Welfare in a Dynamic Limit Order Market *

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

GSIA
Carnegie Mellon University
Pittsburgh, PA 15213

April 28, 2003

*We have benefitted from conversations with Burton Hollifield, Patrik Sandås, Mark Ready, Duane Seppi, Chester Spatt, Tom Tallarini and seminar participants at Madison, Wisconsin and GSIA. We are especially grateful to Tony Smith for ideas generated in a project on simulating limit order books. The current version of this paper is maintained at http://sobers.gsia.cmu.edu/papers/surplus.pdf

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

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

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

Abstract

We provide an algorithm for solving for equilibrium in a dynamic limit order market. We formulate a limit order market as a stochastic sequential game and use a simulation technique based on Pakes and McGuire (2001) to find a stationary equilibrium. Given the stationary equilibrium, we generate artificial time series and perform comparative dynamics. As we know the data generating process, we can compare transaction prices to true or efficient prices. Due to the endogeneity of order flow, the midpoint is not always a good proxy for the true price. We also find that the effective spread is negatively correlated with transactions costs and uncorrelated with welfare. We explicitly determine investor welfare in our numerical solution. As one policy experiment, we evaluate the effect of changing tick size.

1 Introduction

It is well known that order flow is endogenous, and this endogeneity has to be taken into account when drawing inferences from trade data. On observing an order, we can infer something about a trader's willingness to pay. This depends on two things: the common or consensus value of the asset and any private benefits to trade. The first arises because ownership of a security has value because of the dividend stream. The second could arise because of liquidity motives for trade and must exist if we observe trade in the presence of frictions, else agents are not rational. As both jointly determine a trader's willingness to trade, in equilibrium both must be simultaneously inferred. Both are of interest first to determine the "value," of the asset being traded, the second because it allows us to compute welfare. Hence, a structural model of trade in a limit order market can be used to shed light on the relationship between observables (trades and quotes) and both the common value and trader welfare. In effect, such a model could allow one to distinguish between the economic fundamentals and microstructure effects.

We try to disentangle these two effects in a dynamic pure limit order market in which traders choose between buy and sell orders, and market and limit orders. We show that the endogenous choice of orders has important implications for inferences we can draw from the data. In a pure limit order market, agents choose to supply liquidity if the reward is high and demand it if it is cheap. As it is costless for agents to switch between limit and market orders many of the standard intuitions about intermediated markets are inappropriate. In this paper we focus on pure limit order markets (i.e., with no market-making intermediaries). Some exchanges are pure open limit order markets such as Paris, Tokyo, Stockholm, and Vancouver. Other exchanges such as the NYSE or Nasdaq have incorporated limit order books into their market design.

We numerically solve for the equilibrium of the model, and generate time series of trades and quotes. A unique feature of our model is that it enables explicit welfare comparison across different policy regimes. We characterize equilibrium in terms of traders' strategies, transactions costs of market orders, and welfare accruing to both market and limit orders. We generate empirical patterns that have been observed in the previous literature. Our structural models allows us to determine the effect of information on limit order submission and quotes. We find that any measure of realized transaction costs does not proxy well for true transaction costs.

We present a model of trade in which, in addition to a common value for assets, all agents have a private value[1], and a maximal number of shares they can trade. Given the limit

[1] Intuitively, private values may reflect idiosyncratic wealth shocks, tax exposures, or hedging motives.

order book and common value (which are both publicly observed), agents decide whether to buy or sell (or both), and at what prices. In our model traders arrive sequentially and submit orders to maximize their expected surplus given their private value and the current limit order book. Expected surplus for a limit order is computed using beliefs about the order's execution probability and the expected change in the asset's value conditional on the order executing. Market orders execute with probability one while limit orders execute with probability less than one (each limit order on the book is cancelled exogenously with some small probability). In equilibrium, order submission strategies generate actual execution probabilities and "picking off risk" that match traders' beliefs. This "picking off risk" arises as limit buys execute more often when the value drops and limit sells execute more often when the value increases. The equilibrium is symmetric by construction. The equilibrium is also Markov perfect in that a given trader's strategy is a function of only the current book—past traders' actions do not matter, other than in their determination of the current book.

Since the model is analytically intractable, we numerically solve for the equilibrium. The model is essentially a multi-agent dynamic game in which each agent chooses an action only upon entry to the market. Nonetheless, each agent's optimization problem may be expressed recursively using Bellman's equation. Value function iteration then yields the fixed point in beliefs that defines a Bayesian Nash equilibrium. Complicating this approach is the size of the state space, due to the number of permutations of limit order books a trader may face. Consider a market with only seven prices (or ticks) and up to nine buy or sell orders at each price. Suppose the lowest sell is at tick 1. Then, the number of possible books is 10^7. But the lowest sell could be at any of 7 ticks, or there might be no limit sells on the book. Hence the total number of books is 8×10^7. Allowing 19 buys or sells at each tick increases the number of books to more than 10^{10}. Of course, most of these books never arise when traders play equilibrium strategies. Hence, we deal with this curse of dimensionality by obtaining equilibrium values, beliefs, and strategies only on the subset of states in the *recurrent class* of states. Pakes and McGuire (2001) apply the same technique when solving an equilibrium model of industry dynamics.

Having solved for the equilibrium, we characterize it by simulating 500,000 order arrivals for different values of the key parameters. The equilibrium displays order persistence of the sort documented by, e.g., Biais, Hillion, and Spatt (1995) for the Paris Bourse. Some persistence (such as that of small buy or small sell orders) occurs even when there is no change in the consensus value of the asset, suggesting that there are particular sets of states at which such orders are especially preferred.

We further find that on average market buy (sell) orders are submitted when the ask is below (the bid is above) the consensus value of the asset. This reflects the fact that order flow responds to profitable trade opportunities. An implication of this is that conditional on a trade occurring, the midpoint is not a good proxy for the consensus value of the asset. Or rather, that the estimate can be improved by conditioning on the trade. Another implication is that in a pure limit order market transaction costs are negative, as market orders are only submitted when it is beneficial. Thus, measures that were developed for an intermediated market (such as the effective spread) should be interpreted with caution.

The effective spread is defined as the average transaction price minus the midpoint of the contemporaneous bid and ask quotes. It is also used implicitly as a measure of welfare when evaluating policies that affect markets. Using our model, we examine the efficacy of this in a pure limit order market. If two markets have the same distribution over order size and the same supply of liquidity, then effective spread is indeed a good measure of which market generates more consumer surplus. However, if demand or supply schedules, or order sizes are endogenous, then welfare may actually decrease when spreads narrow. In our model the effective spread is negatively correlated with (true) transactions costs, and largely uncorrelated with investor welfare. For example, we find that changing the tick size has little impact on traders' welfare despite increasing the average effective spread.

The literature to date has not provided a tractable model of a dynamic limit order book. Our understanding of the trade-offs involved in submitting limit orders have been enhanced by Cohen, Maier, Schwartz and Whitcomb (1981), Holden and Chakarvarty (1998), and Kumar and Seppi (1993) who analyze a trader's optimal choice between market and limit orders in different trading environments. Biais, Martimort and Rochet (1999), Foucault (1999), Glosten (1994), O'Hara and Oldfield (1986), Parlour (1998), Rock (1996), Seppi (1997) and Foucault, Kadan and Kandel (2002) theoretically analyze prices, trading volumes and efficiency in financial markets with limit order books.

Of these papers, Parlour (1998), Foucault (1999) and Foucault, Kadan and Kandel (2002) are explicitly dynamic. However, these models require restrictive assumptions to obtain analytical solutions. Parlour (1998) assumes a 1-tick market and no volatility in the common value of the asset, Foucault (1999) allows for volatility of the common value of the asset, but truncates the book. Foucault, Kadan, and Kandel (2002) have an interesting interpretation of the cost of immediacy but for tractability, require limit order submitters to undercut existing orders, as opposed to joining a queue. Ideally, for policy work we would like a model with multiple prices and books of varying thickness.

An interesting empirical literature has shed light on both the characteristics of observed limit order books, and on the intuition we have learned from models. In the first category, Biais, Hillion and Spatt (1995) present an analysis of order flow in the Paris Bourse. In the latter category, Sandås (2001) uses data from the Stockholm exchange, a pure limit order market, to develop and test static restrictions implied by Glosten (1994). He tests break even conditions from the static model. Consistent with the notion that market order flow is endogenous, he finds that market order flow depends on state variables. He strongly rejects the restrictions of the static model, suggesting that a dynamic one is needed to explain both price patterns and orders in a limit order market. In short, he suggests that any inference based on models that do not take into account the endogeneity of order flow may be flawed.

Hollifield, Miller and Sandås (2002) test a monotonicity condition generated by a limit order market in which traders cannot split their orders on Swedish data. They reject the condition if they consider both buy and sell orders, and fail to reject when conditioning only on one side of the market. This suggests that one can appropriately view a limit order market as a market for immediacy: i.e., a trader's decision is not merely to buy or sell the security but also immediacy. Hollifield, Miller, Sandås and Slive (2002) use a similar technique to investigate the demand and supply of liquidity on the Vancouver exchange. In keeping with the results in this paper and those in Foucault, Kadan and Kandel (2002) they find that agents supply liquidity when it is dear and consume it when it is cheap.

There is a literature, pioneered by Roll (1984), Glosten (1987), Glosten and Harris (1999) and Hasbrouck, (1991a,1991b, 1993) that considers the relationship between quoted spreads, transaction prices and the true or consensus value of the asset in the presence of an intermediary.[2] Our work is complementary to this as we generate artificial data from a pure limit order market and thus know the true or consensus value. We can therefore consider some of the same issues albeit in a different market environment.

We provide a general framework in Section 2. The details of our solution technique appear in Section 3. We present the results of our model in Section 4 and attempt to disentangle private and common values in section 5 and then present the results of our policy experiment in Section 6. Section 7 concludes.

2 Model

We present an infinite horizon version of Parlour (1995). This is a discrete time model of a pure limit order market for an asset. At each period, t, a single trader arrives at the market place. The trader at time t is represented by a pair, $\{z_t, \beta_t\}$. Here, $z_t \in \{1, \ldots, \bar{z}\}$ is an integer that denotes the maximum quantity of shares the trader may trade. The trader may buy or sell any combination up to z_t shares. Thus, the decision to be a buyer or a seller is endogenous. Let F_z denote the distribution of z_t. The trader's private valuation for

[2] More recently, Huang and Stoll (1997).

the asset, β_t, is drawn from a continuous distribution F_β. Both z and β are independently drawn across time, and their distributions are common knowledge. We normalize the mean of β to zero.

In addition, there is a consensus value of the asset. The consensus (or common) value of the asset at time t, which we denote v_t, is public knowledge. Each period, with probability $\frac{\lambda}{2}$, the consensus value jumps up by one tick, and with the same probability jumps down by one tick. Changes in the consensus value could come about because of new information that is released either about the firm or about the economy. The periodic innovations in this value mean that traders who arrive at $\tau > t$, are better informed than the traders at time t. Thus, this is a model of asymmetric information.

The market place is an open electronic limit order book. Let L_t denote the limit order book at time t. There is a finite set of discrete prices, denoted as $\{p^{-(N)}, p^{-(N-1)}, \ldots, p^{-1}, p^0, p^1, \ldots, p^{N-1}, p^N\}$. The distance between any two consecutive prices p^i and p^{i+1} is a constant, d, and we refer to it as "tick size." For convenience, prices are denoted relative to the consensus value v_t, and p^0 is normalized to 0 at each t. That is, an order to buy 1 share that executes at price p^i requires the buyer to pay exactly $v_t + p^i$.

Associated with each price $p^i \in \{p^{-(N-1)}, \ldots, p^{N-1}\}$, at each point of time t, is a backlog of limit orders, ℓ_t^i. We adopt the convention that buy orders are denoted as a positive quantity, and sell orders as a negative one. The book, therefore, is a vector of limit orders, so that $L_t = \{\ell_t^i\}_{i=-(N-1)}^{N-1}$. At prices outside this set, there is a competitive crowd of traders providing an infinite depth of buy orders (at a price p^{-N}) or sell orders (at a price p^N).[3]

The trader who arrives at time t takes an action X_t. X_t is a vector with typical element x_t^i, that denotes the integer number of shares to be traded at price p^i. An action is feasible if $\sum_{i=-N}^{N} |x_t^i| \leq z_t$. A buy (sell) order at price p^i is denoted by $x_t^i > 0$ ($x_t^i < 0$).

We draw the maximum amount, z_t, that a trader can trade. As we do not restrict a trader to be a buyer or seller, but allow them to respond to profitable trading opportunities, in equilibrium traders demand liquidity when it is cheap and supply it when it is dear. Finally, we could endow traders with a maximum dollar amount they could trade. Thus, when the price goes down, traders can buy more shares, however, for sellers when prices are high, this would restrict the amount that they sell. To retain symmetry, and to focus on demand and supply of liquidity (not shares) we endow the traders with shares.

Market orders submitted at time t execute in that period. Limit orders submitted at time t execute if a counter party arrives at some time in the future. Following Hollifield, Miller, and Sandås (2002), in each period, each share in the book is cancelled exogenously

[3]This truncation is a feature of Seppi (1997) and Parlour (1998).

with some probability. We assume this probability, δ, is constant and independent across shares.

Ideally, one would want a model in which the decision to exit the market is endogenous (i.e., δ is endogenous). However, such a model must take into account the other markets in which an agent trades, i.e., the opportunity cost of trade. As this is a model of trade in one asset, the decision to cancel is reasonably modelled as exogenous.

Suppose $\ell_t^i < 0$; that is, at price p^i, there is a backlog of sell orders at time t. If $x_t^i > 0$ and $x_t^i \leq |\ell_t^i|$, then the agent who arrived at time t has submitted a market buy order. The buy order x_t^i executes immediately against the outstanding limit sell orders. Similarly, a market sell order at price p^i requires $x_t^i < 0$, $\ell_t^i > 0$, and $|x_t^i| \leq \ell_t^i$. Any sell order at price p^{-N}, or any buy order at price p^N, is (trivially) a market order.

Conversely, if the submitted order is on the same side as the depth in the book, it is a limit order at that price. That is, if $x_t^i > 0$, so that the agent submits a buy order at price p^i, and $\ell_t^i \geq 0$, so that there are already outstanding buy orders at that price, then x_t^i represents a limit order. Similarly for limit sell orders.

Finally, agents may submit orders that are in part market orders, and in part limit orders. Suppose $\ell_t^i < 0$, and a buy order $x_t^i > 0$ is submitted. Then, if $|\ell_t^i| < x_t^i$, we say the agent submits $|\ell_t^i|$ as a market order, and the remainder $x_t^i - |\ell_t^i|$ as a limit order.

The agent is allowed to submit no order (i.e., submit an order of 0 shares). An agent may arrive in the market, and decide that, given her type and the current book, she is better off not submitting an order. In addition, she may submit both buy and sell orders. The decision to trade is endogenous, with respect to both the quantity and the direction of the trade.

2.1 Evolution of the Limit Order Book

The limit order book at time t, in conjunction with the orders submitted by the trader at time t and the exogenous cancellation rate, generates the book at time $t+1$. We now determine how the book at time $t+1$ evolves from the book at time t for the arbitrary (not necessarily equilibrium) action of the trader at time t, denoted X_t.

At each time t, the following sequence occurs. First, a trader enters and takes an action X_t. Given this action, the cumulative shares listed at price p^i are now $(\ell_t^i + x_t^i)$. This holds regardless of whether x_t^i represents a limit or market order (that is, even when ℓ_t^i and x_t^i have opposite signs). After the orders x_t^i have been submitted (and executed, if they are market orders), each remaining share at price p^i is cancelled with exogenous probability δ. Consider the cancellations at time t. Fix a price p^i, and suppose that there are ℓ_t^i shares on the book at this price. These may be buy or sell orders. Each share at this price drops out with an independent and exogenous probability δ. Let $\mathbf{1}_{jt}^{i}$ be an indicator function that takes the value 0 if the j^{th} share at this price and time is cancelled, and 1 if it does not.

Then, the evolution of the book from time t to $t+1$ can be represented as follows. Let $L_t = (\ell_t^{-(N-1)}, \ldots, \ell_t^{N-1})$. For each i and t,

$$\ell_{t+1}^{i} \quad = \quad \text{sign}\left(\ell_t^i + x_t^i\right) \sum_{j=1}^{|\ell_t^i + x_t^i|} \mathbf{1}_{jt}^{i}. \tag{1}$$

Now, $L_{t+1} = (\ell_{t+1}^{-(N-1)}, \ldots, \ell_{t+1}^{N-1})$. Figure 1 illustrates the sequence of events for ticks $-1, 0, 1$.

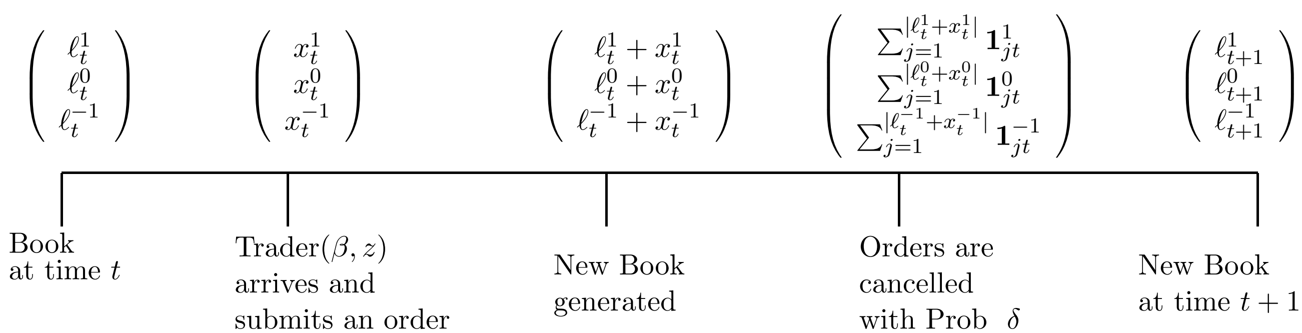

Figure 1: **Evolution of a three tick book**

Now, suppose that at the end of period t, the consensus value of the asset increases from v_t, by one tick. Since all prices are denoted relative to the consensus value, all orders at a price p^i are now listed at price p^{i-1}. That is, such orders are now one tick lower relative to the consensus value. In this process, sell orders at price $p^{-(N-1)}$ will now be listed at p^{-N}, and are automatically crossed off against the crowd willing to buy at that price. Any buy orders that were at $p^{-(N-1)}$ prior to the jump are cancelled.

Similarly, if the consensus value of the asset falls by one tick at the end of period t, all orders previously listed at a price p^i are now listed at a price p^{i+1}, one tick higher relative to the consensus value. Thus, limit orders may execute at a price closer to or further away from the consensus value. In equilibrium this is one of the potential costs of submitting a limit order: orders are more likely to execute if the asset value moves against them.

Limit orders are executed according to time and price priority. Buy orders are accorded priority at higher prices, and sell orders at lower ones. If two or more limit orders are at the same price, time priority is in effect: the one that was submitted first is crossed first. Therefore, an order executes if no other orders have priority, and a trader arrives who is willing to be a counter-party.

Actions of subsequent traders affect the priority of any limit order. Of course, the ultimate change in priority is execution—a counter-party takes the trade. A trader who arrives after an unexecuted limit order can either increase or decrease the price priority. A subsequent trader decreases an existing order's price priority if he submits a competing order closer to the quotes. This moves the unexecuted order further back in the queue. Conversely, a subsequent trader could execute against an order with price priority over the limit order. This moves the limit order toward the front of the queue. Finally, a subsequent order could improve the time priority of the unexecuted order by crossing against orders in the book at the same price, picking off orders with higher time priority. However, it is impossible for a subsequent order to decrease the time priority of a limit order. An agent who submits a limit order is guaranteed a place in the queue at his chosen price. Given that an opposing trade occurs at that price, the agent's order will be executed in sequence.

The per share payoff at time τ to a trader with type β who submits a limit order order at time t at price p^i is

$$\begin{cases} p^i - (v_\tau - v_t) - \beta & \text{if he sells the asset at } p^i \text{ at any time } \tau \geq t \\ \beta + (v_\tau - v_t) - p^i & \text{if he buys the asset at } p^i \text{ at any time } \tau \geq t \\ 0 & \text{if the share is cancelled at any time before it is executed} \end{cases} \tag{2}$$

2.2 Transaction Costs and Welfare Measures

We are interested in both transaction costs and trader welfare. The first is a measure of the distance of the transaction price from the efficient price, and the latter measures the distance between the transaction price and a trader's willingness to pay. We define these for a fixed limit order book. Suppose the book at time t is given by L_t. All measures defined below are conditional on this book; for convenience, we suppress this dependence in the notation.

The bid and ask prices in the market at time t are defined in standard fashion. In any period the ask price is the lowest sell price on the book, and the bid price is the highest buy price on the book. Therefore,

Definition 1 *The current bid and ask prices in the market are given by:*

$$\begin{aligned} B_t &= v_t + \max\{\, p^i \mid \ell_t^i < 0 \,\} \\ A_t &= v_t + \min\{\, p^i \mid \ell_t^i > 0 \,\} \end{aligned}$$

Further, the midpoint of the bid and ask prices is

$$m_t = \frac{A_t + B_t}{2}.$$

Next, consider the transaction costs paid at time t by a market order submitter who submits a market order of size $\bar{x}$. If the market order is large, it may "walk the book," so that different shares execute at different prices. Suppose that x_t^i shares execute at price p^i, with $\sum_{i=-N}^{N} x_t^i = \bar{x}$. Then, the average execution price for the shares is $P_t(\bar{x}) = v_t + \frac{1}{|\bar{x}|} \sum_{i=-N}^{N} p^i x_t^i$. The average execution price for a sell order is found analogously. Since $\bar{x} < 0$ for a market sell order, we use $|\bar{x}|$ in the expression for average price.[4]

Thus, we can define the total transaction costs paid by a market order submitter.

Definition 2 *The true transactions cost faced by a market order of size $\bar{x}$ at time t is*

$$T_t(\bar{x}) \quad = \quad (P_t(\bar{x}) - v_t)\,\mathrm{sign}(\bar{x}). \tag{3}$$

The effective spread, $S_t(\bar{x})$, faced by a market order of size $\bar{x}$ at time t is

$$S_t(\bar{x}) \quad = \quad (P_t(\bar{x}) - m_t)\,\mathrm{sign}(\bar{x}). \tag{4}$$

Notice that in many econometric specifications (see Hasbrouck (2002) for a summary), the transaction price is taken to be the sum of the efficient price (which we call the consensus value) and microstructure effects. Thus, our transaction cost is simply the microstructure effects times the signed order flow in these specifications.

A commonly used proxy for this transaction cost is the effective spread. From our definition, the effective spread is simply the transaction cost with the midpoint of the quotes as a proxy for the consensus value. If a market buy order is small, so that it transacts at the ask and does not go deeper into the book, the effective spread reduces to $(A_t - m_t)$. Similarly, a market sell order that transacts at the bid has an effective spread of $(m_t - B_t)$.

Consider a trade that occurs at time t. We define the welfare of the market order submitter, and the limit order submitter(s) who executed at t. Recall that $\bar{x} > 0$ indicates a market buy order, and $\bar{x} < 0$ a market sell order.

Definition 3 *Consider a trade of $\bar{x}$ shares at t. Then,*
(i) the welfare of the market order submitter is

$$W_t^m \quad = \quad \bar{x}\left(\beta_t + v_t - P_t(\bar{x})\right),$$

where P_t is the average execution price.
(ii) the welfare accruing to limit order submitters taking the other side of the transaction is

$$W_t^l \quad = \quad \bar{x}\left(P_t(\bar{x}) - (v_t + \beta_t^l)\right),$$

where β_t^l is the share-weighted average of the private values of all limit order submitters whose orders trade against the market order at time t.

[4] It is never optimal for a trader to submit both market buy and market sell orders.

If the midpoint of the bid-ask spread, m_t, is equal to the consensus value of the asset, so that $m_t = v_t$, the effective spread is a good proxy for the transactions cost paid by a market order submitter. Further, when $m_t = v_t$, the welfare of a market order submitter can be written in terms of the effective spread. This is the basis for the use of the effective spread to evaluate welfare. However, if $m_t \neq v_t$, this is no longer true.

Proposition 1 *Suppose a trade of size $\bar{x}$ occurs at time t incurring an effective spread of $S_t(\bar{x})$. If (and only if) $m_t = v_t$,*
(i) the welfare of the market order submitter is $W_t^m = \bar{x}\ \beta_t - |\bar{x}|S_t$.
(ii) the welfare accruing to the limit order submitters who trade at t is $W_t^l = |\bar{x}|\ S_t(\bar{x}) - \bar{x}\ \beta_t^l$.

Proof

(i) The welfare of the market order submitter is $W_t^m = \bar{x}(\beta_t + v_t - P_t(\bar{x}))$. From equation (4), if the market order is a buy order, $P_t(\bar{x}) = m_t + S_t(\bar{x})$, and if it is a sell order, $P_t(\bar{x}) = m_t - S_t(\bar{x})$. Hence, for a buy order,

$$W_t^m = \bar{x}\ (\beta_t + v_t - m_t - S_t(\bar{x})).$$

Hence, $W_t^m = \bar{x}(\beta_t - S_t)$ if and only if $m_t = v_t$.

Similarly, for a sell order, $W_t^m = \bar{x}\ (\beta_t + v_t - m_t + S_t(\bar{x})$, and $W_t^m = \bar{x}\ (\beta_t + S_t(\bar{x}))$ if and only if $m_t = v_t$. Putting together the expressions for buy and sell orders, we have $W_t^m = \bar{x}\ \beta_t - |\bar{x}|S_t$ if and only if $m_t = v_t$.
(ii) Next, consider the welfare accruing to the limit order submitters who trade at t. This is $\bar{x}(v_t + P_t(\bar{x}) - \beta^l)$. In a similar fashion as in part (i), we obtain $W_t^l = |\bar{x}|\ S_t - \bar{x}\ \beta^l$. ■

Notice that, even if $m_t \neq v_t$, the aggregate welfare improvement as a result of the trade at t is $\bar{x}(\beta_t - \beta_t^l)$, which is independent of the effective spread, $S_t(\bar{x})$. That is, if one also considers limit order traders in welfare calculations, these transaction costs become irrelevant: in a pure limit order market, these costs are simply transfers between agents. Thus, any measure which determines a cost to one party merely reflects a gain to the counter-party.

We report welfare for both market orders and limit order submitters. For policy purposes, the welfare of limit order submitters should also be considered. Typically, the literature has computed transaction costs for market orders. However, there is no reason why one group of investors should be favored over another. We do not have an intermediary: Every trade in our model consists of a market order executing against a limit order. In a market with an intermediary market-maker, transaction costs may be an important determinant of retail investor (both market and limit order submitter) welfare. While the intermediary provides a benefit by providing liquidity to market orders, it may also deter limit order submission and thus decrease the welfare of such agents (see Seppi, 1997). As Glosten (2000) observes in this case, one should account for the welfare of all parties in the market.

3 Equilibrium

In section 2 we modelled a limit order market as a stochastic sequential game. We now characterize best responses in this game, discuss the existence of a Markov perfect equilibrium, and present an algorithm for numerically finding such an equilibrium.

3.1 Best Responses

In period t, a trader endowed with type (z_t, β_t) arrives to the market and submits an order X_t specifying the number of shares to buy or sell at each price $\{p^{-N}, \ldots, p^N\}$. He observes the current market conditions, which consist of the current consensus value, v_t, and the current limit order book, $L_t = (l_t^{-(N-1)}, \ldots, l_t^{N-1})$. Recall from section 2 that the trader also knows the (exogenous) order cancellation rate, denoted δ, the jump probability for v_t, denoted λ, and the stationary distributions of types given by F_z and F_β.

Of course, the trader does not know the future sequence of trader types, order cancellations, and changes in consensus value. This sequence determines whether his limit orders execute, as well as the value of any such trades (since v_t may change before execution). Hence, the trader forms beliefs about the probability of execution of an order placed at any price p^i and the change in v_t conditional upon execution at this price.[5]

Let $\mu_{t-1}^e(k, i, L_t, X_t)$ denote the period t trader's belief of the probability of execution of his k^{th} share at price p^i given book L_t and order X_t. Similarly, let $\mu_{t-1}^v(k, i, L_t, X_t)$ denote his belief of the (expected) net jumps in the consensus value prior to execution (conditional upon execution). The distinction between market and limit orders is handled by beliefs. Suppose $x^i > 0$, indicating some type of buy order. If $l_t^i \leq -k$ then this share is a market order with $\mu_{t-1}^e(\cdot) = 1$. Furthermore, since market orders execute immediately, $\mu_{t-1}^v(\cdot)$ is necessarily zero for market orders.

[5] Recall that we normalize prices to be relative to the *current* v_t and therefore shift L_t after jumps. Hence, the belief about the "change in v_t conditional upon execution" is equivalently a belief about "the number of shifts in the book between t and execution."

Given these beliefs, the risk-neutral trader optimally chooses[6]

$$X_t \quad = \quad \underset{\tilde{X}=(\tilde{x}^{-N},\ldots,\tilde{x}^{N})}{\arg\max} \quad \sum_{i=-N}^{N} \sum_{k=1}^{|\tilde{x}^i|} \mu^e_{t-1}(k,i,L_t,\tilde{X})\ (\beta_t + \mu^v_{t-1}(k,i,L_t,\tilde{X}) - p^i)\ \text{sign}(\tilde{x}^i)$$

$$\text{subject to:} \quad \sum_{i=-N}^{N} |\tilde{x}^i| \leq z_t.$$

A strategy for an agent at time t, therefore, is a mapping $X_t : \mathcal{L} \times [\underline{\beta}, \bar{\beta}] \times \{1, \ldots, \bar{z}\} \rightarrow \{-z_t, \ldots, z_t\}^{2N+1}$, where $\mathcal{L}$ is the set of all books. Each agent chooses a strategy to maximize his own payoff, given his beliefs about the execution probabilities, $\mu^e_{t-1}(\cdot)$, and changes in v given execution, $\mu^v_{t-1}(\cdot)$. This latter belief reflects the trader's accounting for "picking off risk" and "adverse selection" due to future traders being better informed about future v.

3.2 Existence

In equilibrium, $\mu^e_t = \mu^e$ and $\mu^v_t = \mu^v$ for each t. That is, any two agents facing the same limit order book have the same beliefs about execution probabilities and changes in v conditional on execution. Further, agents' beliefs must be consistent with the actual future course of play. The equilibrium concept we use is Markov perfect equilibrium. The state at any time t is the limit book, L_t.[7] Since time does not enter into the definition of the state, such an equilibrium must be stationary. Thus, we rule out "time of day effects" or equilibria of the form: "Every 333rd period, submit more aggressive orders." The Markov specification requires agents to condition only on the current book, and not on any prior books. In this model it is not restrictive, because the book at any time summarizes the payoff-relevant history of play. Since each agent takes a decision at only one period of time, perfection is immediate.

It follows from existing work that a Markov perfect equilibrium exists to this game. Suppose the support of the β distribution is bounded both above and below, and δ is strictly positive. Then, the payoffs must be bounded. Also, notice that as $\delta > 0$, the probability that a limit order executes at a particular time, $t + \tau$, goes to zero as $\tau \rightarrow \infty$. This is because the order either executes or expires before $t + \tau$. Thus, this game is continuous at infinity (see Fudenberg and Tirole, page 110). It now follows immediately from Theorem 13.2, page 515, of Fudenberg and Tirole (1991) that a Markov perfect equilibrium exists in this game.[8]

[6]Execution, if it occurs, happens at some time $\tau \geq t$. The amount of time that elapses is sufficiently short, however, that we do not discount future surplus.

[7]We exclude v_t from the state since L_t shifts as needed to keep prices relative to the current consensus value.

[8]We do not prove uniqueness. However, in keeping with existing literature, we observe that the equilibrium we exhibit is computationally unique.

3.3 Solving for Equilibrium

Equilibrium is obtained by finding common beliefs, μ^e and μ^v, such that when each trader plays his best response, the means of the distributions of realized executions and changes in v conditional on execution indeed match the expected values for these outcomes, as specified by μ^e and μ^v.

To find this fixed point, we simulate a market session and update beliefs given the simulated outcomes until beliefs converge. We follow Pakes and McGuire (2001), in using a stochastic algorithm to asynchronously update these beliefs. The advantage of this approach is two-fold. Consider the trader's belief for the execution probability of a limit buy for one share at price p^i given the current book, L_t. To update this belief non-stochastically one would integrate over all the possible sequences of future outcomes (of new trader arrivals, order cancellations, and v jumps) that lead to this share either being cancelled or executed. Instead, we simply track whether this share ultimately executes or is cancelled in the market simulation. Upon execution or cancellation, we update the current value of μ^e for the state at which this share was submitted. Updating μ^v is similar: we keep track of the net changes in v since the order was placed. If the share executes, we average in the net changes to μ^v. In essence, this approach uses a single draw to perform Monte Carlo evaluation of a complicated integral.

The second advantage of the stochastic algorithm is that beliefs are only updated for states actually visited. By state we mean a combination of the current book and a share placed at a particular price. The fixed point is computed only for the *recurrent class* of states. As discussed in section 1, the full state space for this game is too large for traditional numerical methods that operate over the entire state space.[9] A natural concern is that false beliefs at points outside the recurrent class may lead players to mistakenly avoid such states. To alleviate this concern, we specify initial beliefs to be overly optimistic: states not in the recurrent class would not be visited even if beliefs for them were correct.

As discussed in Pakes and McGuire (2001), this algorithm may be viewed as a behavioral description of players learning about the game in "real-time" using the same updating rules as the algorithm. Here, we only use the algorithm to characterize beliefs and behavior in equilibrium. We do not use the model to infer how players "arrive" at the equilibrium.

In more detail, the algorithm works as follows. First, we choose a rule, $\{\mu_0^e(\cdot), \mu_0^v(\cdot)\}$, for assigning beliefs to states encountered for the first time. As discussed, this rule must be optimistic. The simplest such rule is $\mu_0^e(\cdot) = 1$ and $\mu_0^v(\cdot) = 0$. A better rule sets $\mu_0^e(\cdot, i, \cdot, \cdot)$ to the probability that a trader, for whom taking the other side of the transaction at price p^i

[9] To be of practical use, an algorithm must only operate on a set of states that can be stored in the computer's memory (without swapping to "virtual" memory on the hard drive).

would yield non-negative surplus, arrives before the order is randomly cancelled. To derive this probability, for a limit buy at p^i, note that the probability of surviving τ periods with no such sellers arriving is $(1-\delta)^\tau(1-F_\beta(p^i))^\tau$. Execution in the $\tau+1$ period, therefore, occurs with probability no more than $(1-\delta)^\tau(1-F_\beta(p^i))^\tau(1-\delta)F_\beta(p^i)$. Since execution can occur in any future period,

$$\mu_0^e(\cdot,i,\cdot,\cdot) = \sum_{\tau=0}^{\infty}\left[(1-\delta)^\tau(1-F_\beta(p^i))^\tau(1-\delta)F_\beta(p^i)\right] = \frac{(1-\delta)F_\beta(p^i)}{1-(1-\delta)(1-F_\beta(p^i))}.$$

The initial belief rule μ_0^e is similarly derived for limit sells.

Next, we choose an arbitrary initial book, L_0. For simplicity we choose L_0 to be empty.[10] We then set $t=1$ and iterate over the following steps.

Step 1: Draw the period t trader's (z_t,β_t) and determine the optimal action X_t, given μ_{t-1}^e,μ_{t-1}^v.

Step 2: For each market order share, update $\mu_t^e(\cdot)$ and $\mu_t^v(\cdot)$ for the initial state of the share crossed-off by the market order. If the crossed-off share was the k^{th} share submitted in period $\tau<t$ at price p^i (relative to v_τ), then

$$\mu_t^e(k,i,L_\tau,X_\tau) = \frac{n}{n+1}\,\mu_{t-1}^e(k,i,L_\tau,X_\tau)+\frac{1}{n+1} \quad (5)$$

$$\mu_t^v(k,i,L_\tau,X_\tau) = \frac{n}{n+1}\,\mu_{t-1}^v(k,i,L_\tau,X_\tau)+\frac{v_t-v_\tau}{n+1} \quad (6)$$

where n is the number of shares submitted at state (k,i,L_τ,X_τ) that have either executed or been cancelled between periods 0 and t.

Step 3: Add each limit order share in X_t to the end of the appropriate queue in L_t. At this point the book is L_t+X_t.

Step 4: Cancel each share in the book with probability δ. Update μ_t^e for the states at which cancelled shares were submitted. The update uses only the first term in equation (5) since the last term has a numerator of zero for cancelled shares (and one for executed shares). Note that μ_t^v is not updated since it corresponds to changes in v conditional on execution.

Step 5: Determine the consensus value for the next period using the following transition kernel:

$$v_{t+1} = \begin{cases} v_t+1 & \text{with probability } \lambda/2 \\ v_t & \text{with probability } 1-\lambda \\ v_t-1 & \text{with probability } \lambda/2. \end{cases}$$

[10] While we have not proven uniqueness of equilibrium, the algorithm converges to the same equilibrium from various initial beliefs and books.

If v changes then shift the book to maintain the normalization of $p^0 = v$ (i.e., to maintain prices being relative to the current consensus value), as discussed in section 2. Consider an increase in v: sell orders at (pre-shift) tick $-(N-1)$ are picked-off by the crowd of buyers at tick $-N$ and buy orders that were at tick $-(N-1)$ are cancelled. The states at which these orders were submitted have beliefs updated in the appropriate manner: executed orders use the update rule in step 2, while cancellations use the update rule in step 4. When v decreases, orders that were at tick $N-1$ are processed similarly.

Step 6: Implicitly set $\mu_t^e = \mu_{t-1}^e$ and $\mu_t^v = \mu_{t-1}^v$ for states not updated in steps 2, 4, or 5. Set $t = t+1$, and return to step 1.

Pakes and McGuire (2001) stop the stochastic updates every million iterations to check for convergence. Their convergence test compares the most recent values with values based on exact evaluation of the integral defining each state's value (instead of the monte carlo approximation of the integral used in the simulations). They stop the algorithm when the correlation in these two measures is 0.995 and the difference between their weighted (by frequency of visiting each state) means is less than 1%.

Performing the same convergence test for our model is computationally not feasible. Our model does not yield a closed form for the integral defining each state's value (equivalently μ^e and μ^v).[11] While one could numerically integrate to a high degree of accuracy for comparison to the single draw implicit in the algorithm, the amount of time needed to do so for each of the 10–20 million states in our model's recurrent class is prohibitive.

Our convergence test, instead, directly compares μ^e and μ^v, respectively, with realized execution frequencies and realized changes in v conditional upon execution. Every 100 million iterations we compute the weighted (by visitation frequency) average of the absolute difference in the realized outcomes over the previous 100 million traders, and beliefs at the end of the 100 million iterations. When the difference is less than 1% for both μ^e and μ^v we stop the algorithm. We then confirm that these weighted averages remain less than 1% when we simulate an additional 100 million periods holding beliefs fixed.

[11] A closed form is obtained when $z_t = 1$ for all traders, and when beliefs are stored for each position in the book that a given share may ever occupy during its life in the book, as opposed to only the initial position upon submission. This latter stipulation corresponds to a model that may be written in the traditional recursive form in which the value of today's state is the sum of current utility plus a (potentially discounted) continuation value. Given that traders in our model choose an action only upon entering the market, the beliefs assigned to "intermediate" states (in route to potential execution) are not needed for decision purposes. By not tracking such states we reduce the algorithm's memory requirements. This allows us to parameterize the model to yield "thicker" books, and books with more ticks.

As in Pakes and McGuire (2001), we reset n (in step 2) to 1 every 1–10 million periods until beliefs have begun to stabilize. This enables the algorithm to quickly correct for the excessive optimism of initial beliefs at most states.

4 Characterization of Equilibrium: Simulation Results

Once the algorithm has converged, we record 500,000 trader arrivals. We provide a detailed description of the outcomes of our equilibrium and compare our outcomes to existing empirical literature. We also perform some comparative statics on the exogenous asset volatility parameter of our model, the innovation rate, λ. This allows us to demonstrate the properties of the equilibrium and to relate it to existing literature. In what follows, given that we have solved for equilibrium, we provide data on means.[12]

4.1 Numerical parameterization

We do not attempt to calibrate the model. Our purpose here is merely to demonstrate some features of equilibrium. In particular, the relationship between the transaction price an both the consensus value and private value.

The following parameterization corresponds to our benchmark case.

- There are nine ticks and the tick size is normalized to $d = \$\frac{1}{16}$. The ticks are denoted $\{-4, -3, -2, -1, 0, 1, 2, 3, 4\}$. The corresponding price vector relative to the common value is $\{-\frac{1}{4}, -\frac{3}{16}, \cdots, \frac{1}{4}\}$. Traders may submit limit orders at ticks $\{-3, \ldots, 3\}$. At ticks -4 and 4, a trading crowd provides infinite liquidity.

- F_β is a normal distribution with mean 0 and standard deviation 3 ticks, or \$0.1875.[13] Hollifield, Miller, Sandås and Slive (2002) estimate for three stocks on the Vancouver Stock exchange that trades "with a valuation within 2.5% of the average value of the stock account for between 32% and 52% of all traders." This implies a standard deviation of the private value distribution approximately equal to 4.5% of the value of the stock. Given our parameterization, this corresponds to a consensus value of approximately $\frac{3}{16}/.045 = \$4.17$.

  The choice of F_β is not motivated by computational need. The algorithm can handle any distribution for F_β.

[12] Robustness checks performed by breaking the sample into smaller sets indicated that our results are not sensitive to the particular part of the infinitely long equilibrium path that we report.

[13] Although existence of equilibrium depends on payoffs being bounded, for convenience we assume that β has a normal distribution. Our numerical results are unchanged if, for example, we truncate this normal so that all values lie between $-n\,\sigma$ and $n\,\sigma$, where n is an arbitrarily large, but finite, number. This restriction would guarantee that payoffs are bounded.

- F_z assigns $z \in \{1, 2\}$ with equal probability. That is, each trader has, with equal probability, either one or two units to trade. The potential trade size distribution is difficult to parameterize by casual observation, because order size is endogenous. For the maximal trade size, we choose the lowest number (2) that allows traders to submit multiple orders.

- Each period, each share is cancelled with probability $\delta = 0.04$. If a share is not executed, the expected time before it is cancelled is 25 order arrival periods.[14] If orders arrive every 120 seconds, this parametrization suggests that limit orders stay on the book for about 50 minutes. This is in keeping with the stylized facts presented in Lo, MacKinlay and Zhang (2002). In a pooled sample of 100 stocks they find that limit orders failing to execute are cancelled on average after 46.92 minutes for buy orders and 34.15 minutes for sell orders.

- We take the innovation to the consensus value to be small, with $\lambda = 0.04$. If the consensus value changes it does so by one tick. The probabilities of increases and decreases are the same (0.02); we do not incorporate a trend for the common value. If trades occur every 90 seconds, it is difficult to envisage information about the cash flows of the firm or aggregate economy arriving with greater rapidity. Notice that λ is the variance of the innovation distribution.

  We choose the innovation in the consensus value of the asset to be one tick. That is, the asset value can increase or decrease by one tick. As agents' private motives for trade are continuous, traders will arrive with varying incentives to undercut the existing book. Thus, the discrete jump coupled with the continuous β distribution captures agents' incentive to undercut.

- For convenience, we report all payoffs in ticks (where one tick represents $\$\frac{1}{16}$).

4.2 The State and Evolution of the Limit Order Book

Since we simulate a symmetric version f the model, we focus primarily on buy orders. 50.091% of the orders were buy orders. Of these, 20.872 % of the orders were market buy orders, and 29.219% were limit buy orders.

The mean of the β distribution, and hence the consensus value of the asset is always zero. Most orders are submitted at this value. In figure 2, we depict the total number of all types of shares submitted at each tick. For convenience, in this figure, we show only orders

[14] With probability $\delta^k(1-\delta)^{k-1}$, the share will last k periods. Hence, the expected time until cancellation (absent market order executions) is $\delta\{1 + 2(1-\delta) + 3(1-\delta)^2 + 4(1-\delta)^3 + \ldots\} = \frac{1}{\delta}$.

submitted by traders, and ignore the infinite liquidity supplied by the trading crowd at ticks -4 and 4. Over the 500,000 trader arrivals, only 18 trades were automatically crossed against the liquidity supplied by the crowd; 11 buys at a tick of $+4$ and 7 sells at -4. A miniscule proportion of traders (3 out of 500,000) chose to not submit an order.

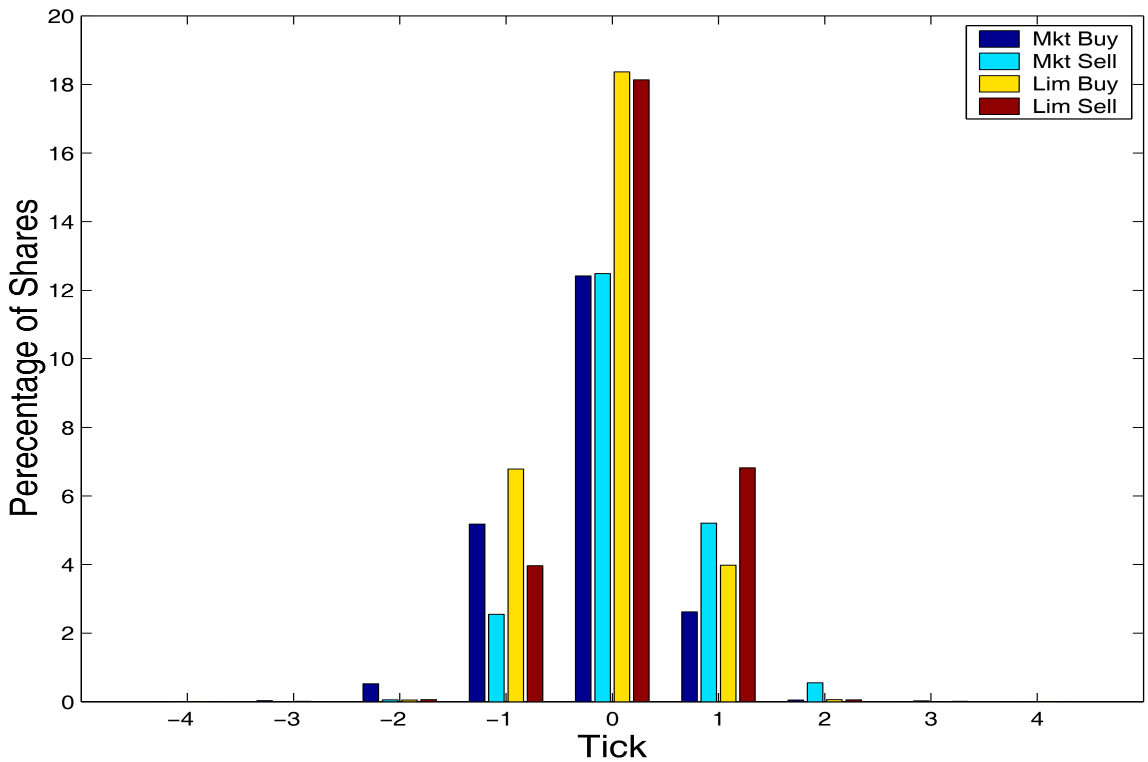

Figure 2: **Total number of shares at different ticks**

While the numbers are large, on average, the book is thin: The open electronic limit order book is effective at consummating trades. We present the average buy and sell sides of the book in figure 3. The average book has a total of 2.86 shares on the buy side, and 2.84 on the sell side. As one might expect, depth is concentrated near the consensus value of the asset, with buy orders more likely to be placed above it and sell orders more likely to be below it. Further, as expected given the symmetry in the model, the book is symmetric. The number of limit order buys one tick below the consensus value of the asset is equal to the number of limit sells one tick above it.

The evolution of the book is determined by a trader's order submission strategies. Following Biais, Hillion and Spatt (1995) and the subsequent literature,[15] we classify the types

[15]See,for example, Ahn, Bae and Chan (2001), Griffiths, Smith, Turnball and White (2000) and Ranaldo (2003).

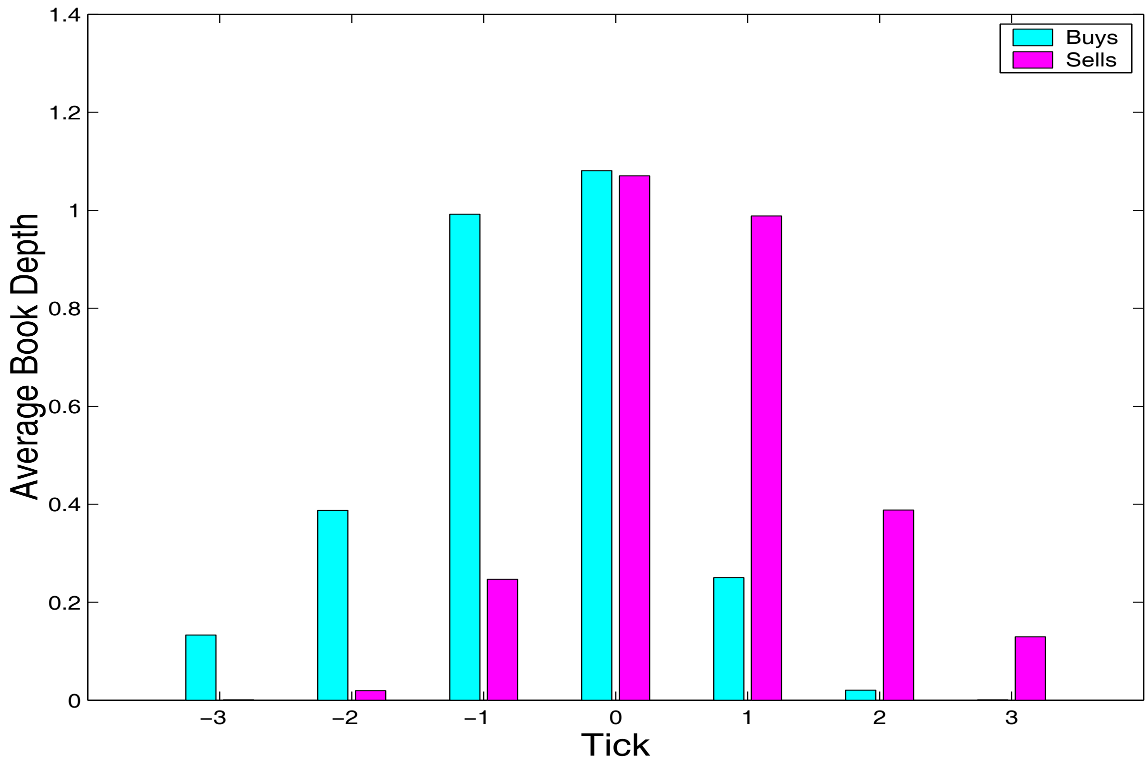

Figure 3: **Average book depth at different ticks**

of orders we observe. The definitions we use (for buy orders) are presented below. The definitions for sell orders are exactly analogous.

Large Buy (LB)	Market buy order that moves the price. This occurs with a market buy for 2 shares, where the second share is bought at a higher price than the first.
Small Buy (SB)	Market buy order where all shares are bought at the same price. Thus, either only 1 share is submitted or 2 if they both execute at the same price.
Aggressive Buy (AB)	Limit buy order at a price that is higher than the current bid. This can be either 1 or 2 shares.
At the Quote Buy (QB)	Limit buy order at the current bid. This can be either for 1 share or 2 shares.
Below the Quote Buy ($< QB$)	Limit buy order at a price lower than the current bid. This is for either 1 or 2 shares.

Biais, Hillion and Spatt (1995) report on the persistence of orders on the Paris bourse. They find systematic patterns of trade that are consistent with information effects. For example, a large buy (at the ask) followed by a limit buy above the previous bid "reflects the adjustment in the market expectation to the information content of the trade." Further, they identify a "diagonal effect": The conditional probability of an order following a similar order is typically higher than the unconditional probability of that kind of order.

In our model, such patterns might emerge for two reasons: either because of a particular pattern in the book or because of information events. First, a trader conditions his actions on the state, as this affects the probability of execution. Thus, if the same action is optimal in similar states of the book, one might expect sequential traders to submit similar orders. Second, the impact of a change in the consensus value, v_t, may induce subsequent traders to take similar actions, until the book has adjusted. For example, we expect to see buy orders following an increase in v_t.

To investigate these sources of persistence, we report the probability of observing an action at time $t + 1$, conditional on the action at time t. In our data, a trader with two shares submits them simultaneously. If he submits two different kinds of orders, to determine the transition probabilities, we need to assign one order to be the "first." We use the following rule: if one share is submitted as a market order, it is the first share. In all other cases, we randomly assign one share to be the first.

Table 1 reports the transition probabilities for buy orders (again reported as percentages). The model is symmetric, so the conditional probabilities of sell orders are similar. The sum across each row is the conditional probability of a buy order at $t + 1$ given the conditioning event in the first column. That is, across each row, the probabilities sum to the conditional probability of observing a buy order at $(t + 1)$, given the event at t.

The Table reveals some persistence in order submission. For many of the defined events, we do find a diagonal effect. Thus, market buys are more likely after market buys than market sells.[16] Further, market orders are frequently followed by aggressive limit orders on the same side of the market. This sometimes happens because traders exhaust the liquidity in the book and then become liquidity providers at the same price.

To determine if information events are the cause of the autocorrelation in the data, we report in Table 2 the transition probabilities for a model in which there are no jumps in the consensus value (thus, there is no informationally motivated trade).

With no information, there are fewer large orders in our data. That is, the existence of large orders per se indicates the importance of information. When the consensus value can

[16] This also accords with the theoretical results of Parlour (1998) and Foucault (1999), and the empirical findings of Hollifield, Miller, Sandås and Slive (2002) and Ranaldo (2002).

	Market Buys			Limit Buys				Total
	Large	Small	Total	Agg.	At Quote	< Quote	Total	Buys
LB	0.400	33.143	33.543	21.439	0.171	0.000	21.610	55.153
SB	1.124	23.442	24.566	29.591	6.555	0.568	36.715	61.280
LS	0.000	0.056	0.056	18.025	20.028	7.900	45.953	46.008
SS	0.005	6.363	6.368	12.892	13.551	6.170	32.612	38.980
AB	0.026	12.188	12.215	4.839	23.454	7.454	35.747	47.962
AS	1.793	35.576	37.369	6.000	6.812	1.997	14.809	52.178
QB	0.010	13.828	13.839	7.177	17.392	8.389	32.958	46.796
QS	0.048	35.980	36.029	7.878	7.781	1.892	17.550	53.579
<QB	0.000	3.906	3.906	2.179	24.509	14.881	41.570	45.476
>QS	3.640	41.118	44.758	4.346	4.777	0.180	9.303	54.061
Overall	0.626	20.246	20.872	12.331	12.357	4.531	29.219	50.091

Table 1: **Conditional frequencies of buy orders at time $t+1$ (column) given the order at time t (row), base case.**

	Market Buys			Limit Buys				Total
	Large	Small	Total	Agg.	At Quote	< Quote	Total	Buys
LB	0.000	34.330	34.330	20.926	0.000	0.000	20.926	55.256
SB	0.187	24.867	25.054	23.430	9.766	0.078	33.274	58.328
LS	0.000	0.000	0.000	9.354	32.883	0.868	43.105	43.105
SS	0.000	7.533	7.533	10.342	21.529	2.050	33.921	41.454
AB	0.000	10.779	10.779	1.906	30.303	5.536	37.745	48.524
AS	1.166	35.559	36.724	3.600	10.484	0.500	14.585	51.309
QB	0.000	15.587	15.587	3.199	25.829	1.535	30.562	46.149
QS	0.013	33.111	33.124	5.451	14.043	1.275	20.770	53.894
<QB	0.000	2.260	2.260	0.600	26.853	11.058	38.511	40.771
>QS	2.550	42.571	45.121	2.033	11.888	0.000	13.921	59.042
Overall	0.192	20.701	20.893	9.172	18.199	1.695	29.065	49.958

Table 2: **Conditional frequencies of buy orders at time $t+1$ (column) given the order at time t (row) with no changes in consensus value**

change, traders are more willing to pay for immediate execution, to avoid picking off risk. Also, limit orders are more dispersed: there are more orders at prices both worse than and better than the current quotes. We comment further on this in Section 5.1 below.

In terms of the conditional probabilities of events, we find greater persistence in small market orders and limit orders at the quote when there is no information. Also, following a large buy, an aggressive buy order is more likely, than in the case with jumps. Therefore, large conditional probabilities for these events do not indicate information effects. However, the persistence of large orders and limit orders away from the current quote does.

4.3 The State and Evolution of the Bid Ask Spread

In a pure limit order market, the quotes are set by limit order traders and changed when either a limit order traders undercuts (competes on price) or a market order trader executes against one. Thus, the incentives of traders to make and take liquidity determine the spread. In our model, in keeping with standard intuition, a market order is less likely when the spread is wide, while limit orders are more likely. In table 3, we report the unconditional probability (reported as a percentage) of observing market or limit orders as a function of the spread. Due to symmetry, we only report buy orders. The results accord with those presented in Foucault, Kadan and Kandel (2002) and Foucault (1999). When spreads are wide, market orders are more "expensive," and thus traders tend to submit aggressive limit orders which narrow the spread. When spreads are narrow, traders tend to "take liquidity," or submit market orders. When the spread is very wide (5–8 ticks), virtually all buy orders submitted are aggressive limit orders.

Spread	Market Buys			Limit Buys				Total
	Large	Small	Total	Agg.	At Quote	$<$ Quote	Total	Buys
1	0.465	28.211	28.677	0.000	16.470	4.998	21.468	50.145
2	0.697	11.599	12.296	25.544	8.251	3.853	37.647	49.944
3	2.321	22.361	24.683	6.959	8.881	9.671	25.511	50.194
4	0.017	15.306	15.322	18.911	13.422	2.264	34.596	49.919
5	0.000	0.647	0.647	48.464	1.119	0.000	49.583	50.229
6	0.000	0.042	0.042	50.681	0.000	0.000	50.681	50.723
7	0.000	0.000	0.000	49.562	0.000	0.000	49.562	49.562
8	0.000	0.000	0.000	50.520	0.000	0.000	50.520	50.520

Table 3: **Frequency (%) of buy orders for different spread sizes.**

The rapid response of traders to profit opportunities ensures that the spreads are narrow. Figure 4 displays the probability density of the quoted spread. Almost half the time, the quoted spread is 1 tick, with an average of 2.20 ticks.

Glosten (1987) shows that one can decompose a quoted spread into order processing and adverse selection costs. His market maker framework is not directly applicable to a limit order market as in the latter context, quotes are set by (possibly stale) limit orders rather than continuously adjusted by market makers. The quotes in a limit order market are thus backward looking not, as in intermediated trade models, the expected value of the asset conditional on a transaction. This does not imply that information plays no role in determining the quoted spread in a limit order market: Limit orders are placed by traders who rationally anticipate changes in the common value. Thus, in equilibrium, the

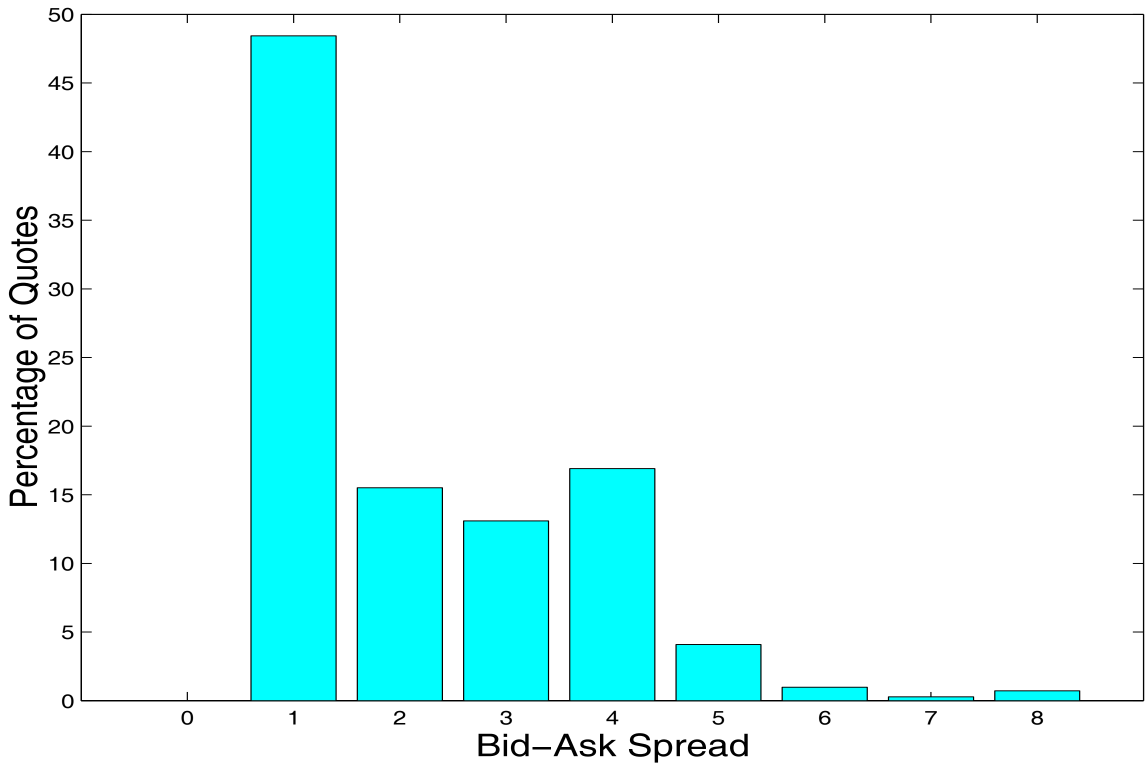

Figure 4: **Frequency distribution of quoted bid-ask spread** $\lambda = 0.04$

spread distribution is affected by information. Figure 5 demonstrates that in the absence of asymmetric information quotes are on average narrower. The mean drops to 1.87 ticks. Thus, volatility of 4% leads to an increase of 17.7% in the average quoted spread.

5 Disentangling the private value and consensus value

All our results derive from the fact that order flow is endogenous. In particular, arriving traders take advantage of profit opportunities on the book. This is the other side of the winner's curse or picking off risk faced by limit order submitters: Losses to limit order traders accrue as benefits to market order submitters. We first document the winner's curse and the market's response to it and then determine the implications for the transaction price and common proxies for the consensus value. For example, we determine if the midpoint of the bid-ask spread is a good proxy for the consensus value.

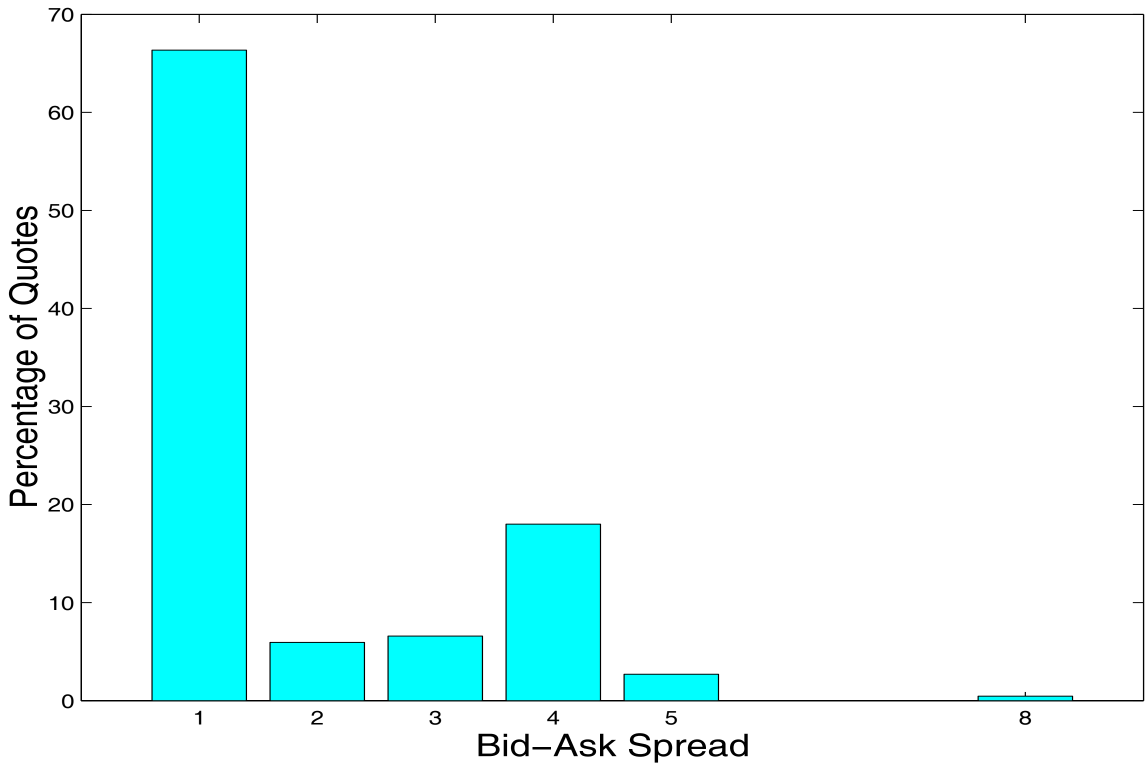

Figure 5: **Frequency distribution of quoted bid-ask spread with no asymmetric information**

5.1 The Winner's Curse and Picking off Risk

Limit orders are more likely to be executed after the consensus value moves against the limit order submitter. That is, a limit buy order is more likely to be executed if the asset value moves down and a sell order is more likely to be executed if the consensus value of the asset moves up. Of course, in equilibrium, agents placing limit orders compensate for this.[17] As we have mentioned, an immediate implication is that the market orders who execute against the 'picked off' limit orders have made advantageous trades. How large are these effects?

In our data, we have 40,154 changes in the consensus value, of which 19,970 were increments and 20,184 were decrements. If there is a change in the consensus value, limit orders may either be executed by incoming market orders (who now possess an informational advantage about the consensus value) or by the trading crowd if they are not picked off by incoming market orders. In both cases, they have a higher probability of execution in states

[17] Nyborg, Rydqvist and Sundaresan find evidence of bidders' compensating for the winner's curse in Swedish Treasury auctions.

in which they receive a lower payoff. This, of course, is the realization of the winner's curse. In figure 6 we illustrate the average payoff to limit orders which executed after a jump.

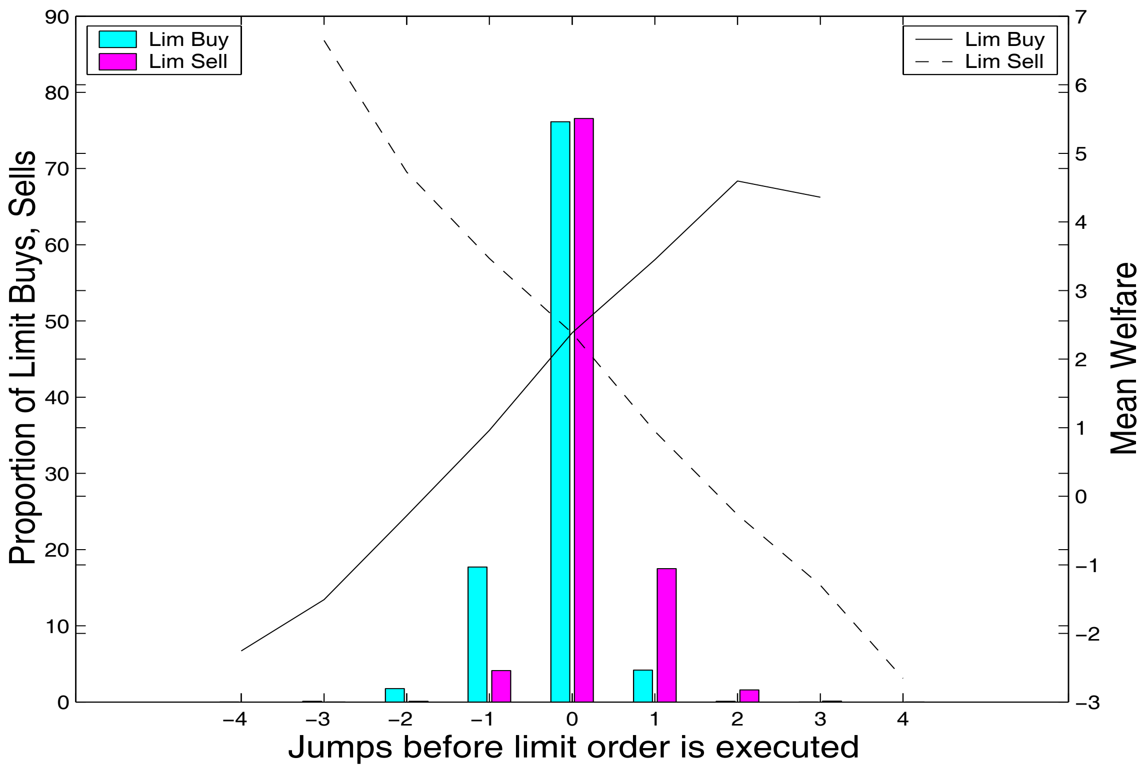

Figure 6: **Number of jumps before a limit order is executed, and trader welfare**

The horizonal axis records the number of jumps before an order is executed. The height bars represent the proportion of executed limit buy or sell orders, as shown on the left hand vertical axis. On the right hand vertical axis the mean welfare of limit orders conditional on execution is recorded. The relationship between this and the number of jumps is depicted by the dashed line for limit sell orders and the solid line for limit buy orders.

Limit order submitters optimally place their orders so that, on average, their payoff is still positive after two jumps down for a buy order and two jumps up for a sell order. Nonetheless, they are subject to picking off risk: conditional on a limit buy (sell) executing, the expected change in the consensus value is -0.17 ($+0.17$) ticks. In total, 19.4% of all limit orders experience an adverse change in the consensus value before execution. However, consistent with ex ante optimization, the number of limit order traders who suffer from the "winner's curse" is small. On only about 4.06% of all limit orders, the submitter makes a loss (relative to his private value). Interestingly, this loss is counterbalanced by some traders (4.2% of all limit orders) who execute after a favorable change in the consensus value.

To determine the equilibrium effect of trader's compensating for the winner's curse, we compare equilibrium behavior in a market with jumps to behavior in a market with a constant consensus value. We document two differences in Table 4. First, limit orders are submitted at more dispersed prices when picking off risk exists. This is because there is a higher probability of orders executing at prices further away from the consensus value. So, each β type submits buy orders at lower prices and thus $\beta - p^i$ increases. Second, the average spread widens as the volatility in the consensus value increases, confirming the prediction of Foucault (1999).

	λ	
Mean of	0	0.04
Quoted spread, all orders	1.87	2.20
$(\beta - p^i)$, limit buys	2.06	2.17

Note: Both spread and $(\beta - p_i)$ are shown in ticks.

Table 4: **Means of quoted spread for all orders, and private value minus submission price for limit buys**

Table 5 shows that, in the absence of an adverse selection component, the book is much thicker at the consensus value.[18] So, higher consensus value volatility means that orders are submitted further away from the consensus value and orders get cleared out more frequently after asset jumps. These two effects reinforce each other and lead to a wider spread.

	Buy Order Depth		Limit Buy Submissions	
Tick	$\lambda = 0$	$\lambda = 0.04$	$\lambda = 0$	$\lambda = 0.04$
-3	0.00	0.13	0.00	2.29
-2	0.00	0.39	0.00	1.79
-1	1.41	0.99	23.84	19.47
0	1.65	1.08	70.15	58.93
1	0.06	0.25	5.99	17.02
2	0.00	0.02	0.02	0.49
3	0.00	0.00	0.00	0.01
Total	3.12	2.86		

Table 5: **Buy orders in an average book (left) and frequency of ticks at which limit buys are submitted (right)**

[18]The effect of volatility shocks on depth has been examined by Ahn, Bae and Chan (2001) and Coppejans, Domowitz and Madhavan (2001) and Hasbrouck and Saar(2002).

5.2 The Winner's Curse and Market Orders

As we have observed, an immediate implication of the winner's curse faced by limit order submitters is that the beneficiaries (market order submitters) are getting "good deals." In particular, when a trader arrives at the market he is more likely to take liquidity when it is cheap (or offered at a subsidy) in the case of limit orders being "picked off."

To confirm that limit orders are more likely to execute against a market order with superior information, we consider the number of buy orders that are submitted after a jump in the asset value. These results are reported in Table 6.

	Market Buys			Limit Buys				Total
	Large	Small	Total	Agg.	At Quote	$<$ Quote	Total	Buys
Down	0.050	6.718	6.768	9.948	12.624	15.710	38.283	45.051
None	0.580	20.085	20.665	12.516	12.658	4.261	29.436	50.101
Up	2.399	38.102	40.501	9.905	4.256	0.255	14.417	54.917
Overall	0.626	20.246	20.872	12.331	12.357	4.531	29.219	50.091

Table 6: **Probability of buy orders conditional on jumps in the asset value**

Clearly, after an increase in the consensus value, there are more market buy orders submitted than after a decrease. This exemplifies picking off risk for limit orders in the book. If the consensus value of the asset shifts up then last period's ask becomes "too low," offering new traders a profitable opportunity. Similarly, after an increase in the asset value, last period's bid is "too low." This leads to fewer limit buys at or below the bid.

On average how much do market order submitters benefit from picking off limit order traders? Such profit opportunities decrease the cost of demanding liquidity. In the simulation, since we know the true value is always the 0-tick, we have a direct measure of the true transactions cost paid by a market order: a market order executing at a price p^i simply pays $|p^i|$. In Table 7, we report statistics on the transaction costs with and with the possibility of jumps in the consensus value.

True Transaction Cost	Mean	Std. Dev.
Jumps ($\lambda = 0.04$)	-0.176	0.685
No Jumps ($\lambda = 0$)	0.070	0.471

Table 7: **True transactions cost and effective spread, base case**

As the table indicates, the true transactions cost when there is "picking off" risk is *negative*. On average, market order submitters execute at prices better than the true value of the asset. The standard deviation in this case is high: profit opportunities are frequently not always available in the book. Hasbrouck (1993) suggests the standard deviation of difference between the efficient price and transaction price as a measure of market quality. In the context of our model this is just the standard deviation of the transaction costs. Even with no jumps (i.e., no "picking off" risk) the transaction cost is close to zero on average and negative for some traders. This result re-emphasizes the endogeneity of order submission.[19] Traders submit market orders when prices are favorable, and limit orders when they are not. This intuition is appropriate in a limit order market (as opposed to an intermediated one) in which it is costless for agents to switch from market to limit orders. If there are substantial costs to flexibility, transaction costs to market orders must be higher. However, it does suggest that caution be exercised in calculating transaction costs in limit order markets.

5.3 Inferences about the Consensus Value

Given the endogeneity of order flow, is the midpoint of the Bid Ask spread a reasonable proxy for the consensus or common value of the asset in a pure limit order market? To investigate this, we first examine the difference between this midpoint and the consensus value $(m_t - v_t)$, across all quotes, over the 500,000 simulated periods. Figure 7 displays the frequency distribution of this difference.

The mean of this measure $(m_t - v_t)$ is 0.003, with a standard deviation of 1.232 ticks. Thus, the midpoint is an unbiased estimator of the consensus value. However, it is frequently incorrect, as shown in the figure. This is because it may be optimal for a trader to submit a limit sell order below the consensus value of the asset if the current bid is "too low." After an innovation in the consensus value of the asset, the midpoint of the contemporaneous bid-ask spread is typically different from the consensus value.

Frequently, the midpoint is used conditional on a transaction. For example, in empirical measures of transaction costs such as the effective spread, the midpoint is used as a proxy for the common value only if a transaction occurs. Thus, we next examine the difference between the midpoint and consensus value conditional on a market order being submitted. Since the effective spread is only measured for market orders, this yields a more direct sense of the validity of the condition $m_t = v_t$. Figure 8 plots the distribution of $(m_t - v_t)$ conditional on a market buy or a market sell in that period.

It is clear from the figure that market buy orders are more likely when the midpoint is below the true value of the asset (representing a profitable buy opportunity), and sell orders more likely when $m_t > v_t$. Conditional on observing a market buy, the true value of the asset is likely to be 1.09 ticks higher than the midpoint, and, conditional on observing a market sell, it is 1.10 ticks below the midpoint.

[19] Hollifield, Miller, Sandås and Slive (2002) estimate changes in the price of immediacy.

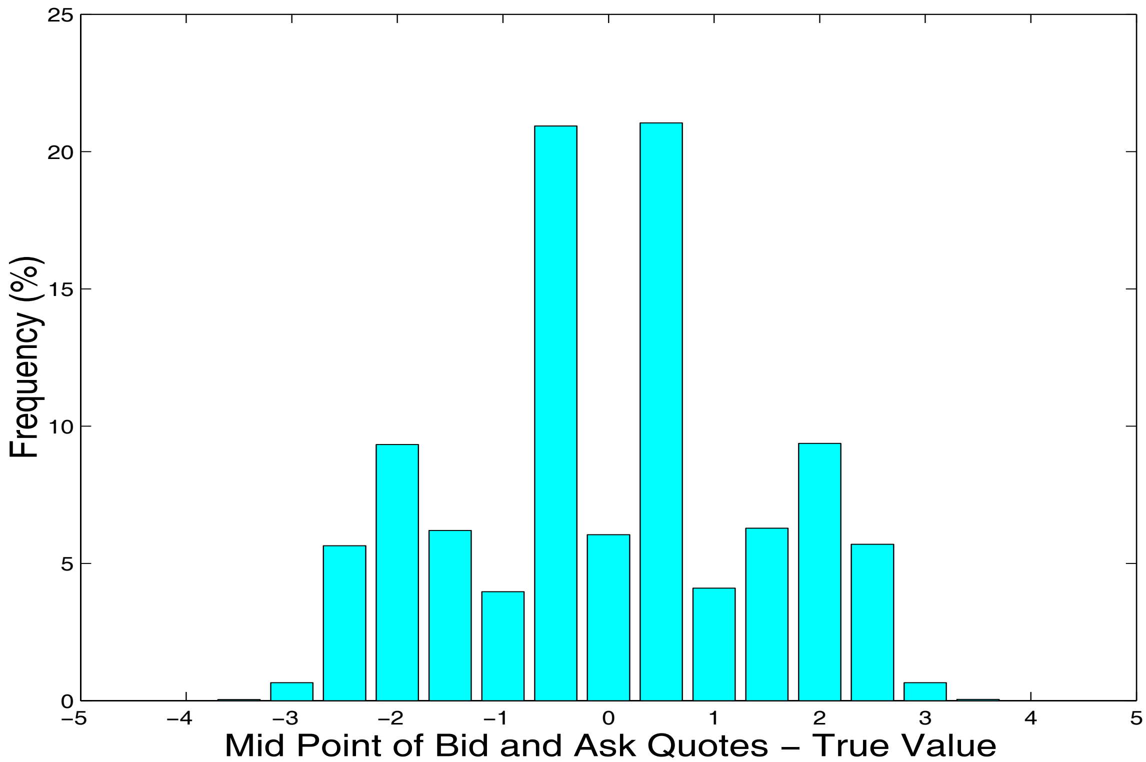

Figure 7: **Frequency chart of midpoint minus true value ($m_t - v_t$)**

To examine the robustness of this result, we checked the corresponding figures for the case when there is no change in consensus value (i.e., $\lambda = 0$). In this case, conditional on a market buy (sell) the true value of the asset is 0.72 ticks higher (lower) than the midpoint of the bid-ask spread. Thus, this result is not solely due to stale limit orders but also because of the endogeneity of orders: Buy orders are more likely when prices are low, and sell orders when prices are high. Thus, we conclude that conditional on a transaction, the midpoint is not a good proxy for the consensus value. Alternatively, a better proxy for the consensus value can be obtained by conditioning on the transaction.

5.4 Inferences about Welfare and Transaction Costs

One empirical measure of the true transaction costs paid by market order submitters is the effective spread, S_t. Given that conditional on a transaction, the midpoint is not a good proxy for the spread, we consider the correlation between true transaction costs and effective spread: -0.23. Thus, when transaction costs are high, the effective spread is low.

A high effective spread is often associated with an imbalance on one side (either the buy

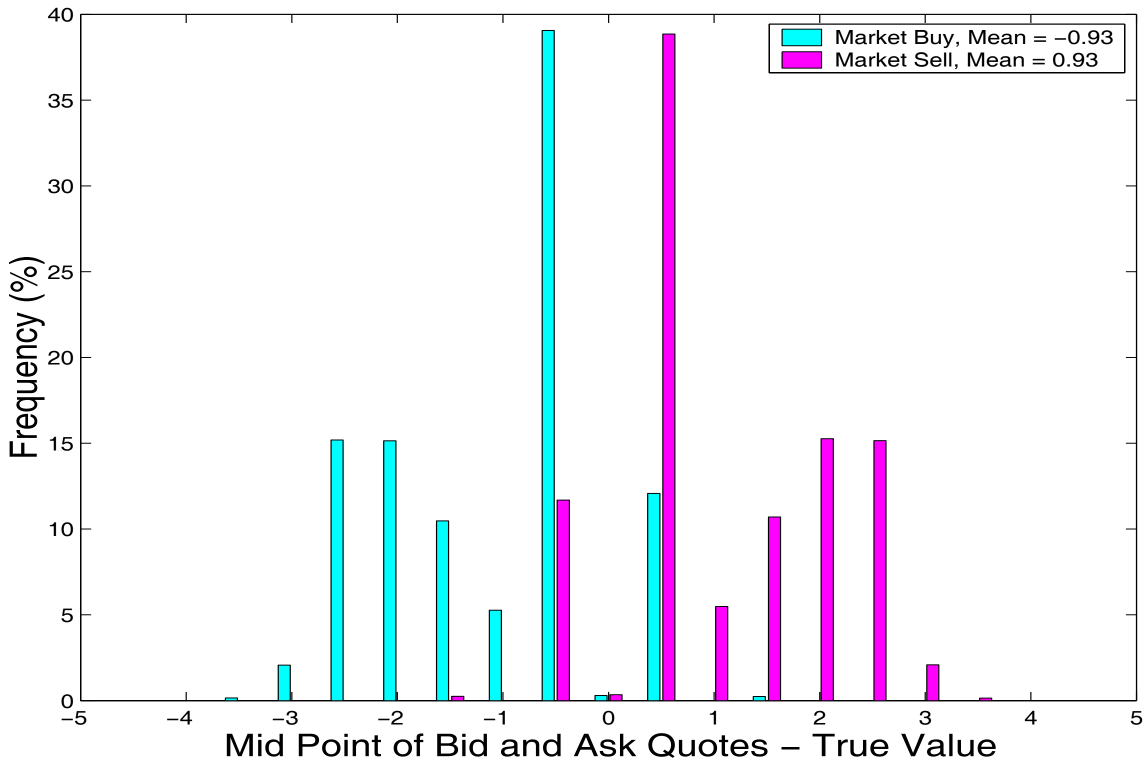

Figure 8: **Frequency chart of midpoint − true value** $(m_t - v_t)$

or sell side) of the limit order book, which skews the midpoint of the bid and ask quotes. As such, a profitable trading opportunity exists on one side of the market, and traders take advantage of this.

Frequently effective spread is used as a proxy for market order submitter welfare. It would be a perfect proxy for the welfare of market order submitters if the correlation between these two were -1. However, since the midpoint is not a good proxy for the consensus value of the asset, we know from proposition 1 that the effective spread will not be perfectly correlated with welfare. We next quantify how well effective spread performs as a proxy for welfare.

To do so, we break our sample into "trading days." Specifically, we have 1000 days (approximately 4 business years) each with 500 trader arrivals.[20] For each day, we calculate the average per-share effective spread, the total effective spread (the total transactions cost paid by market order submitters that day), and the total welfare garnered by market and limit orders. We then examine the day-to-day correlation of these measures, and report this

[20] Alternatively, one could view each subsample as a different stock, under the null hypothesis that trade in each stock is independent and identical. We maintain the interpretation of "days."

in Table 8 below.

	Welfare		
	Market Order	Limit Order	Total
Eff Spread	0.332	0.404	0.444
Avg Eff Sp	0.210	0.310	0.312
Volume	0.526	0.455	0.597

Table 8: **Day-to-day correlations of effective spread and welfare**

Here, the correlation between effective spread and welfare (either market order or total) is actually positive. Ceteris paribus, the more desperate a trader is to trade, the more likely he is to be willing to execute at a worse price. If the endogenous book is such that desperate traders often face bad prices, then effective spreads will be high and so will welfare. That is, the higher the effective spread, the higher the welfare of the market order submitters—a problematic finding for the use of effective spread as a welfare measure.

Notice the high correlation between volume and welfare. A higher trade volume must be correlated with higher welfare, since all trades are individually rational. Indeed, as a rule of thumb, volume appears to be a good proxy for welfare.

6 Evaluating Policy Changes Across Different Regimes

Given that effective spread has sometimes been used to evaluate policy, we perform two such experiments and compute welfare. First, to determine if changes in effective spread are a good proxy for changes in welfare across different regimes and secondly to evaluate directly the policy experiments. We consider two such policy experiments: (i) changing the tick size, and (ii) changing the standard deviation of the β distribution (i.e., the gains from trade).[21]

6.1 Change in the Tick Size

First, we evaluate the effect of changing the tick size. Besides providing an evaluation of the effective spread, a tick size experiment has policy and market design implications. Both the theoretical and empirical literature are mixed on the effects of a tick size change on welfare. Seppi (1997) suggests that small traders are better off under a small tick size, while large traders are at a disadvantage. Cordella and Foucault (1999), in examining competing market makers find that transaction costs are minimized at an optimal, non-zero tick size.

[21] Such a change could come about due to (for example) changes in broker fees. It is possible that a decrease in broker fees could increase the percentage of the population who trade and who have valuations of the asset close to zero.

Empirical evidence on reduction in the tick size both in general and on the NYSE is mixed. There have been a few natural experiments: Toronto moved to decimals in 1996 while Nasdaq and the NYSE have both changed their tick size. The effects have been examined by, among others, Bacidore (1997), Porter and Weaver (1997) and Ahn et al. (1998). Bessembinder (1997), Bonen and Whaley (1998), Ronen and Weaver (1998), and Jones and Lipson (2001) examine the effect on the transaction costs incurred by different parties after the move to "teenies." Goldstein and Kavacejc (2002) and Edwards and Harris (2001) explicitly examine the effect of halving the tick size on liquidity suppliers—the limit order book in the first case and the specialist's ability to "step ahead" in the second. The literature typically finds that a decrease in tick size reduces both effective spread and depth in the book. As evidenced by these papers, the aggregate effect of such changes has defied quantification. Has the reduction in tick size been a Pareto improvement?

To answer this question, we compare two regimes, one with 9 ticks and one with 5. For computational ease in performing this comparative static, we make a slight modification to the base case. We do this so that the dollar magnitude of jumps in the consensus value and the potential gains from trade are the same across the two regimes. Recall, the β distribution is denominated in ticks. In both cases, $\delta = 0.02$, and $\lambda = 0.04$. However, in one case we consider 9 prices in which each jump is two ticks, compared to 5 prices in which the jump is 1 tick. Thus, the dollar magnitude of the jumps is the same. However, notice that the magnitude of the jumps is twice that of the base case model.

The mean and the standard deviation of the β distribution is adjusted so that the same percentage of traders in both cases have valuations more extreme than the trading crowd. In particular, with 9 ticks, we use a mean of 0 and a standard deviation of 3 ticks. In the 5 tick case, we have a mean of 0 and a standard deviation of 1.5 ticks. This ensures that, in both cases, the standard deviation of β is $\frac{3}{16}^{th}$ of a dollar. We illustrate the β distributions in Figure 9.

We report the results of this experiment in table 9. For ease of comparison, all values are reported relative to the tick-size in the 9-tick model. We report the means of welfare and effective spread per available share and per executed share. For welfare, the mean per available share is the most relevant measure. We define the total number of "available" shares to be the sum over all traders of the maximal quantity an agent may trade; that is, $\sum_{t=1}^{500{,}000} z_t$. If a policy change results in fewer trades, the mean welfare per available share will fall, while the mean per executed share need not. For policy prescriptions we should care about foregone trades. For effective spread, the mean per executed share appears to be the most relevant measure, given its prominence in empirical work.

Using the 9-tick regime as a base case, effective spread per executed share rises by 22.0%

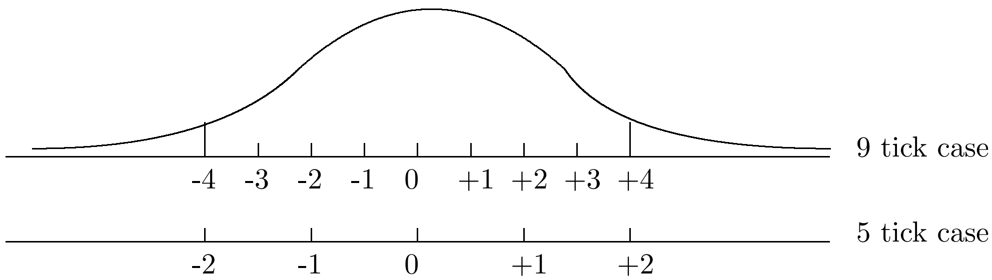

Figure 9: Relationship of ticks to β distribution

	5 ticks Mean Per Share		9 ticks Mean Per Share	
	Available	Executed	Available	Executed
Volume	0.408	1.000	0.419	1.000
Mkt Ord Welfare	1.118	2.736	1.218	2.906
Lim Ord Welfare	0.890	2.178	0.845	2.017
Total Welfare	2.006	4.916	2.063	4.923
Eff Spread	0.526	1.290	0.443	1.057

Table 9: **Results of change in tick size**

when the tick size is doubled. However, the two regimes have roughly the same volume, and hence welfare. Total welfare per available share falls by 2.8%, and the welfare of market order submitters falls by 8.9%. In other words, a large change in the effective spread can occur despite a relatively small change in welfare. Again, the change in volume is a good proxy for the change in welfare; volume per available share changes by 2.7%.

These results allow us to reconcile the empirical literature with the theoretical literature. Most of the empirical literature has found that a reduction in tick size leads to a reduction in spreads, and the inference has been drawn (albeit in intermediated markets) that, ceteris paribus, traders are better off. The theoretical literature has suggested that decreases in tick size are not always Pareto improving. Our results suggest that a decrease in effective spread does improve the welfare of market order submitters, but at the expense of limit order submitters. The change in aggregate welfare is negligible. We interpret our result in the light of order endogeneity. In a pure limit order market, the effect of a tick size change must be of second order. If supplying liquidity becomes too expensive, then agents demand liquidity and vice versa. A change that stopped trades from being consummated would affect welfare. Amending the tick size merely perturbs how the gains from trade are split.

Any decrease in welfare comes about from limit orders that are cancelled unexecuted.

6.2 Change in the Gains to Trade

Even though the effective spread fails, if the bias is systematic, we can still use it to infer welfare. To see if this is the case, we perform another experiment in which we change the gains to trade for agents. We consider a market in which the standard deviation of the β distribution is smaller—2 ticks instead of 3. Effectively, this implies reducing the gains to trade. Gains to trade are larger when traders have more dispersed private valuations. All other parameter values are the same as in the base case.

	$\sigma_\beta = 2$ Mean Per Share		$\sigma_\beta = 3$ Mean Per Share	
	Available	Executed	Available	Executed
Volume	0.405	1.000	0.417	1.000
Mkt Ord Welfare	0.776	1.917	1.234	2.960
Lim Ord Welfare	0.593	1.465	0.887	2.128
Total Welfare	1.370	3.382	2.122	5.088
Eff Spread	0.332	0.819	0.379	0.908

Table 10: **Comparison of two β distributions**

As one might expect, if the gains to trade are larger, the welfare from consummated trade is higher. Indeed, there is a 54.9% increase in total welfare per available share, in moving from $\sigma_\beta = 2$ to $\sigma_\beta = 3$ (consistent with the notion that σ represents the gains to trade). However, the effective spread actually increases when the gains to trade increase. Per executed share, the effective spread increases by 10.9%. In this case, as may be expected, the change volume is a poor proxy for the change in welfare (volume per available share increases by 3.0%). A change in the gains to trade leads to an increase in welfare on every trade, and hence to a corresponding increase in welfare even when volume is held constant.

Thus, in our two policy experiments, effective spread goes in the right direction when the tick size changes, but in the wrong direction as the gains to trade change. Further, in the former case, the magnitude of the change in effective spread (22.0%) bears no relationship to the change in welfare (2.7%). We can only conclude that the effective spread can be a very misleading proxy for welfare. Changes in volume are a good proxy for changes in welfare, provided the gains to trade remain approximately the same.

7 Conclusion

The method we introduce opens the door to a class of more realistic models that are closer to existing institutions. The explicit calculation of investor welfare makes it particularly useful for evaluating policy experiments.

In this paper, we use our model to determine the implication of endogenous order submission for the relationship between transaction prices, transaction costs, trader welfare and some of the commonly used proxies. We find that the midpoint of the quoted spread is an unbiased proxy for the consensus value on average in our symmetric model. However, conditional on a trade occurring it is not. We find that the effective spread is not a good measure of welfare because supply and demand of liquidity are endogenous. Thus, it should not be used to evaluate or motivate policy.

In terms of the model, there are many possible extensions such as including an intermediary, privately informed agents, or competing exchanges. Open questions include: What are reasonable proxies for welfare (to evaluate policy changes), transaction costs (to determine trading strategies), and the consensus value of the asset? Can these be inferred from real data? We hope to answer these questions in future work.

In addition to such market design and policy questions, this method should also be of use to practitioners. In particular, Lo, Mackinlay and Zhang (2002) report that hypothetical limit order executions are poor proxies for actual ones, suggesting the need for a structural model. We suspect that if practitioners work with a calibrated model of liquidity demand and supply that includes endogenous order flow the predicted estimates of price impacts will be better.

References

[1] Ahn, Hee-Joon, K. Bae and K. Chan (2001), 'Limit orders, depth and volatility: evidence from the stock exchange of Hong Kong,' *The Journal of Finance* Vol 61 No 2 p. 767–788.

[2] Bessembinder, Hendrik (1999), 'Trade Execution Costs on NASDAQ and the NYSE: a Post-Reform Comparison,' *Journal of Financial and Quantitative Analysis* 34 p387–407.

[3] Biais, B., P. Hillion and C. Spatt (1995) 'An Empirical Analysis of the Limit order book and the order flow in the Paris Bourse,' *The Journal of Finance* Vol 50 No. 5 p. 1655-1689.

[4] Biais, B. D. Matrimort and J–C Rochet, (2000): Competing mechanisms in a consensus value environment, *Econometrica*, Vol 68 No 4. p 799-837.

[5] Blume, M. and R. Stambaugh (1983) Biases in computed returns: an application to the size effect, *Journal of Financial Economics* 12 p 387-404.

[6] Bollen, Nicholas P.B. and Robert E. Whaley,(1998) 'Are "Teenies Better?" *Journal of Portfolio Management*, p 10–24.

[7] Choi, J. , D. Salandro and K. Shastri (1988) On the Estimation of Bid Ask spreads: theory and evidence, *Journal of Financial and Quantitative Analysis* 23 p 219–230.

[8] Cohen, K., S. Maier, R. Schwartz and D. Whitcomb (1981): Transaction costs, order placement strategy and the existence of the Bid – Ask spread, *Journal of Political Economy*, 89, 287–305.

[9] Coppejans, Mark, I. Domowitz and A. Madhavan (2001) 'Liquidity in an Automated Auction,' working paper.

[10] Demsetz, Harold (1968), "The Cost of Transacting," *Quarterly Journal of Economics* 82, 33-52.

[11] Easley, D. N. Kiefer and M. O'Hara, (1997) 'One Day in the Life of a Very Common Stock,' *Review of Financial Studies* Vol 10, No. 3 805-835.

[12] Edwards, Amy K. and J. Harris (2001), 'Stepping Ahead of the Book', SEC working paper.

[13] Foucault, T. (1999) : Order flow composition and trading costs in a dynamic limit order market, *Journal of Financial Markets* 2, 99-134

[14] Foucault, T. E.Kandel and O. Kadan (2002): The Limit Order Book as a Market for Liquidity, HEC working paper.

[15] Goldstein, Michael A. and K. Kavajecz (2000), 'Eights, sixtennths, and market depth: changes in tick size and liquidity provision on the NYSE,' Journal of Financial Economics 56, p125-149.

[16] Glosten, L. (1987) 'Components of the Bid/Ask spread and the statistical properties of transaction prices,' *Journal of Finance* 42, 1293-1307.

[17] Glosten, L. (1994): Is the electronic limit order book inevitable? *Journal of Finance*, Vol 49, 1127–1161.

[18] Glosten, L. (2000) Mimeo.

[19] Griffiths, Mark, B. Smith, D. Turnball and R. White (2000) 'The costs and determinants of order aggressiveness,' *Journal of Financial Economics* 56 p. 65-88.

[20] Handa, V., and R.A. Schwartz (1996): Limit Order Trading *Journal of Finance*, 51, 1835–1861

[21] Harris, L, and J.Hasbrouck (1996): Market V. Limit Orders: The SuperDot evidence on Order Submission , *Journal of Financial and Quantitative Analysis*, 31, 213–231.

[22] Hasbrouck, Joel and G. Saar, (2002) 'Lmit Orders and Volatility in a Hybrid Market: The island ECN,' Stern School working paper.

[23] Hasbrouck, J. (1991a) Measuring the information content of Stock trades, *Journal of Finance* 46 179–208.

[24] Hasbrouck, J. (1991b) The Summary informativeness of stock trades: An Econometric analysis, *Review of Financial Studies* 4, 571-595.

[25] Hasbrouck, J. (1993), Assessing the Quality of a Security Market: A new approach to transaciton-cost measurement. *Review of Financial Studies* Vol 6. No 1. p 191–212.

[26] Hollifield, B., R. Miller and P. Sandås (1999): An Empirical Analysis of Limit Order Markets, working paper.

[27] Hollifield, B., R. Miller, P. Sandås and J. Slive (2002): Liquidty Supply and Demand on The Vancouver Stock Exchange. Carnegie Mellon working paper.

[28] Huang, Roger D. and Hans R. Stoll (1997), *Review of Financial Studies* Vol 10 No 4 p 995–1034.

[29] Jones, Charles M. and M. Lipson, (2001) 'Sixteenths: direct evidence on institutional execution costs,' *Journal of Financial Economics*, 59 p. 253–278.

[30] Kavajecz, Kenneth A. (1999), 'A Specialist's Quoted Depth and the limit order book,' *Journal of Finance* Vol 54 No., 2 747-771.

[31] Keim, D. (1989) Trading Patterns, Bid-ask spreads and estimated security returns: the case of common stocks at calendar turning points, *Journal of Financial Economics* 25 p 75–97.

[32] Kumar, P. and D. Seppi (1993): Limit orders and market orders with optimizing traders, working paper, Carnegie Mellon University.

[33] Lehman, Bruce N. and David M. Modest (1994) 'Trading and Liquidity on the Tokyo Stock Exchange: A Bird's Eye View," *Journal of Finance* Vol 49 p951-984.

[34] Lo, Andrew, A. MacKinlay and J. Zhang (2002) 'Econometric Models of limit-order executions,' *Journal of Financial Economics* 65 p 31-71.

[35] Milgrom, P. and N. Stokey (1982), "Information, Trade, and Common Knowledge," *Journal of Economic Theory* 26: 17–27.

[36] Nyborg, K., K. Rydqvist and S. Sundaresan (2002), 'Bidder Behaviour in Multi–unit auctions: Evidence from Swedish Treasury Auctions,' *Journal of Political Economy* 110(2): 394-424.

[37] O'Hara,M. and G. Oldfield (1986): The microeconomics of market making *Journal of Financial and Quantitative Analysis*, 21, 361–376.

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

[39] Parlour, C. (1995): Price dynamics in Limit order Markets, Queen's University at Kingston, PhD. thesis.

[40] Parlour, C. (1998): Price dynamics in Limit Order Markets, *Review of Financial Studies* , 11, 789–816

[41] Parlour C. and D. Seppi (2003) Liquidity based competition for Order Flow, *Review of Financial Studies*. 16 No 2.

[42] Ranaldo, A. (2002) Order Aggresiveness in Limit order book markets, *Journal of financial markets* forthcoming.

[43] Rock, K. (1996): The specialist's order book and price anomalies, forthcoming in *Review of Financial Studies*, 10, 103–150.

[44] Roll, R. (1984) A Simple implicit measure of the effective Bid-Ask spread in an efficient market, *Journal of Finance* 39 p1073-1103

[45] Ronen, Tavy and D. Weaver (2001), "Teenies' Anyone?', *Journal of Financial Economics* (forthcoming).

[46] Sandås, Patrik (2001) Adverse Selection and Competitive Market Making: Empirical Evidence from a Limit Order Market, *Reviewe of Financial Studies* Vol 14 No3 p705–734.

[47] Securities and Exchange Commission (2001), 'Report on the Comparison of Order Executions across Equity Market Structures'. Office of Economic Analysis report.

[48] Seppi, Duane J. (1997), 'Liquidity Provision with Limit Orders and a Strategic Specialist,' *Review of Financial Studies* Vol 10., No.1 103-150.