



American Finance Association

Equilibrium in a Dynamic Limit Order Market

Author(s): Ronald L. Goettler, Christine A. Parlour, Uday Rajan

Source: *The Journal of Finance*, Vol. 60, No. 5 (Oct., 2005), pp. 2149-2192

Published by: Blackwell Publishing for the American Finance Association

Stable URL: <http://www.jstor.org/stable/3694745>

Accessed: 23/11/2009 12:17

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/page/info/about/policies/terms.jsp>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Please contact the publisher regarding any further use of this work. Publisher contact information may be obtained at <http://www.jstor.org/action/showPublisher?publisherCode=black>.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.



Blackwell Publishing and American Finance Association are collaborating with JSTOR to digitize, preserve and extend access to *The Journal of Finance*.

<http://www.jstor.org>

Equilibrium in a Dynamic Limit Order Market

RONALD L. GOETTTLER, CHRISTINE A. PARLOUR, and UDAY RAJAN*

ABSTRACT

We model a dynamic limit order market as a stochastic sequential game with rational traders. Since the model is analytically intractable, we provide an algorithm based on Pakes and McGuire (2001) to find a stationary Markov-perfect equilibrium. We then generate artificial time series and perform comparative dynamics. Conditional on a transaction, the midpoint of the quoted prices is not a good proxy for the true value. Further, transaction costs paid by market order submitters are negative on average, and negatively correlated with the effective spread. Reducing the tick size is not Pareto improving but increases total investor surplus.

WE CONSIDER A DYNAMIC PURE LIMIT ORDER market in which rational traders choose between buy and sell orders, and market and limit orders. We numerically solve for the equilibrium of the model and generate time series of trades and quotes. We characterize the equilibrium in terms of traders' strategies, transaction costs of market orders, and welfare accruing to both market and limit orders.

We find that the endogenous choice of orders has important implications for inferences drawn from transactions data. Agents in our model have a private value, or liquidity motive for trade, in addition to the common or consensus value of the asset.¹ Agents with low private values are eager to sell the asset, and often place ask quotes below the consensus value of the asset. Similarly, bid quotes are often above the consensus value. This has two immediate implications: First, many market order submitters benefit from negative transaction costs (i.e., they buy below or sell above the consensus value), and second, conditional on a trade, the midpoint of the bid–ask spread is not a good proxy for the asset's true value. Instead, the transaction price is closer to the true value of the asset.

*Goettler and Parlour are at the David A. Tepper School of Business at Carnegie Mellon University. Rajan is at the Stephen M. Ross School of Business at the University of Michigan. The authors thank Dan Bernhardt, Hank Bessembinder, David Easley, Burton Hollifield, Mark Ready, Patrik Sandås, Duane Seppi, Joshua Slive, Chester Spatt, Peter Swan, Tom Tallarini, and seminar participants at the Bank of Canada, Carnegie Mellon, Houston, Michigan, NYU, North Carolina, Wisconsin, and the EFA (2003), IFM2 (Montréal), NBER Microstructure (2003), SITE (2003), and WFA (2004) meetings. We are especially grateful to Tony Smith for ideas generated in a project on simulating limit order books. The paper has benefitted from the comments of four anonymous referees and the editor, Rob Stambaugh.

¹ Intuitively, the common value represents the true value of the asset (e.g., the present value of the future dividend stream), whereas a private value reflects idiosyncratic motives for trade (such as wealth shocks, tax exposures, or hedging and portfolio rebalancing needs).

A valuable feature of our model is that we can compute the welfare gain from each trade, since all traders maximize utility. This enables explicit welfare comparison across different market designs. As an example, we evaluate the effect of a reduction in the tick size. Consistent with previous intuition, we find that reducing the tick size benefits market order submitters, harms limit order submitters, and increases volume. The welfare gains outweigh the losses, leading to an increase in overall welfare. These findings are driven by the fact that with a reduced tick size, liquidity providers compete more aggressively.

The effective spread, defined as the average transaction price minus the midpoint of the contemporaneous bid and ask quotes, is often used as a proxy for transaction costs and implicitly for changes in the welfare of market participants. In a simulated time series of trades and prices, we find that the effective spread is negatively correlated with the true transaction costs paid by market order submitters. This is because quotes are set by limit order submitters who, in equilibrium, place aggressive orders. Furthermore, the effective spread is positively correlated with the total surplus generated by trade and with the surplus accruing to market order submitters—on average, a trader willing to incur a high transaction cost has a large benefit (i.e., welfare gain) from trade.

Therefore, measures developed for intermediated markets, such as the effective spread, must be interpreted with caution in limit order markets. In microstructure models based on Kyle (1985) or Glosten and Milgrom (1985), prices are set by an intermediary. Imposing a zero-profit condition on the intermediary generates prices that are equal to the expected value of the asset conditional on all public information and on a signed transaction. By contrast, in a pure limit order market, prices are determined by strategic traders and need bear no such relationship to the value of the asset.

Our model has rational liquidity traders who choose optimal submission strategies. Given the limit order book and common value (which are both publicly observed), each trader decides whether to buy or sell (or both), and at what prices. Traders arrive sequentially and submit orders to maximize expected surplus. For a limit order, expected surplus is computed using beliefs about the order's execution probability and the change in the asset's value conditional on execution. In equilibrium, order submissions generate actual execution probabilities and changes in common value before execution that match traders' beliefs. Adverse selection arises as limit buys execute more often when the value drops and limit sells execute more often when the value increases.² To examine how this "picking-off risk" affects our results, we consider alternative specifications of our model with no such risk. The results are qualitatively similar, and hence are driven not by stale limit orders, but rather by competition among limit order submitters.

The model is a stochastic dynamic game in which each agent chooses an action upon entry to the market. Since it is analytically intractable, we numerically solve for the equilibrium. Even a numerical solution using traditional

² Since Copeland and Galai (1983), the option value that a limit order provides to the market has been recognized as a potential cost to limit order submission.

techniques is difficult due to the size of the state space. Following Pakes and McGuire (2001), we deal with the curse of dimensionality by obtaining equilibrium values only on the recurrent class of states.

Our understanding of the tradeoffs involved in submitting limit orders has been enhanced by Cohen et al. (1981), Handa and Schwartz (1996), Chakravarty and Holden (1995), and Kumar and Seppi (1993) who analyze traders' choices between market and limit orders in different environments. Biais, Martimort, and Rochet (2000), Glosten (1994), O'Hara and Oldfield (1986), Rock (1990), and Seppi (1997) theoretically analyze prices and trading volumes in markets with limit order books.

Dynamic limit order markets are modeled by Parlour (1998), Foucault (1999), Foucault, Kadan, and Kandel (2003), and Rosu (2004). However, these models make restrictive assumptions to obtain analytical solutions. Parlour (1998) assumes a one-tick market and no volatility in the common value of the asset. Foucault (1999) allows for such volatility, but truncates the book. Foucault et al. (2003) consider a model in which traders differ in waiting cost, and examine spread dynamics and the resiliency of the limit order book. They require limit order submitters to undercut existing orders, without the option to submit orders at or away from the quotes. Rosu (2004) considers a continuous time version of the latter model with endogenous undercutting. Our solution technique allows us to solve a model flexible enough to characterize the evolution of the book at and away from the quotes.

An interesting empirical literature has shed light on both the characteristics of observed limit order books and the intuition gleaned from models. Biais, Hillion, and Spatt (1995) document the persistence of order flow on the Paris Bourse, and Hamao and Hasbrouck (1995) analyze trades and quotes on the Tokyo exchange. In our model, persistence is observed even when the consensus value of the asset is constant. Hence, microstructure effects unrelated to asymmetric information also contribute to persistence.

Sandås (2001) uses data from the Stockholm exchange to develop and test static restrictions implied by Glosten (1994). He 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. Hollifield, Miller, and Sandås (2004) use Swedish data to test a monotonicity condition generated by the equilibrium of a dynamic limit order market. They reject the condition when considering both buy and sell orders, and fail to reject when examining buys and sells separately. This provides some support for a dynamic model. Hollifield et al. (2004) use a similar technique to investigate the demand and supply of liquidity on the Vancouver exchange, and find that agents indeed supply liquidity when it is dear and consume it when it is cheap.

Our work is complementary to the literature pioneered by Demsetz (1968), Roll (1984), Glosten (1987), and Hasbrouck (1991a,b, 1993) that considers the relations among quoted spreads, transaction prices, and the true or consensus value of the asset in the presence of an intermediary. Since we generate artificial data, we know the true asset value in our limit order market, and consider some of the same issues.

I. 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. In each period, t , a new trader arrives at the market. The trader at time t is represented by a pair, $\{z_t, \beta_t\}$. Here, $z_t \in \{1, 2, \dots, \bar{z}\}$ denotes the maximum number of shares the trader may buy or sell. The decision to buy or sell is endogenous. Let F_z denote the distribution of z_t . The trader's private valuation for the asset, β_t , is drawn from a continuous distribution F_β with support $\mathcal{B}$. Both z and β are independently drawn across time, and their distributions are common knowledge.

The asset's common or consensus value, denoted v_t , is public knowledge at time t . In a frictionless market, all trades should occur at this price. Each period, with probability $\frac{\lambda}{2}$, the consensus value increases by one tick, and with the same probability decreases by one tick. Changes in the consensus value reflect new information about the firm or the economy. The periodic innovations in v_t imply that a trader who arrives at t is better informed than limit order submitters from previous periods. Thus, this is a model of asymmetric information, albeit a nonstandard one.

The market is an open limit order book. The agent in the market at time t can submit a market order, which trades against outstanding orders in the book, or a limit order at a specified price, which enters the book at that price. There is a finite set of discrete prices, denoted as $\{p^{-(N)}, p^{-(N-1)}, \dots, p^{-1}, p^1, \dots, p^{N-1}, p^N\}$, defined relative to the consensus value. The tick size is the constant distance, d , between any two consecutive prices. In our base case, we assume v_t lies between two possible prices, and p^1 is a half tick above the consensus value and p^{-1} is a half tick below the consensus value. An order to buy one share that executes at price p^i requires the buyer to pay $v_t + p^i$. We sometimes refer to price p^i as "tick i ."

Associated with each price $p^i \in \{p^{-(N-1)}, \dots, p^{N-1}\}$, at each time t , is a backlog of outstanding limit orders, ℓ_t^i . Buy orders are denoted as positive quantities and sell orders as negative ones. The limit order book, L_t , is the vector of outstanding orders, so that $L_t = \{\ell_t^i\}_{i=-(N-1)}^{N-1}$. A competitive crowd of traders provides an infinite depth of buy orders at price p^{-N} and sell orders at p^N .³

The quotes are set by traders who have chosen to supply liquidity. The bid and ask prices in the market at time t are defined in standard fashion—the ask price is the lowest sell price on the book, and the bid price the highest buy price. Therefore,

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

$$B_t = v_t + \max \{p^i \mid \ell_t^i > 0\}$$

$$\text{and} \quad A_t = v_t + \min \{p^i \mid \ell_t^i < 0\},$$

respectively. The midpoint of the bid and ask prices is $m_t = \frac{A_t + B_t}{2}$.

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

The trader who arrives at time t takes an action $X_t = \{x_t^i\}_{i=-N}^N$, where x_t^i denotes an integer number of shares to be traded at price p^i . To be feasible, an action must satisfy $\sum_{i=-N}^N |x_t^i| \leq z_t$. A buy (sell) order at price p^i is denoted by $x_t^i > 0$ ($x_t^i < 0$). An agent may submit orders that are in part market orders, and in part limit orders (possibly at different prices). In addition, she may submit both buy and sell orders. Finally, the agent is allowed to submit no order (i.e., submit an order of zero shares). The decision to trade is thus endogenous with respect to both quantity and direction.

Market orders submitted at time t execute in that period. Limit orders submitted at time t execute if a counterparty arrives at some time in the future. Since the consensus value can change, limit orders suffer from adverse selection. For example, suppose a trader at t places a limit buy order. If the consensus value falls, his order may represent a profitable trading opportunity for a new arrival and the trader runs the risk of being picked off. In equilibrium, adverse selection is a potential cost to submitting a limit order: Orders are more likely to execute if the asset value moves against them.

We expect traders to monitor the market, and to cancel their orders if the consensus value moves in an adverse direction. Rather than model this fully, we capture this intuition in reduced form with an exogenous cancellation function. Let τ be the time at which a particular share is placed in the book, and $t > \tau$ some future time at which the order for this share remains unexecuted. If $v_t < v_\tau$, buy orders in the book are likely to be mispriced, and if $v_t > v_\tau$, sell orders are mispriced. We consider a cancellation function, δ_t , that denotes the probability that a given share is cancelled at time t . This function δ_t is (weakly) increasing in $(v_\tau - v_t)$ for buy orders, and in $(v_t - v_\tau)$ for sell orders.⁴

In this paper we use piecewise-linear cancellation functions. Consider a share in the book that remains unexecuted at time $t > \tau$. Let $\underline{\delta} > 0$ denote a minimum cancellation probability, and $\bar{\delta} \geq \underline{\delta}$ a maximum probability (which may be as high as one). The share in question is cancelled at time t with probability $\delta_t(v_\tau, v_t) = \min\{\underline{\delta} + \hat{\delta}, \bar{\delta}\}$. For $\kappa > 0$, the adjustment term $\hat{\delta}$ is defined as

$$\hat{\delta} = \begin{cases} \kappa \max\{(v_\tau - v_t), 0\} & \text{for a buy order} \\ \kappa \max\{(v_t - v_\tau), 0\} & \text{for a sell order.} \end{cases} \quad (1)$$

The notion that $\underline{\delta} > 0$ captures the intuition that the outside opportunity set for a trader (not modeled in our paper) may change even when there is no change in the fundamentals for this particular asset. Similarly, in our base case (described in Section III below), we consider $\bar{\delta} < 1$, which reflects the notion that a trader may imperfectly monitor the market after placing an order, and (with some probability) may find that an order executed before his request to cancel reaches the market.⁵

⁴ Other infinite horizon models also have exogenous cancellation. For example, in Hollifield et al. (2004), each share in the book is cancelled with constant probability in each period. In Foucault (1999), each share is cancelled after one period.

⁵ If δ_t is a constant, the cancellation probability acts as a discount factor: It reduces the trader's expected payoff at a constant rate. In the absence of an incentive for traders to execute early, the model can be unstable. The cancellation function provides such an incentive in our model.




Figure 1. Evolution of a three-tick book, with no change in consensus value. The figure shows the sequence of events at each time t when orders can be submitted at only three prices relative to the consensus value. The depth in the book at time t and price p^i is denoted as ℓ_t^i , for each of the three prices. A trader arrives with type (β, z) , where β is his private value and z the number of shares he can buy or sell. The trader submits an order, where x_t^i denotes the number of shares submitted at price p^i . This leads to an intermediate book, where the depth at p^i is $(\ell_t^i + x_t^i)$. Each share in this book is cancelled, using the cancellation function δ_t . Finally, this leads to the new book before the trader at time $t + 1$ arrives. This figure assumes that the consensus value does not change. If the consensus value changes, it does so after the trader at t has submitted his order, but before any shares are cancelled.

The cancellation function is a proxy for a restricted form of re-entry—it may be interpreted, for each share, as the probability that a trader re-enters and cancels the order for that share. A richer model would also allow traders to submit new orders on re-entry. In such a model, a trader solves an infinite horizon dynamic program. In this paper, as a first cut, we allow re-entry for cancellation purposes only. As long as re-entry is probabilistic, the intuition obtained from our model will hold.

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 $(\ell_t^i + x_t^i)$. After the orders x_t^i have been submitted (and executed, if they are market orders), each remaining share in the book is cancelled with probability $\delta_t(\cdot)$. Figure 1 illustrates the sequence of events with three ticks, when the consensus value is unchanged.

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 increase in consensus value 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.

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.

Actions of subsequent traders may affect the priority of any limit order. The priority of an existing limit order decreases if a subsequent order is placed at a more favorable price. 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, thereby removing orders with higher time priority. Of course, the ultimate change in priority is execution—a counterparty takes the trade.

The payoff to a trader with private value β_t who submits an order for one share at time t and price p^i is

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

Since p^i is relative to the consensus value v_t , the transaction price in dollars is $p^i + v_t$. By the time the order executes at $\tau \geq t$, there may be new information about the asset, and the consensus value may have changed. The agent's valuation at execution is $\beta_t + v_\tau$.⁶

II. Equilibrium

We now discuss best responses and the existence of a stationary Markov-perfect equilibrium in this game, and present an algorithm for numerically finding such an equilibrium.

A. Best Responses

In period t , a trader of type (z_t, β_t) arrives at the market and submits an order X_t that specifies the number of shares to buy or sell at each price $\{p^{-N}, \dots, p^N\}$. He observes current market conditions, which consist of the consensus value, v_t , and the limit order book, L_t . Recall from Section I that the trader also knows the order cancellation function at times $t' > t$, denoted $\delta_{t'}(\cdot)$, the probability that v_t will change in any period, denoted λ , and the type distributions 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. 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.

⁶ This formulation of the payoff function is consistent with a model in which agents value the asset at a liquidation value of v_T in addition to their private value β_t . Conditional on execution, the payoff to a sell order is $(p^i + v_t) - (\beta_t + v_T)$. Adding and subtracting v_τ , and taking expectations conditional on execution at time τ , yields $p^i + v_t - v_\tau - \beta_t - E[(v_T - v_\tau) | v_\tau]$. The last term is zero (since innovations in the consensus value have mean zero), yielding the expression in (2).

Let $\mu_t^e(k, i, L, X)$ denote the period- t trader's belief of the probability of execution of his k^{th} share at price p^i given book L and order X . Similarly, let $\Delta_t^v(k, i, L, X)$ denote his expectation of the net change in the consensus value prior to execution (conditional upon execution). We assume traders are risk neutral. Hence, if this trader submits a buy order for one share at price p^i , his expected payoff is $\mu_t^e(k, i, L_t, X_t) \{\beta_t + \Delta_t^v(k, i, L_t, X_t) - p^i\}$. Similarly, the expected payoff from a sell order for one share at price p^i is $\mu_t^e(k, i, L_t, X_t) \times \{p^i - \beta_t - \Delta_t^v(k, i, L_t, X_t)\}$.

Beliefs are naturally different for market and limit orders. Since market orders execute immediately, $\mu_t^e(\cdot) = 1$ and $\Delta_t^v(\cdot) = 0$. A limit order submitted at t executes only at $(t + 1)$ or later. Since it may be cancelled in the interim, $\mu_t^e(\cdot) < 1$ for any limit order. Similarly, in equilibrium we expect $\Delta_t^v(\cdot)$ to be negative for limit buy orders, and positive for limit sell orders, since a limit order is subject to picking-off risk.

Given these beliefs, the risk-neutral trader optimally chooses

$$X_t = \arg \max_{\tilde{X} = (\tilde{x}^{-N}, \dots, \tilde{x}^N)} \sum_{i=-N}^N \sum_{k=1}^{|\tilde{x}^i|} \mu_t^e(k, i, L_t, \tilde{X}) (\beta_t + \Delta_t^v(k, i, L_t, \tilde{X}) - p^i) \text{sign}(\tilde{x}^i)$$

$$\text{subject to } \sum_{i=-N}^N |\tilde{x}^i| \leq z_t. \quad (3)$$

A strategy for an agent at time t , therefore, is a mapping $X_t : \mathcal{L} \times \mathcal{B} \times \{1, \dots, \bar{z}\} \rightarrow \{-z_t, \dots, z_t\}^{2N}$, 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, $\mu_t^e(\cdot)$, and changes in v given execution, $\Delta_t^v(\cdot)$.

B. Existence

The equilibrium concept we use is Markov-perfect equilibrium. Since time is not a state variable, the equilibrium must be stationary. The Markov specification requires agents to condition only on the current book, and not on any prior books. In our model, this is not restrictive, because the book summarizes the payoff-relevant history of play. In a stationary equilibrium, $\mu_t^e = \mu^e$ and $\Delta_t^v = \Delta^v$ for each t ; any two agents facing the same limit order book have the same beliefs about execution probabilities and changes in v conditional on execution. Of course, agents' beliefs must be consistent with the actual future course of play. Perfection further requires that beliefs of agents in states not visited must be reasonable.

The existence of a Markov-perfect equilibrium follows from standard results. Formally, a state is defined by the four-tuple (k, i, L, X) , which refers to the k^{th} additional share at price p^i added to limit order book L , when the overall action (including orders at prices other than p^i) is represented by X . Now, k , i , and all elements of L and X are integers. Furthermore, i is restricted to $\{-N, -(N - 1), \dots, N\}$ and $k \in \{-\bar{z}, \dots, \bar{z}\}$, where N and $\bar{z}$ are finite. Hence,

the action space is finite and the state space is countable. From Rieder (1979), it follows that this game has a Markov-perfect equilibrium.

We do not prove uniqueness. However, in keeping with the existing literature, we verify that the equilibrium appears to be computationally unique. We start the algorithm at different initial values, and ensure that it converges to the same equilibrium.

C. Solving for Equilibrium

Equilibrium is obtained by finding common beliefs, μ^e and Δ^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 Δ^v . Thus, once the algorithm has converged, we have found a stationary Markov equilibrium—traders' beliefs are consistent with the actual probabilities of future events.

To find this fixed point, we simulate a market session and update beliefs given the simulated outcomes until beliefs converge.⁷ We follow the philosophy of Pakes and McGuire (2001), and use a stochastic algorithm to asynchronously update beliefs.⁸ The advantage of this approach is two-fold. First, the updating of beliefs at a given state is computationally efficient. One way to update beliefs is to integrate over all possible sequences of future outcomes that lead to a share being cancelled or executed. Instead, we track each share in the book until it executes or is cancelled in the market simulation. Upon execution or cancellation, we update the values of μ^e and Δ^v for the state at which this share was submitted, by averaging this outcome with the previous outcomes for shares submitted at this state. In essence, this approach uses a single draw each time a state is visited 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. Hence, the fixed point is computed only for the recurrent class of states.⁹ As discussed in the introduction, the full state space for this game is too large for traditional numerical methods that operate over the entire state space.

A natural concern is that false beliefs at states outside the recurrent class may lead players to mistakenly avoid such states. Consider an extreme case in which all traders believe that limit orders never execute, so $\mu^e = 0$ for all orders. Suppose the book is empty when the trader at time one enters. Given

⁷ The C code for the simulation is available from the authors upon request.

⁸ Pakes and McGuire (2001) solve for equilibrium in a dynamic oligopoly, obtaining convergence in firms' value functions. Our traders take an action only at one point in time, which allows us to reformulate the problem in terms of traders' beliefs.

⁹ A recurrent class is a subset of states with the following properties: (i) regardless of the initial state, the system eventually enters the recurrent class, (ii) once entered, the probability of each state outside the recurrent class is zero, and (iii) each state in the recurrent class is visited infinitely often as t approaches infinity.

this belief, it is a best response for this trader to not submit an order. Hence, the book will be empty when the trader at time two arrives. This trader now faces a decision problem exactly identical to trader one, and hence submits no order. Therefore, no orders are ever submitted in this market.

This perverse situation is a stationary Markov equilibrium, but fails the requirement of perfection. If a trader actually did submit an order at a price close to v_t , it has a positive probability of execution. Hence, traders' beliefs at states outside the recurrent class are incorrect in this case. To ensure our numerical technique does not converge to such an equilibrium, and to ensure we find a perfect equilibrium, we specify initial beliefs to be overly optimistic. In the above example, this means that trader one submits some order, and some trading occurs in the market. Since initial beliefs are optimistic, states not in the recurrent class would not be visited even if beliefs about them were correct. Formally, therefore, we find a stationary Markov equilibrium that leads to the same outcome as a stationary Markov-perfect equilibrium.

In more detail, the algorithm works as follows. We assign optimistic initial beliefs to all states, $\{\mu_1^e(\cdot), \Delta_1^v(\cdot)\}$. We set $\Delta_1^v(\cdot)$ to $(p^N - p^i)$ for a buy order and to $(p^{-N} - p^i)$ for a sell order; these represent the maximum possible favorable changes in common value before execution. To assign μ_1^e for a limit buy at p^i , we consider $F_\beta(p^i)$, the probability that a trader who obtains a positive surplus from selling at p^i will arrive in the next period. The probability that the limit buy will survive until the next period is $(1 - \delta)$. Similarly, the likelihood that such a trader will not arrive in the next period is $(1 - F_\beta(p^i))$. Hence, execution at $\tau + 1$ 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,

$$\mu_1^e(\cdot, i, \cdot, \cdot) = \sum_{\tau=0}^{\infty} [(1 - \delta)^\tau (1 - F_\beta(p^i))^\tau (1 - \delta) F_\beta(p^i)] = \frac{(1 - \delta) F_\beta(p^i)}{1 - (1 - \delta)(1 - F_\beta(p^i))}. \quad (4)$$

The initial belief μ_1^e is similarly derived for limit sells.

We choose the initial book L_1 to be empty. Starting with $t = 1$, we iterate over the following steps.

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

Step 2: For each market order share, update μ_t^e and Δ_t^v for the initial state of the limit order executed against the market order. If the executed limit order is the k^{th} share submitted in period $\tau < t$ at price p^i ,

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

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

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 .¹⁰

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 $\delta_t(\cdot)$, and update μ_t^e for the states at which cancelled shares were submitted. This update uses only the first term in equation (5), since the second term has a numerator of zero for cancelled shares. Note that Δ_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 as follows:

$$v_{t+1} = \begin{cases} v_t + 1 & \text{with probability } \lambda/2 \\ v_t & \text{with probability } 1 - \lambda \\ v_t - 1 & \text{with probability } \lambda/2. \end{cases} \quad (7)$$

If v changes, shift the book to maintain the normalization that prices are relative to the current consensus value. 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 four. When v decreases, orders that were at tick $N - 1$ are processed similarly.

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

Every 10 million periods, we perform two additional adjustments. First, we reset n to one in step 2 above, until beliefs have begun to stabilize. This enables the algorithm to correct quickly for the excessive optimism of initial beliefs.¹¹ Second, there are states that traders chose in the past, but no longer choose due to the sufficiently negative outcomes. Beliefs at such states must also be (weakly) optimistic. To ensure this, we force updates of these states every 10 million periods, by starting the simulation at these states.

We stop the iterative process when beliefs satisfy a probabilistic criterion, similar to the test proposed by den Haan and Marcet (1994). After the algorithm appears to have converged, we hold beliefs fixed, simulate an additional 100 million periods, and record the frequency of limit order executions at each state. Consider the execution frequency of the k^{th} share submitted at price p^i as part of order X given book L , $\mu^e(k, i, L, X)$. Under the null hypothesis that

¹⁰ Since $\Delta_t^v(\cdot)$ refers to the net changes in consensus value conditional on execution, the count used in this updating (i.e., in equation (6)) can alternatively be the number of shares submitted at that state that eventually execute. This leads to faster convergence.

¹¹ The algorithm may be viewed as a behavioral description of players learning about the game in “real-time.” However, we only use the algorithm as a numeric solution technique to find the equilibrium. There is no connection between the rate at which our algorithm “learns” equilibrium beliefs, and the rate at which real traders may learn.

beliefs have converged to a fixed point, execution is a binomial process with probability of success μ^e and of failure $1 - \mu^e$. By the central limit theorem, the limiting distribution of the empirical execution frequency in each state is approximately normal with mean μ^e and variance $\frac{\mu^e(1-\mu^e)}{N_s(k, i, L, X)}$, where $N_s(k, i, L, X)$ denotes how often the state (k, i, L, X) occurs in the 100 million periods. The test statistic standardizes these normal variables and sums their squares. The statistic is chi-square with degrees of freedom equal to the number of states used in the summation. We only use states visited at least 100 times to ensure that the central limit approximation is accurate.¹² If the test statistic is less than the 1% critical value, we deem the algorithm to have converged.¹³

This probabilistic test does not involve Δ^v . Therefore, we also check the absolute differences in the realized outcomes (executions and net changes in v_t) and their respective beliefs (μ^e and Δ^v). We find that whenever the chi-squared test is satisfied, the weighted (by visitation frequency) average absolute differences over the 100 million periods are in both cases less than 1%.¹⁴

III. Base Case Parametrization

We do not formally calibrate the model, but choose parameter values that qualitatively capture salient market features while being consistent with computational tractability. We have experimented with different parameter values, and the results are qualitatively robust.

The following parametrization corresponds to our base case.

- There are eight ticks, with prices represented as $\{p^{-4}, p^{-3}, p^{-2}, p^{-1}, p^1, p^2, p^3, p^4\}$. We normalize the tick size to an eighth (so $d = \$\frac{1}{8}$, or 12.5 cents), as much of the literature on tick-size changes compares the move from eighths to sixteenths. The consensus value is midway between p^{-1} and p^1 . The “tick” of an order is defined to be the difference between the price at which the order is submitted and the current consensus value. Hence, the ticks range from -3.5 to $+3.5$. Traders may submit limit orders at ticks $\{-2.5, \dots, 2.5\}$. At ticks -3.5 and 3.5 , a trading crowd provides infinite liquidity.
- F_β is a normal distribution with mean 0 and standard deviation 2.8 ticks, or \$0.35.¹⁵ Hollifield et al. (2004) estimate for three stocks on the Van-

¹² For the base case, fewer than 1% of the states visited during the 100 million periods are not included in the summation.

¹³ As in den Haan and Marcet (1994), the “tolerance” level of this probabilistic stopping criteria is determined by the number of simulated periods used to construct the test. Eventually, as this number approaches infinity, the variance of the execution frequency approaches zero, and even minute discrepancies between execution frequencies and μ^e would lead to a rejection of the null hypothesis.

¹⁴ The stopping criterion in Pakes and McGuire (2001) compares beliefs about each state’s continuation value with an exact computation of the integral that defines this value. Our stopping criterion is akin to approximating this integral with the average realized values over long simulations with beliefs held fixed. The chi-squared test explicitly accounts for the noise introduced by this approximation.

¹⁵ The choice of F_β is not motivated by computational need. The algorithm can handle any distribution for β .

couver 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 (Table 9, p. 45). This implies a standard deviation of the private value distribution approximately equal to 4.5% of the value of the stock. Given our parametrization, this corresponds to a consensus value of approximately $\frac{2.8}{8}/.045 = \$7.78$.

- F_z assigns $z \in \{1, 2\}$ with equal probability. That is, each trader has, with equal probability, either one or two shares to trade.
- Each period, each share is cancelled with probability $\delta_t = \min\{\underline{\delta} + \hat{\delta}, \bar{\delta}\}$. Let $\tau \leq t$ denote the period in which the share was submitted. We set the minimum cancellation rate $\underline{\delta}$ to be 0.04, and the maximum $\bar{\delta}$ to be 0.64, with

$$\hat{\delta} = \begin{cases} 0.2 \max\{(v_\tau - v_t), 0\} & \text{for a buy order} \\ 0.2 \max\{(v_t - v_\tau), 0\} & \text{for a sell order.} \end{cases} \quad (8)$$

Suppose there is no change in v_t . Since the cancellation probability in this case is 0.04 in every period, if a share is not executed, the expected time before it is cancelled is 25 order-arrival periods.¹⁶ 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.

- The probability of an innovation in the consensus value is $\lambda = 0.08$, with increases and decreases in the consensus value equally likely. Since we do not incorporate a trend for the common value, all aspects of the market are symmetric.

The variance of the innovation distribution is λd^2 , where d is the dollar value of the tick size. Suppose intraday price changes have the same variance as innovations to the common value. If there are M transactions in the course of the day, the variance of daily returns will be $\lambda M d^2$. Setting M to 250 transactions per day, and using estimates of daily return volatility of between 0.2 and 0.4, yields a range for λ between 0.0512 and 0.1024. Our choice of $\lambda = 0.08$ is squarely in this range. The innovation in the consensus value of the asset is one tick. That is, in each period the asset value can increase or decrease by one tick.

IV. Simulation Results

Once the equilibrium is found, we fix beliefs and record a further 500,000 trader arrivals. We consider different parameterizations to evaluate the trade-off between limit and market orders. Conditional on execution, limit orders may receive better prices. However, limit orders might not execute, or may execute after an adverse change in the asset value. We can vary the elements of this tradeoff by varying the parameters of the model.

¹⁶ The probability that a share will last k periods is $\underline{\delta}^k (1 - \underline{\delta})^{k-1}$. Hence, the expected time until cancellation is $\underline{\delta}\{1 + 2(1 - \underline{\delta}) + 3(1 - \underline{\delta})^2 + 4(1 - \underline{\delta})^3 + \dots\} = \frac{1}{\underline{\delta}}$.

Outcomes differ across parameterizations due to both exogenous factors and endogenous behavior. Changes in the consensus value and cancellations lead to exogenous transitions between states. Furthermore, agents in the same state, but across different parameterizations, have different equilibrium strategies. This leads to differences in transition probabilities between states, and hence a difference in the frequencies of particular states. Thus, average market characteristics such as the bid–ask spreads differ. All the measures we report, such as spread frequencies and submission prices for limit orders, incorporate both these effects.

To provide some insight into equilibrium order submission strategies, we first provide an analytic characterization of traders' best responses in partial equilibrium. For a given book, we represent traders' strategies by threshold values of β at which traders are indifferent between two actions. We then show that equilibrium order submission has similar properties in our simulation.

A. Partial Equilibrium Characterization of Traders' Strategies

Consider a one-share trader, with private value β , who faces a given book and wants to sell the asset. The situation faced by a trader submitting a buy order is analogous. At the end of this section, we comment on the buy versus sell decision.

The trader chooses the optimal price at which to submit his sell order. He can submit a market sell at the bid price B , or a limit sell at any higher price. Denote the trader's belief about execution probability for an order placed at price p^i as $\mu_i^e = \mu_i^e(k, i, L_t, X_t)$. Similarly, denote $\Delta_i^v = \Delta_i^v(k, i, L_t, X_t)$. The trader submits a sell order at price p^i if

$$\mu_i^e(p^i - \beta - \Delta_i^v) \geq \max_{j \neq i} \mu_j^e(p^j - \beta - \Delta_j^v). \quad (9)$$

For a market order, $\mu_i^e = 1$ and $\Delta_i^v = 0$.

The payoffs from sell orders at prices p^i and p^j , where $p^i < p^j$, are depicted in Figure 2 as a function of β . The slope of the payoff line for price p^j is $-\mu_j^e$ (since these are sell orders), and the y-intercept is $\mu_j^e(p^j - \Delta_j^v)$. Given these payoff lines, agents with $\beta < \beta_{i,j}$ strictly prefer to submit at the more aggressive price p^i , and those with $\beta > \beta_{i,j}$ strictly prefer to submit at the higher price p^j . In this example, therefore, order submission is naturally characterized by a threshold type $\beta_{i,j}$, who is indifferent between the two orders.

Of course, there are many prices at which orders can be submitted, and hence many such "payoff lines" faced by each trader. To characterize order submission across all prices, we assume that execution probabilities are monotonic: Given a book, the execution probability of a sell order is decreasing in price (so that $\mu_i^e > \mu_{i+1}^e$), so more aggressive orders execute with greater probability.¹⁷ Similarly, conditional on execution, we expect lower-priced sell orders to have less picking-off risk; that is, $\Delta_i^v < \Delta_{i+1}^v$. A sell order at a higher price is likely to take longer

¹⁷ In our simulations, we find that equilibrium beliefs indeed exhibit this monotonicity, though it is not imposed by the algorithm.




Figure 2. Expected payoff to sell orders at different prices. The figure denotes the expected payoff to a seller from sell orders at prices p^i and p^j , where $p^i > p^j$. Here, μ_i^e is the seller's belief that a limit sell submitted at price p^i will eventually execute, β is his private value, and Δ_i^v is his expectation about the change in common value before the share executes. The notation for an order at price p^j is similar. The X -axis has the seller's β , and the Y -axis the expected payoff.

to execute, so there is greater chance of an adverse movement in the common value before it executes.

Let $\beta_{i,i+1}$ denote a trader who is indifferent between submitting an order at price p^i and one at price p^{i+1} . If the thresholds across different prices are ordered (so that, for sell orders, $\beta_{i+1,i+2} > \beta_{i,i+1}$, and so on), the probability of observing a sell order at price p^{i+1} is just $\text{Prob}(\beta \in (\beta_{i,i+1}, \beta_{i+1,i+2}))$. If, on the other hand, $\beta_{i,i+1} < \beta_{i-1,i}$, no order will be submitted at price p^i . If an additional condition on beliefs is satisfied, no price will be dominated. This will then imply that agents with the lowest β s submit market sell orders at price B , and as β increases, agents submit limit sells at prices $B + d$, $B + 2d$, and so on.

PROPOSITION 1: *Suppose execution probabilities are monotonic and, in addition,*

$$\frac{\mu_i^e}{\mu_i^e - \mu_{i+1}^e} (d - \Delta_{i+1}^v + \Delta_i^v) \geq \frac{\mu_{i+2}^e}{\mu_{i+1}^e - \mu_{i+2}^e} (d - \Delta_{i+2}^v + \Delta_{i+1}^v) \quad \text{for all } i. \quad (10)$$

Then, the threshold β s are also monotonic, so that $\beta_{i+1,i+2} > \beta_{i,i+1}$, and the probability of observing a sell order at price p^i is $F_\beta(\beta_{i,i+1}) - F_\beta(\beta_{i-1,i})$.

The amount $d - \Delta_{i+1}^v + \Delta_i^v$ is the gain from having a sell order execute at price p^{i+1} rather than p^i . The condition in the Proposition therefore reflects a

tradeoff between execution probabilities at the three different prices p^i, p^{i+1} , and p^{i+2} , and the relative gain in moving either from price p^i to p^{i+1} or p^{i+1} to p^{i+2} .

A market sell is an order at price B . If the conditions in Proposition 1 hold, market sells will be submitted by agents with the lowest values of β . That is, there will be some agent with $\beta = \beta_{B,B+1}$ indifferent between a market sell and a limit sell at price $B + d$. In this case, the probability of a market sell is $F_\beta(\beta_{B,B+1})$. All else equal, higher execution probabilities and lower picking-off risk for limit orders at some price above the bid will lead to more limit sells. Similarly, higher bid prices will lead to more market sells.

PROPOSITION 2: *Ceteris paribus, (i) a higher execution probability and a lower picking-off risk at some price $B + kd$, where $k \geq 1$, lead to a weakly lower probability of a market sell, and (ii) a higher bid price B leads to a strictly higher probability of a market sell.*

Variations in μ_i^e and Δ_i^v across states are determined in equilibrium. It is reasonable to suppose that thinner books, and an increase in the common value of the asset, increase execution probability. Therefore, we expect these factors to lead to a greater frequency of limit orders. We verify these in our simulations.

Finally, we consider the endogenous choice of buy versus sell orders. Fix a book, as before. Suppose a trader is choosing between a buy order at price p^j and a sell order at price $p^i > p^j$. The payoff to a buy order at price p^j is $\mu_j^e(\beta + \Delta_j^v - p^j)$. The payoffs for these two actions are depicted in Figure 3 below.




Figure 3. Payoff to sell order at price p^i and buy order at price p^j . This figure denotes the expected payoff to a seller from a limit sell at price p^i and a limit buy at price p^j . Here, μ_i^e is the seller's belief that a limit sell submitted at price p^i will eventually execute, β is his private value, and Δ_i^v is his expectation about the change in common value before the share executes. The notation for an order at price p^j is similar. The X-axis has the seller's β , and the Y-axis the expected payoff.




Figure 4. Base case: Optimal order choice for a one-share trader facing the empty book. This figure depicts the optimal action of a trader who can buy or sell one share, and enters the market when the book is empty. Simulation parameters correspond to the base case. The optimal action varies with the trader's private value, β . Here, p^{-1} is the highest price below the consensus value, and p^1 is the lowest price above the consensus value.

Recall that prices are quoted relative to v_t , so that p^i may be positive or negative. When $\beta = 0$, the buy order at p^j earns a payoff $\mu_j^e(\Delta_j^v - p^j)$, shown as positive in the figure (implying that p^j is below the common value v_t). Agents with $\beta > \beta_{i,j}$ choose the buy order at price p^j over the sell order at p^i ; those with $\beta < \beta_{i,j}$ choose the sell order at p^i . Ceteris paribus, an increase in the ask price leads to a lower likelihood of market buy orders, and an increase in the bid price leads to a lower likelihood of market sell orders.

B. Equilibrium Strategies in the Base Case

Now, we turn to a description of the results of the simulation in our base case. We first demonstrate how an agent's actions in the base case depend on the book. The book determines the opportunity set faced by an agent and the execution probabilities. Thus, the same agent may choose different actions depending on the book that he faces. In particular, the same trader type may choose to buy or sell, and to submit limit or market orders. We demonstrate this by considering the optimal action of an agent who only has one share to trade (so $z_t = 1$) given two of the 12,921 different books that arise during the simulation.

The first is the empty book,¹⁸ in which liquidity is only supplied by the trading crowd at the extreme prices, p^{-4} and p^4 . The second book we analyze has two limit buy orders at price p^{-1} . With this book, an arriving trader may submit a market sell order at p^{-1} , as well as a limit sell at $p > p^{-1}$ or a limit buy at any price.

First, consider the empty book. Given beliefs μ^e and Δ^v , we can identify the optimal order for all values of β_t . In our base case, the simulated values of β lie in the interval $(-12, 12)$.¹⁹ Over this interval, the optimal action of a trader faced with an empty book is depicted in Figure 4.

For low values of β , the agent is a seller and submits an aggressive limit order one tick below the consensus value. For slightly higher values of β , the

¹⁸ The empty book is the single most common book, and arises 1.4% of the time in the simulation.

¹⁹ Recall that β is drawn from a normal distribution with mean zero and standard deviation 2.8 ticks.




Figure 5. Base case: Optimal order choice for a one-share trader facing a book with a depth of two limit buys at p^{-1} . This figure depicts the optimal action of a trader who can buy or sell one share, and enters the market when the book has two limit buys listed at price p^{-1} , the highest price below the consensus value. Simulation parameters correspond to the base case. The optimal action varies with the trader's private value, β . Here, p^1 is the lowest price above the consensus value.

trader garners a lower benefit from trade and submits a limit sell above the consensus value of the asset. Such a trader is willing to risk forgone trade in order to extract more from subsequent traders.

Traders with a positive β submit buy orders. The symmetry of the model implies that traders with a large positive β submit aggressive limit buys (at p^1), and those with a low β submit conservative buys (at p^{-1}). Thus, given an empty book and a trader with $z_t = 1$, orders will be submitted only at the prices p^{-1}, p^1 . The states that have an agent submitting an order at any other price (given the empty book) are not in the recurrent class.

The same agent may submit different orders when faced with a different limit order book: Indeed, he may switch from buying to selling. Suppose the book already has two shares on the buy side at p^{-1} , and is empty at every other price. In equilibrium, faced with this book, a trader with one share takes one of four actions. Figure 5 represents the strategy of a trader with $z_t = 1$.

Order submission in equilibrium is therefore endogenous in terms of both the buy/sell decision and the price at which the order is submitted. In particular, facing a book with two limit buys at p^{-1} , a trader with $\beta_t \in [0, 0.16)$ prefers to sell rather than buy. If the book were empty, this trader would submit a limit buy order. In our base case, 4.7% of all market orders and 3.2% of all limit orders involve agents with $\beta_t < 0$ buying the asset, or agents with $\beta_t > 0$ selling the asset. Since order submission strategies differ across these two books, the probability of observing a particular order type (i.e., the probability that a trader arrives with a β valuation in the relevant region) is also different.

The intuition from Figures 4 and 5 carries through to other books and states. Typically, agents with extreme β values submit market and aggressive limit orders, as they are more desperate to trade. Agents with β near zero submit more conservative orders.

C. Average Book and Order Persistence

We present the average buy and sell sides of the book in Figure 6. While the total number of shares traded is large, on average the book is thin, suggesting




Figure 6. Average book depth at different ticks around v_t . The average depth (i.e., total number of shares) on the buy and sell sides of the book is shown. The ticks represent the difference between the price and the consensus value v_t .

that the market is effective at consummating trades. The average book has a total of 2.50 shares on the buy side, and 2.52 on the sell side. As expected, given our symmetric parametrization, the book is symmetric.

The transition between books has been studied empirically both to understand order placement strategies and to infer information-based trade. Order flow persistence is documented by Biais et al. (1995), Hamao and Hasbrouck (1995), and the subsequent literature.²⁰ Persistence is predicted for reasons other than information by Parlour (1998) in a model that exhibits positive serial correlation for market orders and negative serial correlation in limit orders. The probability of a limit sell order is highest conditional on a market buy in the previous period, followed by the conditioning events limit buy, market sell, and limit sell. Using Swiss limit order data, Rinaldo (2004) finds mixed support for her predictions. Ellul et al. (2003) find positive serial correlation in order type over short horizons on the New York Stock Exchange (NYSE). Consistent with their findings, our model exhibits positive serial correlation from transaction to transaction.²¹

²⁰ See, for example, Ahn, Bae, and Chan (2001), Griffiths et al. (2000), and Rinaldo (2004).

²¹ Such correlation may arise in mechanical fashion if large uninformed traders optimally split their orders to minimize their price impact (e.g., Vayanos (2001)). However, positive serial correlation in such a model occurs only if large buyers and sellers do not arrive contemporaneously in the market.

Table I
Conditional Frequencies of Buy Orders at Time $t + 1$ (Column) Given the Order at Time t (Row), Base Case

For the base case, the conditional frequency of each type of buy order at time $t + 1$ is shown down the columns, given the type of order observed at time t , which is shown along the rows. Market orders are classified as *Large* (if the order consists of multiple shares that execute at different prices) or *Small* (if the order executes at a single price). Limit buys are classified as *Aggressive* ("Agg.," if the order improves on (i.e., is priced above) the current bid price), *At-the-Quote* ("At Quote," if the order is at the current bid), or *Below-the-Quote* ("<Quote," if the order is priced away from (i.e., below) the current bid). Limit sells are classified similarly, with respect to the ask price.

Event at t	Market Buys			Limit Buys				Total Buys
	Large	Small	Total	Agg.	At Quote	<Quote	Total	
Large buy	0.59	24.09	24.68	27.12	0.09	0.00	27.21	51.89
Small buy	0.81	20.91	21.72	32.20	6.86	0.34	39.40	61.12
Large sell	0.00	0.09	0.09	27.08	14.59	5.97	47.64	47.73
Small sell	0.01	7.26	7.27	13.99	12.74	4.77	31.50	38.77
Agg. buy	0.02	11.71	11.72	5.63	22.86	7.72	36.22	47.94
Agg. sell	1.76	35.69	37.44	6.41	7.30	1.01	14.72	52.16
At quote buy	0.01	15.19	15.20	7.17	17.35	6.31	30.82	46.02
At quote sell	0.10	35.68	35.79	7.75	10.07	0.45	18.27	54.06
<Quote buy	0.00	1.56	1.56	2.57	26.89	14.10	43.57	45.13
>Quote sell	4.45	45.67	50.13	4.45	0.82	0.00	5.27	55.40
Overall	0.58	20.15	20.73	13.32	12.43	3.56	29.31	50.05

We first classify the types of buy orders submitted; sell orders are defined analogously. A Large Buy (*LB*) is a market order that moves the price. A Small Buy (*SB*) is a market order such that all shares are bought at the same price. Amongst limit orders, an Aggressive Buy (*AB*) is posted at a price higher than the current bid, an At-the-Quote Buy (*QB*) at the current bid, and a Below-the-Quote Buy (<*QB*) at a price below the current bid. Other than large market orders that necessarily require that two shares be traded, all order types may involve one or two shares.²²

We report the probability of observing an order at time $t + 1$, conditional on the action at time t , in Table I. Since we simulate a symmetric version of the model, we report only the frequencies of buy orders. In our base case, 50.0% of the total orders are buy orders, with 20.7% market, and 29.3% limit. The sum across each row is the conditional probability of a buy order at $t + 1$ given the conditioning event in the first column. Note that the frequency is reported as a percentage of all orders (buy and sell).

Biais et al. (1995) find patterns of trade on the Paris Bourse that are consistent with information effects. They identify a "diagonal effect": The conditional

²² 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 share to be the "first." If one share is submitted as a market order, it is the first share. When two different limit orders are submitted by the same trader, we randomly assign one share to be the first.

Table II
Conditional Frequencies of Buy Orders at Time $t + 1$ (Column)
Given the Order at Time t (Row) in a Model with No Changes
in the Consensus Value

Assuming no change in consensus value, the conditional frequency of each type of buy order at time $t + 1$ is shown down the columns, given the type of order observed at time t , which is shown along the rows. Market orders are classified as *Large* (if the order consists of multiple shares that execute at different prices) or *Small* (if the order executes at a single price). Limit buys are classified as *Aggressive* (“Agg.,” if the order improves on (i.e., is priced above) the current bid price), *At-the-Quote* (“At Quote,” if the order is at the current bid), or *Below-the-Quote* (“<Quote,” if the order is priced away from (i.e., below) the current bid). Limit sells are classified similarly, with respect to the ask price.

Event at t	Market Buys			Limit Buys				Total Buys
	Large	Small	Total	Agg.	At Quote	<Quote	Total	
Large buy	0.00	30.06	30.06	23.01	0.00	0.00	23.01	53.07
Small buy	0.08	21.74	21.81	21.33	14.19	0.01	35.53	57.35
Large sell	0.00	0.00	0.00	19.21	22.87	4.57	46.65	46.65
Small sell	0.00	9.61	9.61	9.86	20.65	2.60	33.10	42.71
Agg. buy	0.00	8.17	8.17	1.42	32.58	7.00	41.01	49.18
Agg. sell	0.41	38.21	38.61	1.62	10.06	0.38	12.06	50.68
At quote buy	0.00	19.51	19.51	2.57	23.52	2.76	28.85	48.36
At quote sell	0.00	28.25	28.26	3.66	19.56	0.02	23.24	51.49
<Quote buy	0.00	0.30	0.30	0.86	31.20	13.53	45.60	45.90
>Quote sell	0.61	51.48	52.10	2.08	0.14	0.00	2.23	54.32
Overall	0.06	20.49	20.55	7.93	19.58	1.92	29.43	49.98

probability of an order following a similar order is typically higher than the unconditional probability of such an order. We report similar persistence in Table I. For example, market buys are more likely after market buys than after market sells.²³ Furthermore, market orders are frequently followed by aggressive limit orders on the same side of the market.

In our model, such patterns emerge for two reasons. First, the impact of a change in the consensus value, v_t , induces subsequent traders to take similar actions, until the book has adjusted. For example, following an increase in v_t , sell orders previously on the book are priced “too low.” This should lead to a sequence of buy orders as subsequent traders pick off these limit orders.²⁴ Second, irrespective of asymmetric information, actions could be serially correlated because of persistence in the states.

In Table II, we report the order frequencies for a model in which the consensus value is fixed. Hence, no informationally motivated trade occurs. We still recover order persistence for small market buys and limit buys at the quote.

²³ This also accords with the empirical findings of Hollifield et al. (2002) and Rinaldo (2004).

²⁴ Such a dynamic is similar to that predicted by Kyle (1985) in an intermediated market. Faced with stale quotes, the informed trader submits orders until the price reflects his private information.

Table III
Depth at Bid and Frequency of Buy Orders at Different Spreads

For the base case, the frequency of each type of buy order is shown along the columns, given different depths (i.e., number of orders) at the bid price and different bid–ask spreads. Market buys are classified as *Large* (if the order consists of multiple shares that execute at different prices) or *Small* (if the order executes at a single price). Limit buys are classified as *Aggressive* (“Agg.,” if the order improves on (i.e., is priced above) the current bid price), *At-the-Quote* (“At Quote,” if the order is at the current bid), or *Below-the-Quote* (“Below Quote,” if the order is priced away from (i.e., below) the current bid).

Depth at Bid	Spread (ticks)	Number of Observations	Market Buys			Limit Buys				Total Buys
			Large	Small	Total	Agg.	At Quote	Below Quote	Total	
1–2	1	157,573	0.47	27.72	28.19	0.00	19.52	3.32	22.83	51.02
>2	1	69,201	0.20	31.59	31.79	0.00	13.48	2.99	16.46	48.26
1–2	≥2	52,649	0.22	11.15	11.37	28.88	8.36	1.98	39.22	50.60
>2	≥2	16,856	0.07	6.10	6.17	32.23	7.16	2.18	41.58	47.75

Table IV
Depth Below Bid and Frequency of Buy Orders at Different Spreads

For the base case, the frequency of each type of buy order is shown down the columns, given different cumulative depths (i.e., number of orders) at prices above the bid price and different bid–ask spreads. Market buys are classified as *Large* (if the order consists of multiple shares that execute at different prices) or *Small* (if the order executes at a single price). Limit buys are classified as *Aggressive* (“Agg.,” if the order improves on (i.e., is priced above) the current bid price), *At-the-Quote* (“At Quote,” if the order is at the current bid), or *Below-the-Quote* (“Below Quote,” if the order is priced away from (i.e., below) the current bid).

Depth Below Bid	Spread (ticks)	Number of Observations	Market Buys			Limit Buys				Total Buys
			Large	Small	Total	Agg.	At Quote	Below Quote	Total	
0	1	136,282	0.56	33.28	33.83	0.00	15.55	1.64	17.19	51.02
>0	1	90,492	0.13	22.31	22.44	0.00	20.87	5.60	26.47	48.91
0	≥2	55,905	0.23	11.54	11.77	31.14	6.23	1.17	38.55	50.32
>0	≥2	13,600	0.00	3.29	3.29	23.73	15.62	5.57	44.92	48.21

On the other hand, some patterns are indicative of information events. With no volatility in the consensus value, large orders are rare and not persistent. Furthermore, limit orders at the quotes are less likely to follow orders away from the quotes.

D. Depth in the Limit Order Book and Order Flow

As we have the complete limit order book in the simulations, we can investigate how depth in the book affects order placement. Tables III–VI document the relation between depth and the frequency of buy orders in the base case for

Table V
Depth at Ask and Frequency of Buy Orders at Different Spreads

For the base case, the frequency of each type of buy order is shown down the columns, given different depths (i.e., number of orders) at the ask price and different bid–ask spreads. Market buys are classified as *Large* (if the order consists of multiple shares that execute at different prices) or *Small* (if the order executes at a single price). Limit buys are classified as *Aggressive* (“Agg.,” if the order improves on (i.e., is priced above) the current bid price), *At-the-Quote* (“At Quote,” if the order is at the current bid), or *Below-the-Quote* (“Below Quote,” if the order is priced away from (i.e., below) the current bid).

Depth At Ask	Spread (ticks)	Number of Observations	Market Buys			Limit Buys				Total Buys
			Large	Small	Total	Agg.	At Quote	Below Quote	Total	
1–2	1	157,143	0.56	28.54	29.10	0.00	16.56	3.67	20.23	49.33
>2	1	69,631	0.00	29.71	29.71	0.00	20.19	2.19	22.38	52.09
1–2	≥2	52,435	0.24	9.42	9.66	28.32	8.98	2.33	39.63	49.29
>2	≥2	17,070	0.00	11.48	11.48	33.91	5.28	1.12	40.30	51.79

Table VI
Depth Above Ask and Frequency of Buy Orders at Different Spreads

For the base case, the frequency of each type of buy order is shown down the columns, given different cumulative depths (i.e., number of orders) at prices above the ask price and different bid–ask spreads. Market buys are classified as *Large* (if the order consists of multiple shares that execute at different prices) or *Small* (if the order executes at a single price). Limit buys are classified as *Aggressive* (“Agg.,” if the order improves on (i.e., is priced above) the current bid price), *At-the-Quote* (“At Quote,” if the order is at the current bid), or *Below-the-Quote* (“Below Quote,” if the order is priced away from (i.e., below) the current bid).

Depth Above Ask	Spread (ticks)	Number of Observations	Market Buys			Limit Buys				Total Buys
			Large	Small	Total	Agg.	At Quote	Below Quote	Total	
0	1	134,262	0.00	24.65	24.65	0.00	20.23	4.31	24.54	49.19
>0	1	92,512	0.94	35.08	36.02	0.00	13.96	1.63	15.59	51.61
0	≥2	55,758	0.00	6.71	6.71	31.14	9.03	2.53	42.71	49.42
>0	≥2	13,747	0.93	22.97	23.90	23.81	4.18	0.00	27.98	51.89

different spreads. We restrict attention to books that are nonempty on both the buy and sell sides to provide a cleaner interpretation of the results; otherwise, a low frequency of market orders could merely indicate that a large number of books considered were empty on the sell side.

From Table III, when the spread is one tick, higher depth at the bid increases the likelihood of market buys and decreases the likelihood of limit buys (especially limit buys at the bid). This is intuitive, since shares at the bid provide direct competition for a limit buy order. When the spread is two or more ticks, increased depth at the bid decreases the possibility of market buys and increases the frequency of aggressive limit buys that raise the bid. When the spread is

wide, an aggressive limit order provides an intermediate degree of immediacy, since it offers priority over the rest of the book.

By contrast, from Table IV, increased depth below the bid uniformly reduces the frequency of both market and aggressive limit buys and increases buys at and below the bid. This pattern also occurs in the model with no change in common value, which has no picking-off risk. On the surface, this is puzzling—increased competition in the form of depth below the bid should result in more, not less, aggressive orders.

There are two ways in which depth below the bid can occur. First, a trader submits a limit buy below the bid. In over 99.9% of such cases, the bid is above the common value. Second, a trader submits an aggressive order when the book is nonempty. In approximately 50% of such cases, such an order results in a bid above v_t . Therefore, all else equal, higher depth below the bid suggests that the current bid is “too high,” resulting in more conservative order placement.

Increased depth at the ask leads to a higher frequency of buy orders, as reported in Table V. The proportion of market buys increases slightly; at a spread of one tick, so does the proportion of limit buys at the bid. When the spread is two ticks, increased ask depth results in an increase in aggressive limit buys that raise the bid. Order submission below the quotes is deterred. These results are also counterintuitive as competition among sellers should make buyers less aggressive.

Depth at the ask suggests that the ask is close to the common value. In a one-tick market, this also suggests that the ask and bid are more likely to lie on either side of the true value. Orders below the bid are therefore also far from the true value, leading to lower execution probabilities. In a two-tick market, it implies that the bid is somewhat below the consensus value, leading to a greater frequency of aggressive limit buys.

Corresponding to the intuition that depth below the bid suggests the bid is “too high,” depth above the ask suggests the ask is “too low.” Hence, as shown in Table VI, increased depth above the ask leads to more aggressive order submission.

V. Transaction Costs and Welfare

The endogenous choice between limit and market orders determines how the gains to trade are split among traders. One measure of this split is the transaction costs paid by market order submitters. We define the true transaction cost for a market order as the difference between the average transaction price and the true value of the asset. We also show the nature of the relation between the effective spread and the surplus or welfare gain accruing to the respective market and limit order submitters.

A. Transaction Costs and Welfare Measures

Consider a trader at time t who submits a market buy 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.

In a frictionless market, all trades should occur at the true value of the asset, v_t . In keeping with theoretical work (such as Seppi (1997) and Foucault (1999)), we define the true transaction cost paid by a market order as the difference between the execution price and the true value of the asset. This is also consistent with the price impact cost of Keim and Madhavan (1998). We also define the effective spread, a commonly used proxy for transaction costs. Recall that m_t is the midpoint of the bid and ask quotes.

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

$$C_t(\bar{x}) = (P_t(\bar{x}) - v_t) \text{sign}(\bar{x}). \quad (11)$$

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

$$S_t(\bar{x}) = (P_t(\bar{x}) - m_t) \text{sign}(\bar{x}). \quad (12)$$

The true transaction cost is necessarily unobservable in real data. In many econometric specifications (see Hasbrouck (2002) for a summary), the execution price is decomposed into the sum of the “efficient price” and microstructure effects. In our model, the efficient price is just the consensus value, v_t . Thus, our transaction cost C_t is simply the microstructure effect times the signed order flow. Our sellers pay a cost if the transaction price is greater than the consensus value and receive a discount if it is below.

Consider a trade that occurs at time t . The consumer surplus that accrues to the market order and limit order submitters is a measure of the net change in their welfare.²⁵ 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 surplus accruing to the market order submitter is*

$$W_t^m = \bar{x}(\beta_t + v_t - P_t(\bar{x})), \quad (13)$$

where P_t is the average execution price, and

(ii) *the surplus accruing to limit order submitters taking the other side of the transaction is*

$$W_t^l = \bar{x}(P_t(\bar{x}) - (v_t + \beta_{(t)}^l)), \quad (14)$$

²⁵ Note that, since we examine a market for a single asset, this may be a limited measure of the overall change in welfare resulting from some trading strategy employed by an investor. If there are utility gains to an investor from optimally rebalancing portfolios, or from trading in other financial assets in some way, these are only captured in our measures to the extent that they are reflected in the private value (or β) of the investor.

where $\beta_{(t)}^l$ is the share-weighted average of the private values of all limit order submitters whose orders trade against the market order at time t .

In the absence of an explicit welfare measure to rank market outcomes, one alternative is to rely on observed proxies for transaction costs to estimate changes in welfare. When $m_t = v_t$, the surplus of a market order submitter can be written in terms of the effective spread. However, if $m_t \neq v_t$, this is no longer true.

PROPOSITION 3: *Suppose a trade of size $\bar{x}$ occurs at time t at an effective spread of $S_t(\bar{x})$. If (and only if) $m_t = v_t$,*

- (i) *the surplus of the market order submitter is $W_t^m = \bar{x}\beta_t - |\bar{x}|S_t$, and*
- (ii) *the surplus accruing to the limit order submitters who trade at t is $W_t^l = |\bar{x}|S_t(\bar{x}) - \bar{x}\beta_{(t)}^l$.*

We report the surplus (measured in ticks of $\frac{1}{8}$) for both market order and limit order submitters. For policy purposes, the surplus of limit order submitters should also be considered. Hampered by the lack of order level data, researchers have typically computed transaction costs for market orders. However, there is no reason why one group of investors should be favored over another. Notice that, even if $m_t \neq v_t$, the aggregate surplus improvement from a trade at t is $\bar{x}(\beta_t - \beta_{(t)}^l)$, which is independent of the effective spread, $S_t(\bar{x})$. In a pure limit order market, transaction costs are simply transfers among agents. Hence, any measure that determines a cost to one party merely reflects a gain to the counterparty.

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) surplus. While the intermediary provides a benefit by providing liquidity to market orders, he may also deter limit order submission and thus decrease the surplus of such agents (see Seppi (1997)). As Glosten (1998) observes in this case, one should account for the surplus of all parties in the market.

B. Comparison of Transaction Costs in Three Models

We first summarize our results; the rest of this section explores the intuition underlying them. In the base case, the average true transaction cost to market orders is negative. That is, on average market orders execute at prices better than the current consensus value. As we show in this section, limit order submitters in our base case suffer from picking-off risk.²⁶ However, the result that many market order submitters benefit from negative transaction costs is not simply an artifact of stale limit orders. To demonstrate this, we simulate two

²⁶ In equilibrium, they compensate for this risk by choosing a submission price at which they earn a positive payoff, as we show in the next subsection.

Table VII
Summary Description of Three Models

This table describes the parameters used for the three cases considered in the paper. Here, λ is the probability that the consensus value changes in a given period; $\underline{\delta}$ is the minimum cancellation rate; $\hat{\delta}$ the increase in the cancellation rate for a share when the current consensus value is v_t and the share was submitted when the consensus value was v_τ ; and $\bar{\delta}$ is the maximum cancellation rate.

Model	λ	$\underline{\delta}$	$\hat{\delta}$	$\bar{\delta}$
Base case	0.08	0.04	Buy order: $0.2 \max\{(v_\tau - v_t), 0\}$ Sell order: $0.2 \max\{(v_t - v_\tau), 0\}$	0.64
Zero volatility	0	0.04	n/a	0.64
Immediate cancellation	0.08	0.04	Buy order: $0.96 \max\{(v_\tau - v_t), 0\}$ Sell order: $0.96 \max\{(v_t - v_\tau), 0\}$	1
		0.04		1

Table VIII
True Transaction Costs and Effective Spread

The means of the true transaction cost (C_t) and the effective spread (S_t), and the correlation between the two measures, are shown for the three models considered. Standard errors for the means are in parentheses. Means and standard errors are reported in ticks.

Model	Mean C_t	Mean S_t	Correlation (C_t, S_t)
Base case	-0.10 (0.0009)	0.90 (0.0012)	-0.29
Zero volatility	0.18 (0.0009)	0.81 (0.0009)	-0.84
Immediate cancellation	-0.02 (0.0010)	0.96 (0.0012)	-0.28

models with no picking-off risk. The first of these, the zero volatility model, has no change in the consensus value (i.e., $\lambda = 0$). In this model, the average transaction cost for market order submitters is positive, but 31.3% of market orders obtain a negative transaction cost. In the second model (the immediate cancellation model) limit orders are cancelled with probability one if the consensus value changes in an adverse direction.²⁷ Hence, there can be no picking-off risk. In this model, a limit buy (sell) submitted at t is cancelled if the consensus value falls below (rises above) v_t . Here, transaction costs are closer to zero, but are again negative on average. Table VII outlines the differences in the three models compared in this section.

For each model, the means of the true transaction cost (C_t) and the effective spread (S_t), and the correlation between the two, are reported in Table VIII. The table also displays the standard errors for the mean transaction cost and effective spread. Since we have a large number of traders (500,000), the standard errors on the sample means are sufficiently low such that all differences in means on the order of 10^{-2} or higher are significantly different from zero. Hence, in what follows, we no longer report the standard errors.

²⁷ Such a model was suggested by David Easley.

In all three models, the effective spread is negatively correlated with the true transaction costs, performing worst (in terms of correlation) in the zero volatility model. The true transaction costs are negative in the base case and immediate cancellation models, and positive in the zero volatility model. However, even in the latter case, the average transaction cost (0.18 ticks) is less than half the minimum effective spread (0.5 ticks) on any transaction.²⁸ A large number of market orders (31.3%) transact at a negative cost even in this model.

In all three models, over 90% of the transactions incur costs of ± 0.5 ticks. This suggests that, for the most part, market orders are submitted only when the bid and/or ask prices are near the consensus value. Since transactions are endogenous, observing a market order merely tells us that the trader concerned is willing to incur the transaction cost at the given price, which in turn suggests that this transaction cost is low.

In the base case, even though the average transaction costs are negative, 44.7% of all market orders pay positive transaction costs, though very few (1.6% of all market orders) incur a cost strictly greater than a half tick. The more desperate an agent is to trade, the larger the transaction costs she is willing to bear. On average, market buy orders with $\beta \leq |2.5|$ benefit from a negative transaction cost in the base case; those with $\beta > |2.5|$ pay a positive transaction cost.²⁹

There are two reasons why negative transaction costs occur in our model. In some instances, they reflect stale quotes, due to an adverse change in the common value. However, they may also be an earlier trader's attempt to achieve quick execution at a low cost by submitting an aggressive limit order. For example, suppose a trader with a high β enters the market when the quoted spread is one tick, and the ask is a half tick above the common value. Then, he submits a market buy order. Suppose, instead, he comes in when the quoted spread is two or more ticks, and the current bid is a half tick below the common value. In this case, he has the choice of submitting an expensive market order or an aggressive limit order that increases the bid. The latter also buys quicker execution, since it has higher priority than any existing limit order. While not providing the immediacy of a market buy, it reduces the waiting time in the queue. Such an order potentially provides a negative transaction cost to the next trader who submits a market sell order. In the zero volatility and immediate cancellation models, this is the only source of negative transaction costs.

Greater competition in liquidity provision leads to better terms of trade for market order submitters. Thus, any aspect of the model that fosters such competition will imply lower transaction costs. In the base case, picking-off risk deters undercutting. Such risk is absent in the immediate cancellation model, which leads to more aggressive orders. This effect is reinforced by the higher

²⁸ Since the minimum bid-ask spread is one tick, the minimum effective spread is 0.5 ticks.

²⁹ This is consistent with the evidence of Chan and Lakonishok (1995), who find that money managers with a high demand for immediacy are associated with a larger price impact on their trades. Similarly, Keim and Madhavan (1997) show that index traders (who are essentially forced to trade to match the index) pay a higher transaction cost than value traders, whose trade is fully endogenous.

cancellation probability. In the zero volatility model, the absence of picking-off risk is counterbalanced by the greater patience (i.e., lower cancellation probability) of limit order submitters.

Heightened competition in the immediate cancellation model is exhibited by orders submitted against the empty book. In Figure 7, we display the orders submitted by traders with one share who enter the market when the book is empty. Buy orders at p^1 (above the consensus value) and sell orders at p^{-1} (below the consensus value) are aggressive orders. With immediate cancellation, the




Figure 7. Comparison of three models: optimal order choice for a one-share trader facing the empty book. For each of the three cases considered, this figure depicts the optimal action of a trader who can buy or sell one share, and enters the market when the book is empty. The optimal action varies with the trader's private value, β . Here, p^{-1} is the highest price below the consensus value, and p^1 is the lowest price above the consensus value.

Table IX
Further Comparison of Three Models

For each of the three cases considered, the table shows the average depth (i.e., the total number of shares) listed on the buy side of the book, the average time between a limit order being submitted and its executing (this average ignores shares cancelled before execution), and the proportion of limit buy orders (as a percentage of all limit buys) that are submitted at prices below the consensus value.

Model	Mean Book Depth (Buys)	Mean Time to Execution (Limit Buys)	Proportion of Limit Buys Placed Below v_t (%)
Base case	2.50	4.57	58.1
Zero volatility	3.41	7.40	75.4
Immediate cancellation	2.27	4.12	55.7

range of traders who submit aggressive orders is higher than in the base case, which in turn sees more aggressive order submission than with zero volatility. For example, considering buy orders, the range of trader types submitting aggressive buy orders is $\beta_t \geq 2.91$ with immediate cancellation, $\beta_t \geq 3.19$ in the base case, and $\beta_t \geq 4.23$ with zero volatility.

This variation in actions given a state leads to differences in the distribution of states encountered in the different models. Thus, market aggregates also differ. As Table IX indicates, across all three models, the highest proportion of limit buys placed below the consensus value occurs in the zero volatility model (i.e., there are fewer aggressive orders). Recall that this is the model in which market orders pay the largest transaction costs. On average, the book is thicker and limit orders take longer to execute. Orders are cancelled less frequently, and limit order traders can afford to wait in the hope of extracting surplus from subsequent traders who are keen to trade. Thus, the market order submitters who do submit orders value trade very highly and are willing to accept a positive transaction cost.

Volatility in the consensus value leads to more aggressive order submission, in both the base case and with immediate cancellation. In both these cases, fewer buy orders are submitted below the consensus value (i.e., orders are more aggressive), as shown in Table IX. Conditional on execution, the average time to execution is lowest with immediate cancellation, as is the proportion of buy (sell) orders placed below (above) the consensus value. Orders are also more spread out; more orders are submitted at ticks farther from the consensus value than with zero volatility. Naturally, all this implies lower transaction costs for market order submitters.

The absence of picking-off risk encourages aggressive order submission when there is immediate cancellation. Traders can submit orders closer to their own valuation with no risk of suffering a loss. Conversely, they sometimes benefit from the consensus value moving in a favorable direction (e.g., increasing in the case of a limit buy) before their order is executed. This allows for a high expected payoff despite being an aggressive order. Hence, limit order submitters

are willing to compete aggressively for execution with immediate cancellation. With zero volatility, there is no possibility of a favorable change in consensus value. Therefore, orders are more conservative.

In summary, we find that transaction costs are negative for many market orders. They are also negative on average when the consensus value is volatile, even in the absence of picking-off risk. Volatility encourages competition among limit order submitters, regardless of whether picking-off risk is present. This leads to more aggressive limit order submission and lower transaction costs for market orders.

C. Quantifying Picking-Off Risk

In our base case, over 80% of all executed limit orders face no change in the consensus value before execution. However, on average, limit buy (sell) orders are more likely to execute after a decrease (an increase) in the consensus value of the asset. The average change in consensus value before a limit buy (sell) order executes is -0.08 ($+0.08$) ticks.

In Figure 8, we illustrate the picking-off risk faced by limit order submitters. The horizontal axis records the number of net changes in consensus value before




Figure 8. Net change in v_t before a limit order is executed, and trader surplus. The X-axis of this figure shows the net change in consensus value, v_t , between the time a limit order is submitted and the time it is executed. The columns, measured along the left-hand Y-axis, denote the proportion of limit buy and sell orders that experience each net change in v_t . The lines, measured along the right-hand Y-axis, denote the average trader surplus.

Table X
Frequency (%) of Buy Orders Conditional on Changes in the Consensus Value

For the base case, the frequency of each type of buy order is shown down the columns, given different changes in the consensus value. Market buys are classified as *Large* (if the order consists of multiple shares that execute at different prices) or *Small* (if the order executes at a single price). Limit buys are classified as *Aggressive* ("Agg.," if the order improves on (i.e., is priced above) the current bid price), *At-the-Quote* ("At Quote," if the order is at the current bid), or *Below-the-Quote* ("<Quote," if the order is priced away from (i.e., below) the current bid).

Change in v_t	Market Buys			Limit Buys				Total Buys
	Large	Small	Total	Agg.	At Quote	<Quote	Total	
Down	0.04	6.03	6.07	13.88	10.66	15.19	39.73	45.80
None	0.53	20.12	20.66	13.33	12.83	3.24	29.40	50.05
Up	2.31	35.26	37.57	12.52	3.89	0.20	16.60	54.17
Overall	0.58	20.15	20.73	13.32	12.43	3.56	29.31	50.05

a limit order is executed. The bars represent the proportion of executed limit buys or sells, using the left-hand scale. The lines represent the mean surplus of the limit order trader on execution, using the right-hand scale. The possibility of adverse selection is taken into account by limit order traders; only 2.8% of all executed limit orders result in a loss for the submitter.³⁰

In Table X, we examine the frequency of buy order submission based on changes in the consensus value. As expected, market buy orders are more frequent after an increase in the consensus value than after a decrease.³¹ This exemplifies picking-off risk for limit orders in the book. If the consensus value increases, last period's ask becomes "too low," offering new traders a profitable buy opportunity. Of course, last period's bid is also "too low," resulting in fewer limit buys at or below the bid. The overall frequency of buy orders increases with an increase in the consensus value, implying that some trader types shift from sell to buy orders in this case.

D. The State and Evolution of the Bid-Ask Spread

In Table XI, we report the unconditional probability (as a percentage) of observing market or limit buy orders given the spread. A market order is less likely when the bid-ask spread is wide, while limit orders are more likely. When spreads are wide, market orders are more expensive, and thus traders tend to submit aggressive limit orders that narrow the spread.³² For example, at spreads between five and eight ticks, virtually all orders are aggressive limit

³⁰ This is consistent with Handa and Schwartz (1996) who show that hypothetical limit orders earn positive returns on average, and Harris and Hasbrouck (1996) who find evidence of optimal limit order submission on the NYSE SuperDOT system.

³¹ This is consistent with Harris (1998).

³² This argument was first made and documented by Biais et al. (1995).

Table XI

Frequency (%) of Different Buy Orders for Different Bid-Ask Spreads

For the base case, the frequency of each type of buy order is shown down the columns, given different values of the bid-ask spread. Market buys are classified as *Large* (if the order consists of multiple shares that execute at different prices) or *Small* (if the order executes at a single price). Limit buys are classified as *Aggressive* ("Agg.," if the order improves on (i.e., is priced above) the current bid price), *At-the-Quote* ("At Quote," if the order is at the current bid), or *Below-the-Quote* ("<Quote," if the order is priced away from (i.e., below) the current bid).

Bid-Ask Spread	Market Buys			Limit Buys				Total Buys
	Large	Small	Total	Agg.	At Quote	<Quote	Total	
1	0.42	28.90	29.32	0.00	17.60	3.28	20.87	50.19
2	1.64	14.05	15.69	21.47	7.14	5.58	34.19	49.88
3	0.76	23.40	24.16	6.20	13.55	6.15	25.91	50.06
4	0.00	5.62	5.62	37.53	6.66	0.00	44.19	49.81
5	0.00	0.10	0.10	49.75	0.00	0.00	49.75	49.85
6	0.00	0.00	0.00	50.33	0.00	0.00	50.33	50.33
7	0.00	0.00	0.00	50.18	0.00	0.00	50.18	50.18

orders. When spreads are narrow, traders tend to take liquidity by submitting market orders. The rapid response of traders to profit opportunities ensures that spreads are narrow.

The frequency of market orders is nonmonotone in the bid-ask spread. In particular, market orders are most frequent when the spread is one tick, and are more frequent with a three-tick spread than a two-tick spread. This is mirrored in the submission pattern of limit orders. When the spread is two ticks, there is a large increase in aggressive limit orders. Such orders reduce the spread to one tick, and therefore cannot be undercut in the next period.

Figure 9 exhibits the frequency distribution of the bid-ask spread in the base case ($\lambda = 0.08$) and the zero volatility case ($\lambda = 0$). In the base case, the modal bid-ask spread is one tick (with frequency 43.99%), and the mean is 2.25 ticks. In the zero volatility model, spreads are on average narrower, reflecting the lack of asymmetric information. In the immediate cancellation model, the average spread is even wider: 2.52 ticks. While Foucault (1999) predicts that more volatile assets have wider spreads, and we predict the same, our results are driven by different effects. In Foucault, if assets are more volatile, limit orders are more likely to be picked off and thus require wider spreads as compensation. We find that spreads are widest in the immediate cancellation model in which orders are never picked off. Our effect arises even though limit order submitters are more aggressive (i.e., compete more fiercely), because the exogenous volatility of the asset causes more frequent cancellation of orders, and hence wider spreads on average.

Comparing the base case against the zero volatility model, we find that an increase in λ from 0 to 0.08 leads to an increase of 0.33 ticks or 17.2% in the average quoted spread. Glosten (1987) decomposes the quoted spread into order processing and adverse selection components. While his market-maker framework




Figure 9. Frequency distribution of quoted bid-ask spreads. This is a histogram of the quoted bid-ask spread in two cases: (i) the base case, with $\lambda = 0.08$, and (ii) no change in consensus value, with $\lambda = 0$. Here, λ is the probability that the consensus value changes in each time period.

is not directly applicable to a limit order market, it is interesting to note that information does play a role in determining the quoted spread in a limit order market: Limit orders are placed by traders who are aware that they will execute against market order submitters with superior information.

E. The Midpoint as a Proxy for the Consensus Value

Over the 500,000 simulated periods of our base case, the mean difference between the midpoint of the bid-ask prices and the consensus value (i.e., the mean of $(m_t - v_t)$) is -0.002 ticks, with a standard error of 0.002 ticks. Thus, the midpoint is an unbiased estimator of the consensus value. However, it is frequently incorrect—in 25.9% of the periods, the bid-ask quotes do not contain the consensus value. This happens for two reasons. First, a trader may optimally submit a limit buy (sell) order above (below) the consensus value if the current ask (bid) is “too high” (too low). Second, a change in the consensus value may render the current quotes stale. In the zero volatility model, in which the second effect is absent, the consensus value lies outside the quotes in 13.8% of the periods.

In practice, the midpoint is often used to infer the consensus value when a transaction occurs, as in empirical measures of transaction costs. Thus, we next examine the difference between the midpoint and consensus value conditional




Figure 10. Histogram of midpoint minus true value, conditional on trade. This figure provides a histogram of the difference between the midpoint of the bid–ask spread and the consensus value of the asset, conditional on a transaction occurring in that period. The transaction involves either a market buy order or a market sell order.

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 10 plots the distribution of $(m_t - v_t)$ conditional on a market buy and a market sell in that period. 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 (sell), the true value of the asset is on average 1.02 ticks higher (lower) than the midpoint. An error of one tick is insignificant when measuring, for example, the return on a stock over a week, but is significant when assessing a transaction cost that is of approximately the same magnitude.

Sequential trade models with asymmetric information, in which market makers set quotes (e.g., Glosten and Milgrom (1985) and Easley and O'Hara (1987)) have the property that the transaction price is a better predictor of the true value of the asset than the midpoint of the bid–ask quotes. In such models, market makers choose bid and ask prices conditional on a trade occurring, and hence the quotes reflect the posterior belief of the asset's value. In our model, quotes are less informative, since they are set in previous periods by limit order traders. As we have shown, our limit order traders strategically compete and often set aggressive quotes. Nevertheless, in keeping with sequential models of intermediated markets, trade is informative.

Table XII
Difference in Ticks between Midpoint and Consensus Value,
and Transaction Price and Consensus Value, Conditional on a Market
Sell Order

For the three models considered, the table shows the average difference between the midpoint of the bid–ask spread (m_t) and the consensus value (v_t), and the average difference between the transaction price (p_t) and the consensus value (v_t). These averages are computed only for periods in which a market sell order was submitted.

Model	$m_t - v_t$	$p_t - v_t$
Base case	1.02	0.10
Zero volatility	0.62	−0.19
Immediate cancellation	0.99	0.02

In Table XII, we present the differences between the midpoint and consensus value, and transaction price and consensus value, conditional on a market sell order, across the three models examined. The difference between transaction price and consensus value is simply the negative of the true transaction costs. Even in the zero volatility model, the midpoint is off by over a half tick after conditioning on the signed order. A transaction is more likely when the midpoint is misaligned with the true asset value. Therefore, conditional on a transaction, the transaction price is a better proxy than the midpoint for the asset value.³³

In the light of Proposition 3, since the midpoint of the bid–ask spread is not a good proxy for the consensus value of the asset in periods in which a transaction occurs, we do not expect the effective spread to be a good proxy for consumer surplus.

VI. Policy Evaluation: Effect of Tick Size

A valuable feature of our model is its usefulness in evaluating policy changes. By computing the equilibrium and associated consumer surplus before and after a policy change, we can determine the overall change in consumer welfare. In this section, we specifically consider a reduction in the tick size. The theoretical and empirical literature are both mixed on the effects of a tick-size change on surplus. Seppi (1997) suggests, in the context of an intermediated market, that small traders are better off under a small tick size, while large traders are at a disadvantage. Cordella and Foucault (1999) examine competing market makers and find that transaction costs are minimized at a nonzero tick size.

Nasdaq and the NYSE, both hybrid markets, have changed their tick size in recent years. Empirical evidence on the effects of these reductions is mixed.³⁴ In

³³ Comparing mean absolute errors leads to the same conclusion. For example, in the base case the mean absolute error using m_t is 1.05, compared to 0.58 using p_t .

³⁴ Among others, Ahn et al. (1998), Bessembinder (1999), Bollen and Whaley (1998), Jones and Lipson (2001), and Ronen and Weaver (2001) examine the effect on the transaction costs incurred by different parties after the move to “teenies.” Goldstein and Kavajecz (2000) 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.




Figure 11. Relationship of ticks to β distribution. This figure depicts the distribution of private values (β) used in the tick experiment. The distribution is held constant in dollar terms. Measured in eighths, the ticks range from -2.5 to $+2.5$. These translate to ticks in sixteenths ranging from -4.5 to $+4.5$.

pure limit order markets, as well, there have been a few natural experiments, such as Toronto moving to decimals in 1996. This change is analyzed by Bacidore (1997) who finds that spreads fell but trading volume did not increase.

Has the reduction in tick size been a Pareto improvement? To address this, we compare two regimes—one with 6 ticks (and tick size $\frac{1}{8}$) and one with 11 (and tick size $\frac{1}{16}$). The mean and standard deviation of the β distribution are chosen so that the same percentage of traders in both cases have valuations more extreme than the trading crowd. In the sixteenths case, we use a mean of zero and a standard deviation of four ticks. For the eighths case, we have a mean of zero and a standard deviation of two ticks. This ensures that, in both cases, the standard deviation of β is $1/4$ of a dollar. We illustrate the β distributions in Figure 11, where each tick on the axis denotes a tick at which orders may be submitted.

With a tick size of a sixteenth, one of the ticks falls on v_t , so that traders are able to trade directly at the current consensus value of the asset. This is one of the benefits of a smaller tick: Trades can occur closer to the consensus value.

A proper comparison across these two markets also requires an adjustment to innovations in the consensus value. Since innovations are in whole numbers of ticks, we set them to one tick in the eighths case, and two ticks in the sixteenths case. Hence, innovations have the same dollar magnitude. The probability of an innovation, λ , is set to 0.08 in both cases. Finally, the minimum cancellation rate is $\underline{\delta} = 0.04$. For limit buy (sell) orders, the cancellation rate increases by 0.3 for each eighth of a dollar decrease (increase) in the consensus value, up to a maximum of 0.64. The cancellation function is the same in both cases.

We report the results of this experiment in Table XIII. For ease of comparison, all values are reported relative to the tick size of a sixteenth. For surplus, 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 surplus per available share will fall, while the mean per executed share may rise. For policy prescriptions, we should care about forgone

Table XIII
Results of a Reduction in Tick Size

The table shows the average volume, investor surplus, and effective spread in a market with a tick size of $\frac{1}{8}$ and in one with a tick size of $\frac{1}{16}$. The surplus to a buyer (seller) is $\beta - p(p - \beta)$, where β is the private value of the investor and p the transaction price. *Mkt Ord Surplus* refers to the surplus accruing to market order submitters, and *Lim Ord Surplus* the surplus accruing to limit order submitters. *Total Surplus* is the sum of these two. *Available Shares* are the total shares available to be bought or sold, and *Executed Shares* are the total shares that actually transact. Each traded share represents one executed share, but two available shares (one each for the buyer and seller). *Eff Spread* is the effective spread. For ease of comparison, surplus and spread measures are quoted in sixteenths.

	Mean Per Share			
	Tick Size = $\frac{1}{8}$		Tick Size = $\frac{1}{16}$	
	Available	Executed	Available	Executed
Volume	0.387	1.000	0.406	1.000
Mkt Ord Surplus	1.544	3.986	1.623	3.998
Lim Ord Surplus	1.166	3.010	1.143	2.816
Total Surplus	2.708	6.994	2.765	6.815
Eff Spread	—	1.384	—	1.313

trades. For the effective spread, the mean across executed shares is the relevant measure.³⁵

In our experiment, the change in tick size affects the welfare of all market participants. As measured per available share, the surplus of market order submitters increases 5.1%, and that of limit order submitters falls 2.0%. Hence, the change is not Pareto improving. Traders with high $|\beta|$, who tend to submit market orders, benefit from a reduced tick size; traders with low $|\beta|$, who tend to supply liquidity, are harmed by the change. Averaging over all traders, total surplus generated by the market increases by 2.1%. In overall terms, therefore, a reduced tick size is better.

For this particular policy experiment, measures such as volume and the effective spread both capture the direction, though not the magnitude, of the change in welfare across regimes.³⁶ The effective spread falls by 5.1% with the reduction in tick size. Volume per available share is 4.9% higher, suggesting an increased willingness to trade.³⁷

³⁵ For example, de Jong, Nijman, and Röell (1995) and Venkataraman (2001) use the effective spread to measure execution quality of orders on a pure limit order market, the Paris Bourse, with those on an intermediated market (respectively, SEAQ and the NYSE).

³⁶ We also examined the effects of a change in the gains to trade (captured by the standard deviation of the β distribution). While welfare increases with an increase in gains to trade, volume remains approximately the same and the effective spread actually increases.

³⁷ Volume per available share is proportional to the ratio of limit to market orders. In the eighths market, a volume per available share of 0.387 implies a ratio of market to limit orders of 0.613. The corresponding ratio for the sixteenths market is 0.684.

Table XIV
Depth, Quoted Spread, and Time to Execution When Tick Size Changes

The table considers a market with a tick size $\frac{1}{8}$ and one with a tick size $\frac{1}{16}$, and shows the average depth (i.e., total number of shares) on each side of the book, the average bid–ask spread (Quoted Spread) and average time between a limit order being submitted and its executing (cancelled shares are ignored in computing this average time). For ease of comparison, the spread is quoted in sixteenths in both cases.

Tick Size	Book Depth on Each Side (Shares)	Quoted Spread (Sixteenths)	Time to Execution of Limit Orders (Trader Arrivals)
1/8	2.54	3.44	4.66
1/16	2.24	3.32	3.97

Market characteristics differ, and are consistent with the empirical literature.³⁸ The sixteenths market has a smaller quoted spread, as shown in Table XIV. In addition, the average depth on each side is lower. Limit orders execute more quickly in the sixteenths market, as shown in the last column of Table XIV. The time to execution shown is conditional on an order executing (i.e., is only computed for limit orders that executed). If limit orders on the book are snapped up quickly, a high volume of trade is generated. However, this also means that there are fewer orders remaining in the book. Thus, the reduced liquidity reflects successful consummation of trade, not a lack of trading opportunities.

Consider the equilibrium strategies in the respective markets. Again, we examine the strategy of a one-share trader who arrives and finds the book empty. His optimal strategy is depicted in Figure 12, where the β axis is in sixteenths in both cases, to ensure comparability across the markets. Recall that, when the tick size is an eighth, traders cannot submit orders directly at v_t . The closest tick above v_t is at $v_t + \frac{1}{16}$, and the closest tick below v_t is at $v_t - \frac{1}{16}$. Conversely, when the tick size is a sixteenth, traders can submit orders at v_t .

As Figure 12 shows, agents with $|\beta| \leq 1.8$ take the same action whether the tick size is $\frac{1}{8}$ or $\frac{1}{16}$. Those with $|\beta| \in (1.8, 5.18)$ submit more aggressive limit orders when the tick size is a sixteenth. In particular, agents with a positive β in this range submit limit buys at v_t , rather than at one tick below v_t (which corresponds to $v_t - \frac{1}{16}$, a half tick below v_t , at the tick size of an eighth).

Since β is drawn from a normal distribution with a mean of zero and a standard deviation of four ticks (as measured in sixteenths), 45.7% of the traders have $|\beta| \in (1.8, 5.18)$. Conversely, we find that agents with $|\beta| > 5.18$ submit more conservative orders when the tick size is smaller. However, only 19.5% of traders have β in this range. Hence, on average, reducing the tick size results in more aggressive orders being submitted against the empty book.

³⁸ See, for example, Goldstein and Kavajecz (2000), Jones and Lipson (2001), and Ronen and Weaver (2001).




Figure 12. Comparison of tick sizes: Optimal order choice for a one-share trader facing the empty book. For each of the two tick sizes considered, this figure depicts the optimal action of a trader who can buy or sell one share, and enters the market when the book is empty. The optimal action varies with the trader's private value, β . Here, v_t is the consensus value at time t .

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 the tick size decreases the effective spread and improves the surplus of market order submitters, at the expense of limit order submitters. Furthermore, the change in aggregate surplus is of a smaller order of magnitude than the corresponding change in volume or the effective spread.

We interpret our result in the light of order endogeneity. Amending the tick size in a limit order market primarily perturbs the division of the spoils. If supplying liquidity becomes too expensive, agents demand liquidity and vice versa. A policy change, such as a tax that directly changes the volume of trade would have a larger effect on surplus.

VII. Conclusion

The method we introduce opens the door to a class of models that embody many of the features of existing markets. The explicit calculation of investor surplus 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 relations among transaction prices, transaction costs, trader surplus, and a commonly used proxy, the effective spread. 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 surplus because supply and demand of liquidity are endogenous.

The basic model we introduce can be augmented with intermediaries, privately informed agents, or competing exchanges. Open questions include: What are reasonable proxies for surplus (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, our method should also be of use to practitioners. In particular, Lo et al. (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 more accurate.

Appendix: Proofs

Proof of Proposition 1: Suppose execution probabilities are monotonic. A trader of type β is indifferent between submitting a sell order at price p^i and price $p^{i+1} = p^i + d$ if

$$\begin{aligned}\mu_i(p_i - \beta - \Delta_i^v) &= \mu_{i+1}^e(p_{i+1} - \beta - \Delta_{i+1}^v), \\ \text{or, } (\mu_i^e - \mu_{i+1}^e)\beta &= (\mu_i^e - \mu_{i+1}^e)(p^i - \Delta_i^v) - \mu_{i+1}^e(p^{i+1} - p^i - \Delta_{i+1}^v + \Delta_i^v) \\ \beta &= p^i - \Delta_i^v - \frac{\mu_{i+1}^e}{\mu_i^e - \mu_{i+1}^e}(d - \Delta_{i+1}^v + \Delta_i^v).\end{aligned}$$

Define $\beta_{i,i+1} = p^i - \Delta_i^v - \frac{\mu_{i+1}^e}{\mu_i^e - \mu_{i+1}^e}(d - \Delta_{i+1}^v + \Delta_i^v)$. Then, all traders with $\beta < \beta_{i,i+1}$ strictly prefer to submit at p^i rather than p^{i+1} , and all traders with $\beta > \beta_{i,i+1}$ strictly prefer to submit at p^{i+1} .

Now, consider the thresholds $\beta_{i,i+1}$ and $\beta_{i+1,i+2}$. We have $\beta_{i+1,i+2} > \beta_{i,i+1}$ if and only if

$$\begin{aligned}p^{i+1} - \Delta_{i+1}^v - \frac{\mu_{i+2}^e}{\mu_{i+1}^e - \mu_{i+2}^e}(d - \Delta_{i+2}^v + \Delta_{i+1}^v) &> p^i - \Delta_i^v - \frac{\mu_{i+1}^e}{\mu_i^e - \mu_{i+1}^e}(d - \Delta_{i+1}^v + \Delta_i^v) \\ \left(1 + \frac{\mu_{i+1}^e}{\mu_i - \mu_{i+1}^e}\right)(d - \Delta_{i+1}^v + \Delta_i^v) &> \frac{\mu_{i+2}^e}{\mu_{i+1}^e - \mu_{i+2}^e}(d - \Delta_{i+2}^v + \Delta_{i+1}^v) \\ \frac{\mu_i}{\mu_i - \mu_{i+1}^e}(d - \Delta_{i+1}^v + \Delta_i^v) &> \frac{\mu_{i+2}^e}{\mu_{i+1}^e - \mu_{i+2}^e}(d - \Delta_{i+2}^v + \Delta_{i+1}^v).\end{aligned}$$

If this condition holds for all prices p^i , the thresholds across all prices are monotone. Then, the traders who submit a sell order at price p^i must have a β in the range $[\beta_{i-1,i}, \beta_{i,i+1}]$, so that the probability of such an order is $F_\beta(\beta_{i,i+1}) - F_\beta(\beta_{i-1,i})$. Q.E.D.

Proof of Proposition 2: Consider the choice between a market sell order at price B and a limit sell order at price $B + kd$. Let μ_k^e denote the probability that an order placed at price $B + kd$ will execute, and let Δ_k^v be the associated change in common value conditional on execution. The market order will be submitted over the limit order at $B + kd$ if and only if

$$B - \beta \geq \mu_k^e (B + kd - \Delta_k^v - \beta)$$

$$\text{or, } \beta \leq B - \frac{\mu_k^e}{1 - \mu_k^e} (kd - \Delta_k^v).$$

Define $\beta_{B,B+k} = B - \frac{\mu_k^e}{1 - \mu_k^e} (kd - \Delta_k^v)$.

- (i) Now suppose, all else equal, either μ_k^e increases for some price k , or Δ_k^v increases for some price k (since we are considering sell orders, an increase in Δ_k^v constitutes an increase in picking-off risk). Then, $\beta_{B,B+k}$ decreases, so that fewer traders prefer a market sell order at B over a limit sell at $B + kd$. However, a trader submits a market order only if it is preferred to limit sell orders at $B + kd$ for all $k > 0$. Hence, the probability of a market sell order decreases weakly.
- (ii) Suppose, all else equal, B increases. Then, $\beta_{B,B+k}$ increases for all $k > 0$. The probability of a market sell order is $F_\beta(\min_{k>0} \{\beta_{B,B+k}\})$. This probability unambiguously increases when $\beta_{B,B+k}$ increases for all $k > 0$. Q.E.D.

Proof of Proposition 3:

- (i) The surplus of the market order submitter is $W_t^m = \bar{x}(\beta_t + v_t - P_t(\bar{x}))$. From equation (12), 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,

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

Thus, $W_t^m = \bar{x}(\beta_t - S_t)$ 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 surplus accruing to the limit order submitters who trade at t . This is $\bar{x}(v_t + P_t(\bar{x}) - \beta_{(t)}^l)$. Similar to part (i), we obtain $W_t^l = |\bar{x}|S_t - \bar{x}\beta_{(t)}^l$. Q.E.D.

REFERENCES

- Ahn, Hee-Joon, Kee-Hong Bae, and Kalok Chan, 2001, Limit orders, depth and volatility: Evidence from the stock exchange of Hong Kong, *Journal of Finance* 61, 767–788.
- Bacidore, Jeffrey, 1997, The impact of decimalization on market quality: An empirical investigation of the Toronto Stock Exchange, *Journal of Financial Intermediation* 6, 92–120.
- Bessembinder, Hendrik, 1999, Trade execution costs on NASDAQ and the NYSE: A post-reform comparison, *Journal of Financial and Quantitative Analysis* 34, 387–407.
- Biais, Bruno, Pierre Hillion, and Chester Spatt, 1995, An empirical analysis of the limit order book and the order flow in the Paris Bourse, *Journal of Finance* 50, 1655–1689.
- Biais, Bruno, David Martimort, and Jean-Charles Rochet, 2000, Competing mechanisms in a common value environment, *Econometrica* 68, 799–837.
- Bollen, Nicholas P. B., and Robert E. Whaley, 1998, Are “teenies” better? *Journal of Portfolio Management* 25, 10–24.
- Chakravarty, Sugato, and Craig Holden, 1995, An integrated model of market and limit orders, *Journal of Financial Intermediation* 4, 213–241.
- Chan, Louis K. C., and Josef Lakonishok, 1995, The behavior of stock prices around institutional trades, *Journal of Finance* 50, 1147–1174.
- Cohen, Kalman J., Steven F. Maier, Robert A. Schwartz, and David K. Whitcomb, 1981, Transaction costs, order placement strategy and the existence of the bid-ask spread, *Journal of Political Economy* 89, 287–305.
- Copeland, Thomas E., and Dan Galai, 1983, Information effects on the bid-ask spread, *Journal of Finance* 38, 1457–1469.
- Cordella, Tito, and Thierry Foucault, 1999, Minimum price variation, time priority and quote dynamics, *Journal of Financial Intermediation* 8, 141–173.
- de Jong, Frank, Theo Nijman, and Ailsa Röell, 1995, A comparison of the cost of trading French shares on the Paris Bourse and on SEAQ International, *European Economic Review* 39, 1277–1301.
- Demsetz, Harold, 1968, The cost of transacting, *Quarterly Journal of Economics* 82, 33–52.
- den Haan, Wouter J., and Albert Marcet, 1994, Accuracy in simulations, *Review of Economic Studies* 61, 3–17.
- Easley, David, and Maureen O'Hara, 1987, Price, trade size, and information in securities markets, *Journal of Financial Economics* 19, 60–90.
- Edwards, Amy K., and Jeffrey Harris, 2001, Stepping ahead of the book, Working paper, SEC.
- Ellul, Andrew, Craig W. Holden, Pankaj Jain, and Robert Jennings, 2003, Determinants of order choice on the New York Stock Exchange, Working Paper, Indiana University.
- Foucault, Thierry, 1999, Order flow composition and trading costs in a dynamic limit order market, *Journal of Financial Markets* 2, 99–134.
- Foucault, Thierry, Ohad Kadan, and Eugene Kandel, 2003, The limit order book as a market for liquidity, Working Paper, HEC.
- Goldstein, Michael A., and Kenneth Kavajecz, 2000, Eighths, sixteenths, and market depth: Changes in tick size and liquidity provision on the NYSE, *Journal of Financial Economics* 56, 125–149.
- Glosten, Larry, 1987, Components of the bid-ask spread and the statistical properties of transaction prices, *Journal of Finance* 42, 1293–1307.
- Glosten, Larry, 1994, Is the electronic limit order book inevitable? *Journal of Finance* 49, 1127–1161.
- Glosten, Larry, 1998, Competition, design of exchanges and welfare, Working paper, Columbia University.
- Glosten, Larry, and Paul Milgrom, 1985, Bid, ask, and transaction prices in a specialist market with heterogeneously informed traders, *Journal of Financial Economics* 13, 71–100.
- Griffiths, Mark, Brian Smith, D. Alasdair Turnbull, and Robert W. White, 2000, The costs and determinants of order aggressiveness, *Journal of Financial Economics* 56, 65–88.
- Hamao, Yasushi, and Joel Hasbrouck, 1995, Securities trading in the absence of dealers: Trades, and quotes on the Tokyo Stock Exchange, *Review of Financial Studies* 8, 849–878.

- Handa, Puneet, and Robert A. Schwartz, 1996, Limit order trading, *Journal of Finance* 51, 1835–1861.
- Harris, Lawrence, 1998, Optimal dynamic order submission strategies in some stylized trading problems, *Financial Markets, Institutions and Instruments* 7, 1–76.
- Harris, Lawrence, and Joel Hasbrouck, 1996, Market vs. limit orders: The SuperDOT evidence on order submission strategy, *Journal of Financial and Quantitative Analysis* 31, 213–231.
- Hasbrouck, Joel, 1991a, Measuring the information content of stock trades, *Journal of Finance* 46, 179–208.
- Hasbrouck, Joel, 1991b, The summary informativeness of stock trades: An econometric analysis, *Review of Financial Studies* 4, 571–595.
- Hasbrouck, Joel, 1993, Assessing the quality of a security market: A new approach to transaction-cost measurement, *Review of Financial Studies* 6, 191–212.
- Hasbrouck, Joel, 2002, Stalking the “efficient price” in market microstructure specifications: An overview, *Journal of Financial Markets* 5, 329–339.
- Hollifield, Burton, Robert Miller, and Patrik Sandås, 2004, Empirical analysis of limit order markets, *Review of Economic Studies* 71, 1027–1063.
- Hollifield, Burton, Robert Miller, Patrik Sandås, and Joshua Slive, 2004, Estimating the gains from trade in limit order markets, Working paper, Carnegie Mellon University.
- Jones, Charles M., and Marc Lipson, 2001, Sixteenths: Direct evidence on institutional execution costs, *Journal of Financial Economics* 59, 253–278.
- Keim, Donald B., and Ananth Madhavan, 1997, Transactions costs and investment style: An inter-exchange analysis of institutional equity trades, *Journal of Financial Economics* 46, 265–292.
- Keim, Donald B., and Ananth Madhavan, 1998, The cost of institutional equity trades, *Financial Analysts Journal* July/August, 50–69.
- Kumar, Praveen, and Duane Seppi, 1993, Limit orders and market orders with optimizing traders, Working paper, Carnegie Mellon University.
- Kyle, Albert, 1985, Continuous auctions and insider trading, *Econometrica* 53, 1315–1336.
- Lo, Andrew, A. Craig MacKinlay, and June Zhang, 2002, Econometric models of limit-order executions, *Journal of Financial Economics* 65, 31–71.
- O’Hara, Maureen, and G. Oldfield, 1986, The microeconomics of market making, *Journal of Financial and Quantitative Analysis* 21, 361–376.
- Pakes, Ariel, and Paul McGuire, 2001, Stochastic algorithms, symmetric Markov perfect equilibrium, and the ‘curse’ of dimensionality, *Econometrica* 69, 1261–1281.
- Parlour, Christine A., 1995, Price dynamics in limit order markets, PhD. thesis, Queen’s University at Kingston.
- Parlour, Christine A., 1998, Price dynamics in limit order markets, *Review of Financial Studies* 11, 789–816.
- Ranaldo, A., 2004, Order aggressiveness in limit order book markets, *Journal of Financial Markets* 7, 53–74.
- Rieder, U., 1979, Equilibrium plans for non-zero-sum Markov games, in Otto Moeschlin and Diethard Pallaschke, eds.: *Game Theory and Related Topics* (North Holland Publishers, Amsterdam), pp. 91–101.
- Rock, Kevin, 1990, The specialist’s order book and price anomalies, Working paper, Harvard Business School.
- Roll, Richard, 1984, A simple implicit measure of the effective bid-ask spread in an efficient market, *Journal of Finance* 39, 1073–1103.
- Ronen, Tavy, and Daniel Weaver, 2001, “Teenies” anyone? *Journal of Financial Markets* 4, 231–260.
- Rosu, Ioanid, 2004, A dynamic model of the limit order book, Working paper, MIT.
- Sandås, Patrik, 2001, Adverse selection and competitive market making: Empirical evidence from a limit order market, *Review of Financial Studies* 14, 705–734.
- Seppi, Duane J., 1997, Liquidity provision with limit orders and a strategic specialist, *Review of Financial Studies* 10, 103–150.
- Vayanos, Dimitri, 2001, Strategic trading in a dynamic noisy market, *Journal of Finance* 56, 131–171.
- Venkataraman, Kumar, 2001, Automated versus floor trading: An analysis of execution costs on the Paris and New York exchanges, *Journal of Finance* 56, 1445–1485.