of other agents—either all agents with this β and ρ will acquire information, or none of them will.

In practice, we solve for the equilibrium of the model for different cases, assuming an information acquisition strategy for each (ρ, β) pair. We then compute the change in expected payoff for each pair from deviating at the information acquisition stage. This provides bounds for the information acquisition cost, c , for which the assumed information acquisition strategies constitute an equilibrium.

Each trader is allowed to trade exactly one share of the asset. However, they may choose to buy or sell a share. Further, traders who have entered the market previously, but have not executed yet, re-enter the market at some random time. As a result, 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 for ever. Thus, at any point of time, there will be a random number of agents who have not yet traded. Each unexecuted trader re-enters the market according to a Poisson distribution with parameter λ_R . This captures the idea that agents monitor the market, but not continuously. The re-entry times are independent across agents. Let G denote the distribution over re-entry time, with g the associated density. At any particular instant, there is at most one agent (either a new or returning trader) who takes a decision.

Upon re-entry, a trader may leave an existing order on the book, or cancel it and submit a new order. The benefit of retaining the existing order is that he maintains his time priority (his place in the queue). The cost is that the asset value may have moved in a manner that affects the expected payoff from the order. For example, if he submitted a buy order and the asset value has fallen since then, his order may be picked off. Conversely, if the asset value has risen since then, his offer may be at too low a price, and there may be little chance

of it executing. Of course, a trader may also find that the priority of a previous order has improved by the time he re-enters the market. Suppose a trader submitted an order at price p . Given book $L(\tau)$, let $q(\tau)$ denote the number of orders with higher priority at the same price, among all orders at price p at time τ .

When he is in the market at time τ , a trader may submit an order $x^i(\tau)$, which denotes an order at price p^i , and

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

If there is an existing order at price p^i on the other side of the market, the submitted order executes immediately and is called a market order. 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 forward in the queue. Buy orders are accorded priority at higher prices, and sell orders at lower ones. Therefore, an order executes if no other orders have priority, and a trader arrives who is willing to be a counter-party.

The perceived trading opportunities for each trader depend on his information set. Let $s(\theta)$ be the state observed by the trader. For a fully-informed trader in the market at time τ , $s(\theta) = \{L(\tau), v(\tau)\}$. That is, he knows the entire book $L(\tau)$ and the current consensus value $v(\tau)$. In addition, the trader also knows the price p and priority at that price q of his previous order.

Consider instead a trader who is not fully informed, but only views the consensus value of the asset with a lag of Δ_τ trader arrivals. Let $\tilde{\tau}$ denote the (random) time before τ at which the previous Δ_τ^{th} trader arrival was recorded. In addition, this agent observes the transactions that have taken place since $\tilde{\tau}$. To limit the size of the state space, we allow the trader to condition on the difference between the total number of market buy orders and market sell orders that have taken place since then, n_{Δ_τ} . This difference may be positive or negative. Hence, the state space for an uninformed trader is $s(\theta) = \{L(\tau), v(\tilde{\tau}), n_{\Delta_\tau}\}$. For computational tractability, for all agents (informed and uninformed) we restrict order submission to a finite set of prices centered around the agent's expectation of $v(\tau)$. That is, the agent is allowed to submit a limit order at prices up to k ticks above or k ticks below his expectation of $v(\tau)$. He is allowed to submit a market buy (sell) order at the current ask (bid), regardless of his expectation of $v(\tau)$.

Consider the problem faced by a trader who is in the market at time t . Suppose this trader is re-entering the market (the problem faced by a new trader is similar to the problem faced by a re-entering trader who did not submit an order on his previous entry), and, on

his previous entry (at $t' < t$), 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 as is, and taking no further action.

To capture this possibility, let $a = (p, q, x)$ denote an action taken by a trader. Here, p is the price at which an order is submitted, and q is the priority of his order among all orders on the same side at the market at price p . As before, x is -1 or 1 depending on whether the order is a sell or a buy order. If no order is submitted, let $x = 0$. In the latter case, p and x correspond to the values chosen at the time the retained order was submitted. When $x = 0$, and no order is submitted, the values of p are irrelevant.

Let $\mathcal{A}(\theta, s(\theta), a)$ denote the action set of the trader. Here, $s(\theta)$ denotes the state faced by the trader when he is in the market at this time, and $a = (p, q, x)$ is the status of the previous action of this trader. If he is a new trader, we let $x = 0$. The feasible action set for the trader depends on his type. As mentioned, each trader can submit a market order (at the prevailing quotes), or a limit order within k ticks of his expectation of the common value. Since the latter depends on his type, so does the action set.

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 $\theta = (\rho, \beta, I)$, the instantaneous utility at time τ may be defined as

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

If a trader submits a limit order, his execution time, \tilde{T} , is random. The limit order executes only if another trader submits a market order that executes against it. Let $F_{\tilde{T}}(\cdot | s, a)$ denote the distribution of execution time of his order, given that he takes action a and the state he faces is s . Execution times for limit orders are endogenous: they depend on the actions of this trader as well as those of future traders.

The probability distribution over execution times is different for different prices and orders. In particular, a market order submitted at t executes at time t . Formally, this is represented as a probability distribution that has a mass of 1 at t . Intuitively, we expect limit buy orders at higher prices to execute sooner than orders at lower prices.

Consider a trader in the market at time t . Suppose he faces state s , and his previous action is given by a . When the trader submits an order, he has to consider the distribution over execution times for that order, as well as the distribution of his own re-entry time into the market. Upon re-entry, 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.

We denote the value to trader type θ of being in the market in state s , given that his previous order is a , as $J(s, a \mid \theta)$. On entry into the market, the trader has a finite action set, $\mathcal{A}(\theta, s, a)$. 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 re-enters the market, and second, the value associated with re-entering the market in some new state s' . For convenience, we normalize the time at which the trader has to take a decision to be time 0. Let $\phi(\tau, v_\tau; s, \tilde{a})$ be the probability that an action \tilde{a} taken in state s at time 0 executes at time $\tau > 0$ when the common value is v_τ . Then, the value $J(s, a \mid \theta)$ to trader type θ of entering the market in state s , given that his previous action was a , may be written as follows:

$$J(s, a \mid \theta) = \max_{\tilde{a} \in \mathcal{A}(\theta, s, a)} \left[\int_{w=0}^{\infty} \left\{ \int_{\tau=0}^w \int_{v_\tau=-\infty}^{\infty} e^{-\rho\tau} \tilde{x}(\beta + v_\tau - \tilde{p}) \phi(\tau, v_\tau; s, \tilde{a}) f_v(v_\tau \mid v, \tau) dv_\tau d\tau \right. \right. \\ \left. \left. + (1 - F_{\tilde{T}}(w \mid s, \tilde{a})) e^{-\rho w} \int_{(s', \tilde{a}') \in \mathcal{S} \times \mathcal{A}} J(s', \tilde{a}' \mid \theta) h(s', \tilde{a}' \mid s, \tilde{a}, w) d(s', \tilde{a}') \right\} g(w) dw \right]$$

The first term on the RHS indicates the payoff from execution before re-entry at the random time w . The distribution of execution time is denoted $F_{\tilde{T}}$. This distribution is defined as $F_{\tilde{T}}(w) = \int_{\tau=0}^w \int_{v_\tau=-\infty}^{\infty} \phi(\tau, v_\tau; s, \tilde{a}) f_v(v_\tau \mid v, \tau) dv_\tau d\tau$. This depends on the action taken (for example, for a market order, $F_{\tilde{T}}(0) = 1$, since the order executes immediately), and on changes in the common value. The latter affects the actions taken by subsequent agents.

Suppose the agent takes an action $\tilde{a} = (\tilde{p}, \tilde{q}, \tilde{x})$. Suppose further that his order executes at a time $\tau \in [0, w]$. Then, the payoff to the order depends on the common value at time τ , which we denote v_τ . As noted, the instantaneous payoff of this order at time τ is $\tilde{x}(\beta + v_\tau - p)$. This payoff must then be discounted back to time 0, at the rate ρ . The innermost integral of the first term on the RHS is over the different common values that can obtain at time τ . Picking off risk is manifested in $\phi(\cdot)$, which is higher when v has moved in an adverse direction (for example, v has decreased after a limit buy was submitted).

The second term captures the payoff to the trader if his order remains unexecuted at time w . The probability of this is $(1 - F_{\tilde{T}}(w \mid s, \tilde{a}))$. The agent re-enters the market at the random time w . If his order is still unexecuted, he can choose to instead submit an order at a different price $\tilde{p} \neq p$ or in a different direction, $\tilde{x} \neq x$. This implies a cancellation of the previous order. Alternatively, he can choose to leave his previous order on the books, by setting $\tilde{p} = p$ and $\tilde{x} = x$. In this case, we have $\tilde{q} = q$, so that the order retains its status. Of course, the state may have changed since he first submitted the order, to s' . This could happen for exogenous reasons (for example, 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 \tilde{p} (so that $\tilde{q}_w < \tilde{q}_0$), or it could reduce the overall priority, if other agents submitted limit orders at prices more aggressive than \tilde{p} . Hence, the action \tilde{a} taken at time 0 evolves to \tilde{a}' by the time the trader re-enters at time w . Recall that g is the density of re-entry time; the outermost integral is over this random re-entry 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 the first term on the RHS is zero. Since a trader is never forced to submit an order, and, in this model, there is no cost to re-entering 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.

2.1 Existence

The equilibrium concept we use is stationary Markov-perfect equilibrium. The existence of a Markov perfect equilibrium follows from standard results. 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.

Perfection requires that agents' beliefs about payoffs to actions off the equilibrium path be correct. Numerically, this requires the computation of beliefs about actions that are not chosen. To do this, we introduce the notion of trembles. With probability $(1 - \epsilon)$ close to 1, agents play best responses whenever they enter the market. However, on each entry, there is a small probability ϵ that an agent may tremble to a sub-optimal action. For informed traders, the payoff of a market order is known with certainty, so trembles only need to be to limit orders.

known to the trader, they are easily determined in the simulation, and can be used to update traders' beliefs even when market orders are not chosen. Hence, again, trembles by these agents only need to be to limit orders. The trembles enable the determination of payoffs for all limit orders a trader may submit. Of course, the probability of trembles has to be sufficiently small to not affect the strategies along the equilibrium path (in the simulation, traders will respond to the possibility that other agents may tremble, and their

equilibrium strategies may change).

2.2 Solving for Equilibrium

The simulation algorithm follows Pakes and McGuire (2001) in that it is asynchronous, and the values for each state are only obtained on the recurrent class of states. We depart from Goettler, Parlour and Rajan (2004) as we directly solve for the equilibrium values for every state encountered by every type of trader. Since traders may return to the market several times, they are solving true dynamic programs. Even if there were no changes in common value in this model, therefore, it is not possible to reduce the problem to one of determining execution probabilities for every order they can submit.

There are three sets of exogenous events that drive the simulation—arrival of new traders, re-entry of old traders who have not yet executed, and changes in the common value.

The algorithm is as follows:

1. At time 0, we start with an empty book. At this time, there are no returning traders. Hence, the only two exogenous events that can occur are a new trader arrival or a change in the common value. For each of these two events, we draw a random time from the respective random process. The time interval between events for a Poisson process has an exponential distribution, so we use the latter to draw these times. Let t_v denote the additional time until v changes, and t_n the additional time until a new trader arrives.
2. At $t = \min\{t_v, t_n\}$, an exogenous event occurs. Suppose $t_v < t_n$. 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. The time before a new trader arrives is now $t_n - t_v$. When the common value changes, we draw a new random time t_v for the next change in common value.

Suppose, instead, $t_n < t_v$. A new trader arrives to the market. His type is denoted as $\theta = \{\rho, \beta, I\}$. The discount factor ρ is the same for all traders, and β is drawn independently from the distribution F_β . The trader's information may depend on his β (details are specified in Section 4).

The trader observes the state s and his previous action a . Since he is a new trader, his previous action has $x = 0$, and the state when the first new trader enters is (L_{t_n}, v_{t_n}) , where L_{t_n} is the empty book, and v_{t_n} is the current value of v . The trader takes an action a_n . At this time, the common value will change after a further lapse of

time given by $t_v - t_n$. We also draw a new random time interval t_n before the arrival of a new trader, and a random time interval t_o before the current new trader returns to the market.

3. After the first trader has entered the market, either a returning trader or a new one may come to the market. When an event occurs, let t_v, t_n, t_o be the elapsed time before the common value changes, a new trader arrives, or an old trader returns. The next event occurs at $t = \min\{t_v, t_n, t_o\}$. If the event that occurs at t is a change in common value or a new trader arrival, the process in the previous step is repeated. Suppose, instead, it is an old trader that returns (so that $t_o < t_v$ and $t_o < t_n$). The old trader observes the current state s_{t_o} and the status of his previous action a , and 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, the interval for a change in common value is redefined to be $t_v - t_o$, and for a new trader arrival it is $t_n - t_o$. We then draw a new time interval t_o before this old trader returns to the market yet again.
4. Suppose a trader is in the market. To choose a payoff-maximizing action, he has beliefs about the discounted expected utility from every action he can take in that state. The trader has a type θ , and observes the current state $s(\theta)$ (recall that the state is defined to depend on trader type), and the status of his previous action, a . Let $U(\tilde{a} \mid \theta, s, a)$ be the (actual) discounted expected utility to a trader of type θ from taking action \tilde{a} in state s , when his previous action is a . Then, $J(s, a) = \max_{\tilde{a} \in \mathcal{A}(\theta, s, a)} U(\tilde{a} \mid \theta, s, a)$.

At the start of the simulation, traders have beliefs over these payoffs. Let $U^k(\tilde{a} \mid \theta, s, a)$ be the trader's belief about the expected payoff, after the action \tilde{a} has been taken k times in state s , given previous action a . The initial belief is denoted $U^0(\tilde{a} \mid \theta, s, a)$, since this action has never been taken.

5. Suppose a trader is a new trader at time t , and takes an action \tilde{a} that does not represent a market order. This action could involve not submitting an order at time t . At some future point of time, t' , the trader re-enters the market. He finds that his action has evolved to \tilde{a}' , and the new state is s' . The payoff from the action \tilde{a} on this visit (i.e., the continuation value of this action) is then defined to be $e^{-\rho(t'-t)} J(s', \tilde{a}')$. This payoff is “averaged in” to the belief $U^k(\tilde{a} \mid \theta, s, a)$ in the following manner. We define

$$U^{k+1}(\tilde{a} \mid \theta, s, a) = \frac{n+k}{n+k+1} U^k(\tilde{a} \mid \theta, s, a) + \frac{1}{n+k+1} e^{-\rho(t'-t)} J(s', \tilde{a}'), \quad (3)$$

where n is some integer chosen at the beginning of the simulation that determines how fast beliefs are updated.

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 the limit order is $x(\beta + v_{t'} - \tilde{p})$, where β denotes the private value of the trader. In this case, we update

$$U^{k+1}(\tilde{a} | \theta, s, a) = \frac{n+k}{n+k+1} U^k(\tilde{a} | \theta, s, a) + \frac{1}{n+k+1} e^{-\rho(t'-t)} x(\beta + v_{t'} - \tilde{p})$$

6. Whenever a trader has to take a decision, 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 state s and previous action a . In the simulation, we 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 sub-optimal for the trader. For this updating, we use equation (4), with $t' = t$ and $v_{t'} = v_t$.
7. Suppose the trader in the market at some time t is a returning trader, with an order on the book. To determine the payoff to a new order he submits, he must hypothetically cancel his existing order, thereby altering the book (and thus the state). He then compares the optimal action given the new state, and its associated payoff, to the action of retaining his order on the book. The overall optimal action is the one that yields maximal payoff across these two.
8. In the simulation, most traders take the optimal action given their 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 sub-optimally.

To ensure that beliefs are updated for all actions in every state, we introduce trembles. Specifically, with probability ϵ a trader trembles over all sub-optimal actions available to him. He chooses among sub-optimal actions with equal probability. The algorithm will then naturally update the beliefs about payoffs to this action.⁵

The initial beliefs are chosen to ensure that traders converge to the optimal action in a given state relatively quickly. Given that we allow traders to tremble, any initial belief

⁵When a player trembles, the payoff of the optimal action is used to update $U^k(\cdot)$ to $U^{k+1}(\cdot)$. Thus, traders do not anticipate behaving sub-optimally in the future.

can eventually lead to an equilibrium. The choice across initial beliefs is driven more by computational considerations (in particular, converging to equilibrium more quickly) than by a theoretical need.

2.3 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 periods, by computing $|U^{k_2}(\tilde{a} | \theta, s, a) - U^{k_1}(\tilde{a} | \theta, s, a)|$ for each \tilde{a}, θ, s, a . Here, k_1 is the number of times the action \tilde{a} has been chosen given θ, s, a at the start of the current 100 million periods, and $k_2 > k_1$ the number of times it has been chosen at the end of the current 100 million periods. 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 value functions fixed and simulate the model for a total 100 million more trader arrivals (new and returning). Let $U^*(\tilde{a} | \theta, s, a)$ be the fixed beliefs.

We compare the empirical payoffs from different actions to the fixed beliefs. This comparison is done at two levels: first, limit orders may transition to different states by the time a trader re-enters. The updating process here is given by equation (3); the convergence test involves comparing the actual utility $e^{-\rho(t'-t)}J(s', a' | \theta)$ to the belief $U^*(\tilde{a} | \theta, s, a)$. Second, eventually every trader in this model executes, and leaves the market. At the time he executes, he obtains a realized payoff $e^{-\rho(t'-t)}x(\beta + v_{t'} - \tilde{p})$ to his belief at the time of initial entry, $U^*(\tilde{a} | \theta, s, a)$.

We use three convergence criteria. The most stringent of these is a χ^2 test similar to that in Goettler, Parlour and Rajan (2004).⁶ Suppose $U^*(\cdot)$ indeed represents equilibrium values. Since the computed values $U^k(\cdot)$ are averages, the central limit theorem implies that the empirical distribution of values for each action in each state is approximately normal with mean U^* and a variance that is empirically determined from the simulation. The test statistic standardizes these normal variables and sums their squares. The statistic is χ^2 with degrees of freedom equal to the number of states used in the summation. We only use states visited at least 100 times to ensure that the central limit approximation is accurate. 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 $U^*(\cdot)$ and realized outcomes. 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

⁶The theoretical properties of this test were derived by den Haan and Marcet (1994).

0.01.

2.4 Parametrization

- We let $\Delta_\tau = 24$, where this represents the number of trader arrivals (returning and new) that occur between the time the uninformed trader observed v and the time that he is in the market.
- ρ , the continuous discount rate is the same for all agents and set to 0.05.
- F_v , the distribution for changes in common value, is a Poisson distribution. The expected time between changes in v is 12 units of real time.
- New traders arrive at the average rate of 1 trader per unit of real time.
- Agents re-enter the market at an average rate of 6 units of real-time per re-entry. Re-entries are independent across traders and entries.
- The support of the discrete β distribution is $\{-4, -2, -0.1, 0.1, 2, 4\}$. The probability of $-4, -2, 2, 4$ are all 20% and the probabilities of $-0.1, 0.1$ are 10% each. The traders with $\beta = -0.1, 0.1$ constitute traders who may be willing to buy or sell, depending on the state of the market when they come in. The traders with $\beta = 2, 4$ are likely to be buyers overall, and those with $\beta = -2, -4$ are likely to be sellers (this is borne out in our simulations).
- Limit orders may be submitted up to three ticks above or below an agent's expected value of v . For an informed trader, this is just his last observation on value. For traders who observed v with a lag, this consists of their best estimate given the common value they observed, the current book, and the difference between the cumulative market buys and sells observed in the interim.
- The probability that an agent trembles to a sub-optimal order is 0.01.

3 Trading strategies in different information acquisition regimes

A consistent feature of our results is that information is most valuable to the agents with $|\beta| = 0.1$. These agents have a low private motive for trade relative to the common value. Hence, when fully informed about the common value, they are only willing to trade if the transaction price is sufficiently above (if they are selling) or below (if they are buying) the common value. If other agents in the model are uninformed, these traders therefore have the most to gain from knowing the exact common value. Conversely, even if other agents in

the model are informed, they have the most to lose if they are uninformed. An agent with a β of 4, for example, who buys the asset 2.5 ticks above the common value, still earns a positive payoff overall. However, an agent with $\beta = 0.1$ needs to buy (sell) the asset at a price below (above) the common value to avoid a loss. Hence, these agents have the highest value for up-to-date information about v . It also follows that agents with $|\beta| = 2$ value information more than those with $|\beta| = 4$.

We therefore consider four basic variants of our model:

1. All agents are informed about the current value of v .
2. Only the traders with $|\beta| = 0.1$ are informed.
3. The traders with $|\beta| = 0.1$ or 2 are informed.
4. All agents observe v with a lag.

In each variant, we consider the payoff to each agent if they deviated at the information acquisition stage, but played optimally thereafter. For example, in the second case above, we compute the payoff to the $|\beta| = 0.1$ agents if they chose to not acquire information, and to the $|\beta| = 2, 4$ agents if they did acquire information. A comparison of the equilibrium payoffs with the payoffs to such deviations determines the value of information to each type of trader.

To determine how information acquisition and the belief that others have acquired information affects order submissions, we first compare two distinct cases: equilibrium in which all agents are informed and equilibrium in which only the $|\beta| = 0.1$ are informed. We consider three different sets of actions. First, we compare the equilibrium actions for the different trader types between the two equilibria to see how beliefs about other agents' strategies affect optimal order submission. Second, we consider optimal deviations within the equilibrium to determine how an agent's action would change if he acquired information holding the actions of the other agents fixed. Notice, that if $\beta = |0.1|$ are informed, then these traders deviate to being uninformed, while all others, $|\beta| \in \{2, 4\}$ deviate to acquiring information (becoming informed).

As our market and equilibrium are symmetric, we compare order submission strategies for $\beta \in \{0.1, 2, 4\}$. Recall, traders with positive β s value the asset more than the common value and thus are the natural buyers in the market. Of course, this decision is endogenous which we comment on later. We first provide information on the ultimate disposition of shares in Table 1. If a trader eventually executed via a market order, he demanded immediacy, since his order executed against an existing limit order. If a trader eventually

executed via a limit order, his order had previously been submitted to the market as a price contingent order and thus supplied immediacy.

β type	All informed		$ \beta = 0.1$ informed			
	Equilibrium Strategies % Market	% Limit	Equilibrium Strategies % Market	% Limit	Optimal Deviation % Market % Limit	
0.1	56.38	43.62	37.72	62.28	62.35	37.65
2	39.91	60.10	44.2	55.78	41.26	58.75
4	56.33	43.67	61.58	38.42	62.10	37.90

Table 1: **Eventual execution by order type**

If all agents are equally informed, then $\beta = 0.1$ types execute more frequently by market orders than limit orders. We hypothesize that such traders, who not have a large private incentive to trade, wait to find aggressively priced limit orders and take advantage of them.⁷ We investigate this intuition further when we consider order placement strategies.

In the equilibrium in which the $|\beta| = 0.1$ traders are the only ones who are informed, they are more likely to execute by limit orders, compared to the case in which all traders are informed. The contrast is especially strong when looking at the optimal strategy of these traders if they deviated and chose to be uninformed about the current value of v . In the latter case (shown in the “Optimal Deviation” column in Table 1), these traders are far more likely to execute via market orders. These observations are consistent with informed trade being conducted with limit orders. In an experimental setting, Bloomfield et al. (1999) find that informed traders are more likely to submit limit orders.

A unique feature of this model is that an agent may enter the market more than once. The value of reentering differs across β type and information structure, as well as the book encountered upon entry. In Table 2 we consider the average number of times that agents enter the market. The minimum number of entries is 1. Intuitively, when $|\beta| = 0.1$ have private information, then they are more likely to reenter the market than if all are informed. This is consistent with the idea that agents who have few gains to trade monitor the market for profitable trading opportunities. Interestingly if they deviate, then they are even more likely to enter the market. In this case, not knowing the true value of v , they incur higher waiting costs before finding a profitable transaction. By contrast, with the same discount rate, waiting costs are higher for $\beta = 4$, and across all the information acquisition regimes execute most quickly.

The equilibrium in which only the $|\beta| = 0.1$ traders are informed has the following curious feature. If this type chose to remain uninformed (i.e., observe v with a lag), it

⁷An aggressively priced sell order is one below the common value, and an aggressively price buy order is one above the common value.

β type	All Informed	$ \beta = 0.1$ Informed	
	Equilibrium Entries	Equilibrium Entries	Deviation Entries
0.1	2.08	3.40	4.47
2	1.75	1.82	2.14
4	1.28	1.25	1.47

Table 2: **Average number of market entries by type.**

re-enters the market more often (and thus takes longer to execute following initial entry into the market). The traders with $|\beta| = 2, 4$ are uninformed in equilibrium; if they choose to become informed, they similar extend the number of re-entries and the eventual time to execution.

In Table 3 we present more detail on the average actions taken by traders. Notice that the sum of order submissions adds up to the average entries less any non orders. The latter occurs if an agent does not have an order on the book (retained order) and chooses not to submit any, but to wait until he returns to the market.

β type	All Informed			$ \beta = 0.1$ Informed					
	Equilibrium			Equilibrium			Deviation		
	Retained	Market	Limit	Retained	Market	Limit	Retained	Market	Limit
0.1	0.55	0.56	0.98	1.25	0.38	1.76	1.18	0.63	2.41
2	0.57	0.40	0.78	0.43	0.44	0.88	0.54	0.46	1.00
4	0.20	0.57	0.51	0.13	0.62	0.48	0.18	0.60	0.56

Table 3: **Optimal Order Submissions Per Trader**

The β types who are closer to the common value of the asset are more likely to supply liquidity (submit limit orders). This effect occurs even if the trader is uninformed. Consider the last three columns of the table, which indicate the optimal strategies followed by traders if they deviate at the information acquisition stage. The deviation strategy for the 0.1 trader (if he chooses not to become informed) is to submit more limit orders. Interestingly, when not informed, this trader type also submits more market orders, suggesting that informed traders are more likely to submit limit orders. If $\beta = 2$ or $\beta = 4$ deviate and become informed, they will also submit more limit orders. The difference between order submission strategies and ultimate execution reflects the fact that execution is endogenous.

In Table 4, we consider the decision to buy or sell, and show that it is indeed endogenous. $|\beta| = 2, 4$ are invariably buyers, but $\beta = 0.1$ both buy and sell, depending on the state of the limit order book when they enter. As these traders do not derive large private benefits from trade, they appear on both sides of the market. They submit market orders if the more

extreme traders have submitted aggressive orders, and they supply liquidity at conservative prices when necessary.

Information		Order Type	β Type		
			0.1	2	4
All	Buy	Limit	0.715	0.778	0.518
		Market	0.269	0.397	0.563
Informed	Sell	Limit	0.273	0.002	0.00
		Market	0.295	0.002	0.00
$ \beta = 0.1$ Informed Equilibrium	Buy	Limit	1.00	0.877	0.482
		Market	0.260	0.442	0.616
	Sell	Limit	0.748	0.004	0.000
		Market	0.117	0.000	0.000
$ \beta = 0.1$ Informed Deviation	Buy	Limit	1.405	0.992	0.536
		Market	0.364	0.397	0.621
	Sell	Limit	1.235	0.002	0.000
		Market	0.259	0.015	0.000

Table 4: **Optimal Buying/selling and limit and market order submission per Trader**

This is reflected in the placement of limit orders relative to the common value. Notice that if agents are uninformed, then they place their limit order relative to the expected consensus value. These are reported in Table 5.

Informed	β	Limit Buys			Limit Sells		
		above v	$= v$	$< v$	below v	$= v$	$> v$
All	0.1	0.00	0.00	100.00	0.00	0.00	100.00
	2	50.04	0.00	49.96	0.00	0.00	100.00
	4	77.77	0.00	22.23	0.00	0.00	0.00
$ \beta \in \{0.1, 2\}$	0.1	0.00	0.00	100.00	0.00	0.00	100.00
	2	45.55	0.00	54.45	0.00	0.00	100.00
	4	63.06	0.00	35.94	0.00	0.00	0.00
$ \beta = 0.1$	0.1	0.00	0.00	100.00	0.00	0.00	100.00
	2	40.25	0.00	59.75	4.41	0.00	95.59
	4	59.55	0.00	40.45	0.00	0.00	0.00
None	0.1	25.05	0.00	75.95	17.23	0.00	82.77
	2	48.05	0.00	51.95	10.00	0.00	90.00
	4	60.76	0.00	39.24	0.00	0.00	0.00

Table 5: **Submission of Limit orders relative to v**

The β types with the lowest gains to trade (those with $|\beta| = 0.1$), submit very conserva-

tive orders. That is, they sell above the consensus value of the asset and seek to buy below it. This pattern is consistent across all information regimes. By contrast, those types with the highest gains to trade, $|\beta| = 4$, are more likely to post orders that are very aggressive. These trader types, essentially make price concessions to increase the probability that their order will execute.

A final measure of order submission strategy that we consider is the difference between the price of the executed order and the consensus value; that is,

$$\begin{aligned} p - v(\tau) & \text{ for a buy order} \\ v(\tau) - p & \text{ for a sell order} \end{aligned}$$

If this measure is negative then the trader bought the asset below the average valuation (but not necessarily his), and sold it above.

β type	All Informed	$ \beta \in \{0.1, 2\}$ Informed	$ \beta = 0.1$ Informed	None Informed
4	0.0957	0.3740	0.3818	0.0669
2	-0.1871	-0.1714	0.0976	-0.1215
0.1	-0.6866	-0.8309	-1.0996	-0.6497

Table 6: **Difference between the transaction price and consensus value by β type**

As Table 6 demonstrates, the traders with $|\beta| = 0.1$, on average execute on the advantageous side of the common value. This is consistent with the notion that if there are a large number of traders who monitor the market, they will trade whenever the posted price deviates from its fundamental value. Their ability to extract surplus is highest when they are the only ones who are informed. However, they still only execute advantageously when both everyone is informed and all are uninformed.

4 Market Efficiency across different information acquisition regimes

There are two notions of market efficiency we consider. In this section, we look at informational efficiency—is information about v reflected in market outcomes? In the next section, we consider efficiency in the welfare sense—how well does the market perform, relative to welfare benchmarks?

Is information acquired by investors reflected in market outcomes? Information in our model is potentially public (since it is available to all investors). Thus, we examine whether market outcomes are semi-strong efficient. Of course, in equilibrium, some (or every) subset of investors may choose to not acquire information. If no investor chose to acquire information, it remains private (we do not model the information-generation process).

The rational expectations literature (e.g. Grossman, 1976, and Admati and Pfleiderer, 1987) has typically focussed on whether transaction prices reflect privately held information. Since our market is inherently dynamic, other than the transaction price itself, other factors such as the direction of the transaction, or the number of transactions in a given period, may be informative about changes in the common value. As this is a strategic model in which agents have a private motive for trade, transaction prices may occur above or below the common value of the asset. Indeed, a trade is more likely to be initiated by a buyer (seller) if the transaction price is below (above) the true value of the asset.

On prices as well, one could consider quotes or transaction prices. Informed traders submit both limit and market orders in this model, depending on the state they face on arrival. For now, we examine transaction prices. We examine both a static measure of efficiency (the standard deviation of transaction price minus common value) and a dynamic one (after a change in the common value, how quickly do transaction prices come close to the common value).

One way to measure whether information is impounded in price is to compare transaction prices to the true common value. In a frictionless market with only informed agents, one would expect trades to occur at or close to the common value.⁸

In Table 7 we report the average time between transactions and the average characteristics of transaction price minus the common value. That is, suppose a transaction occurs at time τ . We consider the instantaneous difference between the transaction price and the common value, v_τ . Given the discrete prices and our assumption that v falls between two ticks, the smallest absolute value of this difference is $\frac{1}{2}$ tick.

Measure		All Informed	$ \beta \in \{0.1, 2\}$ Informed	$ \beta = 0.1$ Informed	All Uninformed
Time between transactions	Mean	2.005	1.994	2.006	2.006
	Std.Dev.	1.751	1.707	1.731	1.731
$\tilde{p} = p_\tau - v_\tau$	Mean	0.003	0.005	0.001	0.008
	Std. Dev.	0.748	0.903	1.169	1.543
$\Delta\tilde{p} = \tilde{p}_{\tau+1} - \tilde{p}_\tau$	Mean	0.000	0.000	0.000	0.000
	Std Dev.	0.616	0.771	0.807	0.745

Table 7: **Transaction Frequency and the difference between transaction price and the common value**

The time between transactions has a very similar mean and standard deviation across all four models. That is, there is a certain robustness to the market—regardless of information, there are gains to trade, and agents have an incentive to consummate these quickly. The

⁸All prices between -0.1 and 0.1 of the common value constitute Walrasian equilibria in this model.

average transaction price is very close to the true value of the asset, which is to be expected in a symmetric model (that is, roughly as many transactions occur above the common value as below it). The standard deviation of $(p_\tau - v_\tau)$ is decreasing in the number of informed agents. Prices in this model, therefore, do reflect information available to investors. The greater the number of informed people, the less dispersed prices are around the true common value. Finally, the standard deviation of price changes (relative to v) from transaction to transaction is also lowest when all agents are informed. However, this measure is higher when a subset of agents are informed, than when all agents are uninformed.

4.1 Speed with which information is impounded in price

A second measure of informational efficiency is the speed with which changes in the common value are reflected in transaction prices. To measure this, we do the following. Let every change in v_t be an information event. Consider isolated information events only; that is, consider information events that have the property that, for a defined interval of time both before and after the event, no other information event occurs. For this section, we consider information events that occur no sooner than ten units of real time from another change in common value. Then, a change in the common value may be interpreted as a shock to the system.

After every information event, we can track how informative the transaction price is about the event. That is, for the next n transactions, we measure $|p_{\tau(n)} - v_{\tau(n)}|$, where $\tau(n)$ reflects the time of the n^{th} transaction after the change in common value. This provides a guide to how quickly the system recovers from shocks. All four models are similar in terms of the approximate time between transactions following an information event. Hence, this measure also serves as a proxy for how close the transaction price is to the common value in real time.

Figure 1 displays the absolute difference of transaction price and common value for the first ten transactions after an information event.

The results are intuitive. With each transaction following an information event, the prices converge towards the consensus value. Recall that the smallest $|p_\tau - v_\tau|$ can be is 0.5 ticks. When all agents are informed, after five transactions, the price is within 0.6 ticks of the common value. Further, prices occur closer to the common value when there are a larger number of informed agents. When no agent is informed about v , even after ten transactions prices are over a tick away from v . When some or all agents are uninformed about the latest common value, a higher value of $|p_\tau - v_\tau|$ could reflect either uncertainty about v_τ or an unwillingness to post an order close to v_τ (to prevent being picked off by informed agents).




Figure 1: **Response of Transaction Price after Information Event**

Another measure of the speed with which information is reflected in prices is the time following an information event until $|p - v|$ first reaches 0.5 ticks. We restrict attention here to information events that are isolated enough to allow a transaction at 0.5 ticks from the common value before there is another change in common value. The results are reported in Table 8. The mean time until price is first within 0.5 ticks of the common value is lowest when all agents are informed, and increases as the number of informed agents decreases.

	All Informed	$ \beta \in \{2, 0.1\}$ Informed	$ \beta = 0.1$ Informed	Uninformed
Mean	2.503	2.727	3.762	5.550
Max	29.570	36.203	32.254	32.155

Table 8: **First time before $|p - v| = 0.5$**

4.2 Beliefs of Uninformed Traders

Another way to gauge whether information is reflected in market outcomes is by examining the beliefs of uninformed traders about the true common value. These traders observe market outcomes (specifically, in our simulation, they observe (i) the limit order book (ii) cumulative net buys and sells since their last observation of v (iii) the most recent transaction price), and update their beliefs about common value. If market outcomes reflect all privately held information about v , the beliefs of uninformed traders should be correct.

All Informed	$ \beta \in \{2, 0.1\}$ Informed	$ \beta = 0.1$ Informed	All Uninformed
0.3875	0.4369	0.6469	1.0323

Table 9: **Average absolute Difference between Belief of Uninformed Traders and True Common Value**

In Table 9, we report the mean absolute difference between trader belief about the current consensus value and the true consensus value. This is reported for the uninformed traders. In the model in which all agents are informed, these are the beliefs agents have if they deviate at the information acquisition stage. The more the number of informed agents in the market, the better the estimates of the current common value. Hence, the more agents who acquire information, the easier it is for the uninformed to predict it.

The ease with which agents can infer the common value, v , is similar to the rational expectations notion of a “revealing price.” In our simulations, therefore, market outcomes are partially revealing. In a trading game, the important difference is that individuals have traded at prices before they became revealing. This implies that the welfare properties of revealing prices in a rational expectations equilibrium and a trading game are different. Hence, the incentive to acquire information is different.

5 Endogenous information acquisition

If prices are fully revealing, then information has no value. This is Grossman–Stiglitz (1980) paradox. If agents have an unobservable private motive for trade, then price may not reveal information. The degree to which it does reveal information affects the amount strategic agents are willing to pay for information. Thus, for a given information acquisition strategy, we consider the expected consumer surplus that each agent makes in the trading game holding the information structure fixed. This is determined in the equilibrium of the trading game: that is, his trading strategy and those of the other agents when all are correct in their beliefs about others’ information acquisition strategies.

5.1 The Demand for Information

Let $\sigma_\beta = \{0, 1\}$ denote the information acquisition strategy of an agent of private type β . If $\sigma_\beta = 0$, then he does not acquire information, and if $\sigma_\beta = 1$, then he does. $\Sigma_{-\beta}$ denotes the information acquisition strategy of all agents other than θ . Ex ante, before he enters the market, a trader has beliefs about the payoff he will receive for participating in the market. Let $E[W(\beta) | \sigma_\beta, \Sigma_{-\beta}]$ be the gross expected utility (consumer surplus) of an agent of type θ and information acquisition strategy σ_β who participates in the market when the other agents follow an information acquisition strategy of $\Sigma_{-\beta}$.

All surplus numbers we report in this section are gross, in the sense that they do not account for the cost of acquiring information. This then allows us to determine the optimal information acquisition strategy for different costs.

We run the model under different information acquisition strategies for each type. To determine the marginal value of being informed holding the actions of all other agents fixed, we need to calculate $E[W(\beta) | 1, \Sigma_{-\beta}] - E[W(\beta) | 0, \Sigma_{-\beta}]$. To hold the actions of other agents fixed, for each such simulation, we allow a negligible number of agents to deviate (i.e., acquire information if their β type is uninformed or not acquire information if their β type is informed). Welfare is then computed given that such agents follow an optimal trading strategy. Thus, a deviation in the information acquisition game means both changing the information an agent acquires and changing his trading strategy, holding the information acquisition and trading strategies of other agents fixed.

We first ensure that introducing the negligible number of deviators does not substantially alter equilibrium outcomes. To investigate this, we compare such outcomes with a model in which no agents deviate. In Table 10, we demonstrate that the ex ante welfare of agents with full information, in the equilibrium in which a small mass of agents deviate (1% of each type) is similar to that in which none deviate.⁹

β type	Equilibrium with Deviation Informed	Equilibrium with no Deviation Informed	Difference
4	3.503	3.478	0.025
2	1.699	1.689	0.010
0.1	0.500	0.521	-0.021

Table 10: **Ex ante Payoffs if all agents are informed but a small number deviate**

In Table 11 we present the difference in ex ante welfare for the different β types in markets with different exogenous information structures. Notice, that the incremental value to being informed is decreasing in the absolute value of β . For $|\beta| = 0.1$ there is little intrinsic

⁹In future drafts, the no-deviation payoff will be obtained from a model in which there are no deviators.

benefit to trade. Thus, information is valuable to them as it allows them to expropriate surplus from the extreme $|\beta|$ types. Second, the benefit to information does not decrease monotonically in $|\beta|$. Third, the more agents that are informed, the lower the benefit to acquiring information for any agent.

β type	All Informed	$ \beta \in \{0.1, 2\}$ Informed	$ \beta = 0.1$ Informed	All Uninformed
4	0.151	0.054	0.071	0.028
2	0.183	0.232	0.149	0.208
0.1	0.190	0.266	0.384	0.625

Table 11: **Net Welfare gain to being informed**

If agents incur a cost, c , to be informed, then agents acquire information if the difference in payoffs between being informed and uninformed in a specific information regimes is greater than the cost. Or, if $E[W(\theta) | 1, \Sigma_{-\theta}] - E[W(\theta) | 0, \Sigma_{-\theta}] > c$. The net expected utility takes into account the information acquisition cost, c , if agents choose to acquire information. The cost can be interpreted as either the opportunity cost of acquiring information, for example reading corporate reports etc., or the direct cost of subscribing to an information service or newsletter.

Our derived demand for information differs from the literature on the value of information to a Bayesian decision maker. Radner and Stiglitz (1984) demonstrated that the marginal value of information for the uninformed is zero, resulting in a non-concavity. Their result is derived in a different framework, however it broadly relies on the idea that if information does not cause an agent to change his actions, then it has no value.¹⁰ Chade and Schlee (2000) generalize the Radner–Stiglitz result.¹¹ They argue that non-concavity of information is difficult to avoid in a model with generality. In our model, information is always less valuable to the agents with extreme β s than to those with $|\beta| = 0.1$ as the former do not change their trading strategies as much as the latter in response to having information or the existence of different information acquisition strategies. However, as all our traders are strategic, their strategies always change in response to different information allocations.

5.2 The information acquisition game

Before entering the trading market, each agent decides on his information acquisition strategy as a function of β . If he chooses to be informed, he pays a cost c . Every time he reenters

¹⁰They require that the set of states and signal realizations is finite, the optimal decision is continuous in the information partition and invariant to signal at the uninformative partition.

¹¹They show the importance of strict concavity of the utility function to generate a unique action given beliefs and weak convergence of the posterior to the prior so that beliefs do not change too much with changes in information to generative non-concavity.

the market, he will be informed. We restrict attention to equilibria in which all β types behave identically. We focus attention on subgame perfect equilibria.

The incentive to deviate from an information acquisition strategy is determined by the demand for information. Across all our simulations, the incentive for $|\beta| = 0.1$ to acquire information is highest.

Observation 1 *If any agent acquires information, then $|\beta| = 0.1$ acquires information.*

Different equilibria obtain depending on the cost of acquiring information. Further, for the same cost, different equilibria are possible. This follows immediately from the fact that the demands for information are not monotonic. We can determine equilibria up to a numerical error. We ensure that both sides of the market do not deviate.

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

$$c \in \begin{cases} [0, 0.15] & \text{everyone acquires information} \\ [0.07, 0.22] & |\beta| \in \{0.1, 2\} \text{ acquire information} \\ [0.19, 0.34] & |\beta| = 0.1 \text{ acquires information} \\ [0.62, \infty] & \text{no one acquires information} \end{cases}$$

It is immediate that these ranges can be overlapping. Thus:

Observation 3 *For $c \in [0.07, 0.15]$ there are at least two equilibria*

(i) $|\beta| \in \{0.1, 2\}$ acquire information

(ii) All agents acquire information.

Further, for $c \in [0.19, 0.22]$, there are at least two equilibria

(i) $|\beta| \in \{0.1, 2\}$ acquire information

(ii) $|\beta| = 0.1$ acquires information.

Thus, when $c \in [0.07, 0.15]$, no individual trader with $|\beta| = 4$ will deviate to being the only $|\beta| = 4$ trader acquiring information, however when all of his type acquire information then he will also do so.

The existence of multiple equilibria is important because agents behave differently in the trading game depending on how many other informed agents they are competing against. Thus, the speed with which information is incorporated into price differs.

5.3 Welfare Ranking across different Equilibria

First, we consider overall surplus, or welfare improvement, generated across different market regimes. We then examine how the gains to trade are split among agents with different private values.

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. This clearly provides an upper bound to the welfare any market mechanism can generate. Let W_f be the surplus (or welfare improvement) per trader in a frictionless market. The maximum surplus each investor can obtain when all trades are instantaneously consummated at a price of v is $|\beta|$.

There are three 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 thus submit orders strategically.

It is straightforward to account for the frictions related to the timing of agent arrivals, and to discrete prices. Ideally, a benchmark for welfare should also account for the friction introduced by incentive compatibility. We know, for example, from Myerson and Satterthwaite (1982) that, in a simple double auction setting, there is no efficient, incentive compatible, budget-balanced mechanism that exhausts all gains from trade. At any one point of time, there is a relatively small number of agents (approximately five) active in our market.

Determining 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 trade at a price $p^* = v$. However, the planner must respect the arrival times of agents, and discount accordingly. A simple trading rule is as follows: all traders with $\beta > 0$ are buyers, and those with $\beta < 0$ are 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 misrevealing his β . The LIFO rule prevents excessive discounting (by comparison, a FIFO allocation rule performs much worse). Let W_ℓ be the surplus generated per trader by such a LIFO mechanism.

Both these measures are straightforward to compute given a particular arrival sequence of traders. The frictionless benchmark can also be determined from the probability distribution of β . Given this distribution, the expected value of W_f is

$$W_f = 0.2(4 + 2 - (-2 - 4)) + 0.1(0.1 - (-0.1)) = 2.42.$$

We compute W_ℓ from the actual arrival of traders in the all-informed case. While this measure will vary slightly from simulation to simulation, over a large number of trader arrivals, it will be approximately the same.

Frictionless: W_f	LIFO: W_ℓ	All Informed	$ \beta = \{0.1, 2\}$ Informed	$ \beta = 0.1$ Informed	All Uninformed
2.42	2.09	2.19	2.17	2.12	2.18

Table 12: **Welfare per trader, and benchmarks**

The welfare per trader in each market, and the benchmarks, are shown in Table 12. For ease of comparison, we continue to report all surplus numbers gross of the cost of acquiring information. As expected, in all cases the market performs worse than the frictionless benchmark W_f . However, even in the worst case (when only agents with $|\beta| = 0.1$ are informed), it recovers 87.6% of the surplus generated in the frictionless case. Further, the market outperforms the LIFO benchmark in each case. That is, traders with extreme β values have an incentive to submit aggressive orders, 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.

Welfare does vary a little across the four markets we consider. The results suggest that adverse selection is the important friction to consider, as opposed to a lack of information itself. When all traders are uninformed about the latest common value, the surplus generated is approximately the same as when all are informed. However, it is lower when some subset are informed—in these cases, the uninformed traders suffer from adverse selection, resulting in a loss of welfare.

Next, we turn to the surplus acquired by each type of agent in our markets. For each of the possible information acquisition strategies, in Table 13 we report the average gross payoffs obtained as a percentage of the frictionless maximum.

β type	All Informed	$ \beta \in \{2, 0.1\}$ informed	$ \beta = 0.1$ informed	None informed
4	86.95	85.28	85.1	87.43
2	84.45	85.75	79.9	84.25
0.1	521.00	584.00	711.00	510.00
Average	90.14	88.33	89.46	90.23

Table 13: **Welfare accruing to each trader type as a percentage of theoretical maximum**

To determine if information helps traders in a market, we consider net payoffs; $E[W(\beta) | \sigma_\beta, \Sigma_{-\beta}] - \sigma_\beta c$. We report these in Table 14.

We compare these with the payoffs that agents receive if all agents are uninformed. We

β type	All Informed	$ \beta \in \{0.1, 2\}$ Informed	$ \beta = 0.1$ Informed	All Uninformed
4	$3.483 - c$	3.388	3.388	3.469
2	$1.694 - c$	$1.704 - c$	1.582	1.676
0.1	$0.518 - c$	$0.579 - c$	$0.689 - c$	0.472

Table 14: **Net Payoffs for different information acquisition equilibria**

observe that there is an element of the prisoner’s dilemma to the information acquisition game. For some possible cost ranges, all agents are strictly worse off if they acquire information. From inspection of the net payoffs, it is immediate that any uninformed agent in a market in which others have acquired information is worse off. Thus, any uninformed agent strictly prefers the trading game in which all others are uninformed. However, given a positive cost of information, it is possible for the agents who acquire information to be worse off. Thus, for each equilibrium information acquisition strategy, we present the possible ranges of information cost which could support it as an equilibrium and we present the minimum information price above which the informed agents would prefer to commit ex ante to not knowing information. These are presented in Table 15.

Information Partition	Cost Range	Rank Relative to All uninformed
All Informed	$c < 0.15$	$c > 0.01$, 4 prefers uninformed All others prefer uninformed
$ \beta \in \{0.1, 2\}$ Acquire information	$0.07 < c < 0.22$	$c > 0.09$, 0.1 better off uninformed 2 prefers uninformed
0.1 informed	$0.19 < c < 0.34$	$c > 0.21$, 0.1 prefers uninformed

Table 15: **Equilibrium payoffs relative to uninformed equilibrium.**

That information could hurt agents was first observed by Hirschleifer (1971) who showed that risk averse agents who participated in risk sharing markets became worse off with more information (i.e., about the realization of the state), a result extended by Schlee (2000).¹² In such models, if payoffs depend on the state, and if the Nash equilibrium is not Pareto efficient, then agents may be better off if do not know the realization and may reach the Pareto efficient outcome. These papers observe that adverse selection can lead to possible market breakdown. By contrast, in our model trades are still consummated. However, the existence of informed traders changes trading strategies, leading to longer execution time and thus reducing payoffs.¹³

¹²In any endowment economy where agents are weakly risk averse, some agents are risk neutral or there exists a representative agent, more information is Pareto inferior.

¹³Bassan, Gossner, Scarsini and Zamir (2003) provide conditions under which all agents prefer more

6 Conclusion

Watching the stock market and researching markets have become a national pastime, presumably because the cost of acquiring information and the cost of access to the market have declined. Our results suggest that this might not be Pareto efficient: Too much information might be collected in financial markets. Agents would be better off if they did not invest in information, but given that others have done so, to acquire information is a best response. Thus, agents are faced with the prisoner's dilemma.

Who has acquired information affects agents trading strategies, and how the gains from trade are split. This can be measured by the discount or premium agents pay over the common value of the asset. Further, the speed with which transaction prices reflect information depends on how many agent have acquired information. We demonstrate how the decision to acquire information depends on agents' beliefs about other agents' information acquisition strategies. Thus, for the same information cost, different groups of investors may acquire information. As trading strategies are jointly determined by private gains to trade and agents beliefs about the information others have, the information content of orders differs by these unobservables. This suggests that econometric strategies that seek to infer how information enters a market place are difficult to identify.

information. If a game with a particular information structure has a unique Pareto efficient equilibrium, then in all coarser partitions more information is preferred.

References

- [1] Admati, Anat, P. Pfleiderer (1986) “A Monoplist Market for Information,” *Journal of Economic Theory*, 39, 400–438
- [2] Admati, Anat, P. Pfleiderer (1987a), “Viable Allocations of Information in Financial Markets,” *Journal of Economic Theory* 43 p 76-115.
- [3] Admati, Anat, P. Pfleiderer (1987b), “Selling and Trading on Information in Financial Markets,” *American Economic Review* Vol 78 no 2 Papers and Proceedings, p96–103.
- [4] Admati, Anat R. and Paul Pfleiderer (1990), “Direct and Indirect Sale of Information,” *Econometrica* Vol 58 No 4. 901-928.
- [5] Back, Kerry, C. Henry Cao, Gregory A. Willard (2000), “Imperfect Competition among informed traders,” *Journal of Finance* Vol 55, No 5, p 2117 –2155.
- [6] Barlevy, Gadi and Pietro Veronesi (2000), “Information Acquisition in Financial Markets,” *Review of Economic Studies*.
- [7] Baruch, Shmuel, (2003), “Who benefits from an Open Limit Order book?” *Journal of Business* forthcoming.
- [8] 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.
- [9] Bernardo, Antonio and Kenneth Judd, (1997) “Efficiency of Asset Markets with Asymmetric Information,” CMS working paper #16-97
- [10] Bernardo, Antonio and K. Judd (1997), “Efficiency of Asset Markets with Asymmetric Information,” CMS working paper #16-97.
- [11] Biais, Bruno (1993), “Price Formation and Equilibrium Liquidity in Fragmented and Centralized Markets,” *Journal of Finance* Vol 48 No 1. p157–185.
- [12] Bloomfield, Robert and M. O’Hara (1999): “Market Transparency: Who wins and who loses?” *Review of Financial Studies* 12(1) 5-35.
- [13] Bloomfield, Robert and Maureen O’Hara (2000): “Can Transparent Markets Survive?” *Journal of Financial Economics* 55, 425–59.
- [14] Blume, Lawrence and David Easley (1991): “Implementation of Walrasian Expectations Equilibria,” *Journal of Economic Theory* 51: 207–227.

- [15] Boehmer, Ekkehart, L. Yu and G. Saar (2004) “Lifting the Veil: An Analysis of Pre-Trade Transparency at the NYSE,” *Journal of Finance* forthcoming.
- [16] Chade, Hector and E. Schlee (2000), “Increasing returns in the Value of information,” world congress contributed papers.
- [17] den Haan, Wouter J., and Albert Marcet (1994), Accuracy in Simulations, *Review of Economic Studies*, 61, 3–17.
- [18] Foucault, Thierry, Eugene Kandel and Ohad Kadan (2002), The Limit Order Book as a Market for Liquidity, HEC working paper.
- [19] Foucault, Thierry, Sophie Moinas and Eric Thiessen (2003), “Does Anonymity Matter in Electronic Limit Order Books?” working paper HEC Paris.
- [20] Glosten, Larry and Paul Milgrom (1985), Bid, Ask, and Transaction Prices in a Specialist Market with Heterogenously Informed Traders, *Journal of Financial Economics*, 13, 71–100.
- [21] Goettler, Ronald, Christine Parlour and Uday Rajan (2003), “Equilibrium in a Dynamic Limit Order Market,” working paper Carnegie Mellon University.
- [22] Grossman, Sanford and Joseph Stiglitz (1980), “On the impossibility of Informationally efficient markets,” *American Economic Review* Vol 70 p 393–408.
- [23] Hakansson, Nils H., J. Gregory Kinkel and James A. Ohlson (1982), “Sufficient and Necessary Conditions for Information to have Social Value in Pure Exchange,” *Journal of Finance* Vol 37, No 5. P 1169–1181.
- [24] Hirshleifer, Jack (1971), “*American Economic Review* Vol 61 No 4 p 561-574.
- [25] Holden, Craig and A. Subrahmanyam (1992), “Long-lived Private information and imperfect competition,” *Journal of Finance*, Vol 47, No 1 p 247–270.
- [26] Kyle, Albert (1985), Continuous Auctions and Insider Trading, *Econometrica* 53 p 1315-1336.
- [27] Leschinskii, Dima (2001), “Does it pay to voluntarily disclose private information?” HEC working paper.
- [28] Mendelson, Haim and Tunay Tunca (2004), “Strategic Trading, Liquidity, and Information Acquisition,” *Review of Financial Studies* Vol 17 No 2. p 295-337.

- [29] Moscarini, Giuseppe and Lones Smith, “The Law of Large Demand for Information,” *Econometrica*
- [30] O’Hara, Maureen (1995), “Market Microstructure Theory,” Blackwell.
- [31] Pagano, Marco and Ailso Roell (1996), “Transparency and Liquidity: A comparison of Auction and Dealer Markets with informed trading,” *Journal of Finance* Vol 51 No2. 579-611.
- [32] Pakes, Ariel, P. McGuire (2001), ‘Stochastic Algorithms, Symmetric Markov Perfect Equilibrium, and the ‘Curse’ of dimensionality,’ *Econometrica* Vol 69 No 5. 1261-1281.
- [33] 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.
- [34] Simonov, Andrei (1999), “Competition in Markets for Information,” Insead working paper.
- [35] Schlee, Edward, E., ‘The Value of Information in Efficient Risk Sharing Arrangements,’ *American Economic Review* 91 (3) p501-524.
- [36] Schlee, Edward E. and Hector Chade, “Another Look at the Radner–Stiglitz Non–concavity in the Value of Information,” *Journal of Economic Theory*
- [37] Spiegel, Matthew, Avandhar Subrahmanyam (1992), “Informed Speculation and Hedging in a Noncompetitive Securities Market,” *Review of Financial Studies* Vol 5., No 2. 307-329.
- [38] Taub, Bart, D. Bernhardt and P. Seiler, (2004) “Cladistic Asset Pricing,” University of Illinois working paper.